

A STUDY OF NUCLEAR RECOIL
BACKGROUNDS IN DARK MATTER
DETECTORS

SHAWN S. WESTERDALE

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
PHYSICS
ADVISER: FRANK CALAPRICE

NOVEMBER 2016

© Copyright by Shawn S. Westerdale, 2016.

All rights reserved.

Abstract

Despite the great success of the Standard Model of particle physics, a preponderance of astrophysical evidence suggests that it cannot explain most of the matter in the universe.

This so-called dark matter has eluded direct detection, though many theoretical extensions to the Standard Model predict the existence of particles with a mass on the 1–1000 GeV scale that interact only via the weak nuclear force. Particles in this class are referred to as Weakly Interacting Massive Particles (WIMPs), and their high masses and low scattering cross sections make them viable dark matter candidates. The rarity of WIMP-nucleus interactions makes them challenging to detect: any background can mask the signal they produce. Background rejection is therefore a major problem in dark matter detection.

Many experiments greatly reduce their backgrounds by employing techniques to reject electron recoils. However, nuclear recoil backgrounds, which produce signals similar to what we expect from WIMPs, remain problematic.

There are two primary sources of such backgrounds: surface backgrounds and neutron recoils. Surface backgrounds result from radioactivity on the inner surfaces of the detector sending recoiling nuclei into the detector. These backgrounds can be removed with fiducial cuts, at some cost to the experiment’s exposure. In this dissertation we briefly discuss a novel technique for rejecting these events based on signals they make in the wavelength shifter coating on the inner surfaces of some detectors.

Neutron recoils result from neutrons scattering off of nuclei in the detector. These backgrounds may produce a signal identical to what we expect from WIMPs and are extensively discussed here. We additionally present a new tool for calculating (α, n) yields in various materials.

We introduce the concept of a neutron veto system designed to shield against, measure, and provide an anti-coincidence veto signal for background neutrons.

We discuss the research and development that informed the design of the DarkSide-50 boron-loaded liquid scintillator neutron veto. We describe the specific implementation of this veto system in DarkSide-50, including a description of its performance, and show that it can reject neutrons with a high enough efficiency to allow DarkSide-50 to run background-free for three years.

Acknowledgements

In some sense, this dissertation is meant to encapsulate the work I did over the last several years. However, none of this work would have been possible without the many people, organizations, and fundamental forces that got me through all of this. While an adequate acknowledgement of everyone I owe thanks to could be a whole dissertation by itself, I would like to use this space to give a special thank-you to people whose support was particularly valuable.

First and foremost, I would like to thank my advisor, Frank Calaprice. I would not have gotten where I am today without the generous guidance and support he gave me from my first day at Princeton, and I would not have been able to begin writing this document without all he has taught me. I cannot thank Frank enough for his patience and encouragement as I progressed through this journey, and for the hours he spent explaining various ideas to me.

I also owe an indescribable level of gratitude to Peter Meyers. The amount of time he spent listening to me rambling on about crazy ideas without breaking out laughing is truly an impressive feat. I wouldn't have made it through these past few years of graduate school without the support and advice that he so generously gave me, and the analysis presented in this document would be of much lower quality without all that I learned from him.

These acknowledgements would, of course, be incomplete without giving ample thanks to Cristiano Galbiati for the support he gave me through my graduate career. He encouraged me when I needed encouragement and pushed me to strive harder. Thanks to Cristiano's support, this document achieved both the breadth and depth that it has today.

In the summer of 2015, I attended the TRISEP summer school at Perimeter Institute. This is where I first met Mariangela Lisanti, who gave some of the best lectures on dark matter phenomenology during the program. Her brilliant lectures

elucidated the physics underlying my entire field, and motivated me to pursue these questions with a clear target in mind. Back at Princeton, her further discussions on the direction of the field of dark matter detection and the motivating theoretical space in which we may search for dark matter continued to deepen my understanding and appreciation of the field. I want to thank her for teaching me so much, and for graciously agreeing to serve on my defense committee. I hope the phenomenological details in my introductory chapter are up to snuff.

Throughout my time at Princeton, Helen Ju has been a constant help as I've had to figure out how to navigate the department and the ever-changing paperwork and rules related to traveling for work. I would therefore like to thank Helen for all that she has done for me, from booking travel to helping with important paperwork.

In reality, I owe thanks to the entire DarkSide collaboration for working with me these past few years and making this experiment possible. However, a brief look at the author list of any of our recent publications will make it clear that it would be impractical for me to thank each member individually for their specific contributions. However, there are certain collaborators I would like to mention in particular.

At the top of this list is Stefano Davini, my neutron veto co-conspirator. The work he did on the outer detector data acquisition system was crucial to the operation of the vetoes; but, more than that, I want to thank him for the time he spent working with me as we mapped out and coordinated the neutron veto analysis. Stefano has patiently waded through my all-too-long emails full of analysis ideas, and he has helped weed out the bad ideas, replacing them with better suggestions. Much of the neutron veto analysis presented in this document would not have been possible without Stefano's help. More importantly, I want to thank Stefano for convincing me to read *Lord of the Rings*; our journey through the neutron veto analysis was paralleled only by Frodo's journey through Mordor.

I also want to thank Luca Pagani and Laura Marini for all of the hours they poured into helping with the neutron veto analysis, the great work they did on the hardware and software systems of the outer detectors, and for being overall excellent people. I want to thank them for showing me around Rome and Genova as well as for working with me on DarkSide these past few years.

And, of course, I owe endless gratitude to Emily Shields, for introducing me to the bell jar studies we worked so tirelessly on for so many years and for the Herculean amount of work she did in the early days of prototyping and developing the neutron veto. Emily showed me the ropes with this research when I first arrived at Princeton, and I want to thank her for the great momentum with which she set me moving. I also want to thank her for that great photograph of us looking awesome in our clean room suits after installing Lumirror in the neutron veto. I also owe her many thanks for the work we did together studying neutron recoils in sodium iodide and liquid scintillator at the University of Notre Dame neutron beam.

Furthermore, I would like to acknowledge the industrious and tireless work contributed by Shawn Westerdale. This document would never have been written without all of the work done by him.

I also owe many thanks to Jingke Xu for the work we did together studying neutron recoils at the University of Notre Dame neutron beam and for his guidance and collaboration studying surface background scintillation in wavelength shifter. The experience and knowledge that he has passed down to me were crucial to the work I did in this dissertation.

I also want to thank Francis Froberg for her help surviving the chemistry lab. Her work was especially heroic when FedEx broke the scintillator cells we had meant to test at the University of Notre Dame neutron beam, and she stepped in at the last minute to prepare a fresh batch of cells to use in their place.

On this note, I want to also thank Mike Souza and Allen Nelson for helping me build the apparatuses used for the prototype studies. I especially want to thank Mike Souza for blowing the glass used to build a lot of the equipment we used, including the scintillator cells themselves. Mike graciously gave his time, even coming in on the weekends to teach me how to sparge the scintillator and help me seal the cells once they were full. It was an honor to work with a glassblower as skilled and as generous as Mike.

I would also like to thank Charles Cao, who generously lent his wisdom and expertise gained from the SCENE experiment to help guide us with our similar studies done at the University of Notre Dame. Similarly, I would like to thank the entire crew at the University of Notre Dame proton accelerator complex who helped us perform these experiments, including Weiping Wu who generously let us use much of his resources and the shifters who tirelessly worked around the clock to help us make the best use of our beam time.

I would like to thank the Academy for allowing me the honor of being here today.

Masayuki Wada is another hero in the story of this thesis whom I would like to give my thanks. The countless hours we have spent together discussing physics and analysis topics has taught me so much and greatly improved the quality of the work presented here. His mastery of ROOT has saved me many hours of struggling with my code.

I would also like to thank Xinran Li and Parth Signh for their hard work on FLUKA simulations of cosmogenic neutrons in the DarkSide-50 outer detectors, and Toni Empl for lending us his vast expertise in this field and for teaching us so much about the subject matter.

Further credit goes to Xin Xiang for his work studying the cosmogenic signals in the outer detectors and his work studying systematics of the outer detector recon-

struction with his electronics Monte Carlo. I would furthermore like to thank him for helping me maintain my sanity during the thesis-writing process.

I also want to thank Chris Stanford for his collaboration and hard work in our surface background studies discussed in this text, for the great waveform viewer tool that he has developed, for his discussions about board games, and for his heroic efforts to catch 'em all.

Thanks also go to Kenan Diab, Chaney Lin, Guangyong Koh, Katie Spaulding, and Hao Qian for putting up with me as a roommate over these past few years and for being great friends. I additionally want to thank Hao Qian for the great work he has done analyzing neutron veto data with me and for his hard work for the neutron veto group. His tireless hours spent studying the neutron veto spectrum and simulating and analyzing the calibration data have especially been crucial to much of the analyses presented here, and much of this work could not have been done without him. I also owe many additional thanks to Guangyong for all of the work I have done in collaboration with him.

I must also thank Ben Loer for writing the daqman code used to record and analyze much of the data taken in the research and development projects discussed here. I also want to thank him for passing on the burden that is argus to me.

Over the last few years, I've had the great experience of working with other members of the DarkSide collaboration here at Princeton, with whom I have learned much and formed valuable friendships. I would specifically like to thank Jason Brodsky, Chengliang Zhu, Pablo Mosteiro, Richard Saldanha, and Kirsten Randle for our time together.

While at Princeton, I've formed several valuable friendships with many of my fellow graduate students. Their camaraderie and companionship has kept me going these past few years, and I would not have made it through this program without their support and friendship. I would like to make a special thanks to some of these

friends not already mentioned, including Sara Simon, who guided me through the whole postdoc application and graduating processes by staying a few steps ahead of me and shining light back in my direction and fueled my thesis-writing process with a steady supply of puns; Max Hirschberger, who helped me spoil the entire third floor about Game of Thrones; Ed Young; Anne Gambrel; Mallika Randeria; Lauren McGough; Peace Sangtawesin; Zach Sethna; Joaquín Turiaci; Aitor Lewkowycz; Nikolay Dedushenko; and Debayan Mitra. I would also like to thank Vuk Brajušković for the steady stream of humor, companionship, and anthropomorphic animal pictures that kept me going during this process. Further acknowledgements are due to the many friends not listed here that I know I am forgetting; if you are not on this list, it is not because your friendship is any less valuable to me, but rather that I accidentally missed your name in the sleep-deprived, caffeine-fueled haze in which I wrote this.

Speaking of which, I would also like to thank Lavazza for supplying me with the requisite dosage of caffeine needed to write this document.

I also want to thank my many MIT friends whom I have kept in touch with. Instead of listing names, I'll mention all of Zephyr and the Order of the Lepton as groups. In addition to the long distance friendships we've sustained over the years with occasional visits and trips together, the technical expertise of my friends with far more technological competence than I possess has saved me during many long nights that were full of terrors and me banging my head against my desk, stuck on various problems.

In addition to the friends I've mentioned above, I want to give a special thank-you to Mai Nguyen. Without her by my side, supporting me through this process, helping me cope with the stress, and aiding me in figuring out how to juggle preparation for my upcoming postdoc with trying to graduate, I would not have made it through this process.

This entire process would have been impossible to get through without the support of my family. I want to thank my parents, Don and Ellen Westerdale; to begin with, I would not have been born without them, and then this whole thing would never have gotten done. Beyond that, their love and support through this process made it possible for me to start and get through graduate school, and continues to help me as I prepare for my postdoc. My brothers Andrew and Adam Westerdale were also there for me through this whole process, and our conversations and visits sustained me through this journey.

Last but not least, I would also like to thank you. Yes, you. I want to thank you for reading even this far into my dissertation and for giving me an audience and a reason to write this thing.

To me.

Contents

Abstract	iii
Acknowledgements	v
List of Tables	xix
List of Figures	xxvi
1 Introduction	1
1.1 Evidence for Dark Matter	1
1.1.1 Galactic Rotation Curves	2
1.1.2 Gravitational Microlensing	3
1.1.3 Cosmic Microwave Background	5
1.2 Dark Matter Candidates	9
1.2.1 Axions	9
1.2.2 Sterile Neutrinos	12
1.2.3 Asymmetric Dark Matter	14
1.2.4 Weakly Interacting Massive Particles	15
1.3 WIMP Detection Experiments	25
1.3.1 Searching Parameter Space	26
1.3.2 Current State of the Search	35
1.3.3 The Importance of Low Background	36
1.4 Outline	45

2	The DarkSide-50 Detector	48
2.1	TPC Detector Design	49
2.2	TPC and Cryostat Composition	50
2.3	LAr Scintillation	57
2.4	Photomultiplier Tubes	64
3	Radioactive Contaminants	67
3.1	Radioisotopes of Interest	67
3.1.1	Argon-39	68
3.1.2	Krypton-85	70
3.1.3	Cobalt-60	72
3.1.4	Potassium-40	74
3.1.5	Tritium	76
3.1.6	Carbon-14	78
3.1.7	Thorium-232	80
3.1.8	Uranium-238	83
3.1.9	Uranium-235	88
3.2	DarkSide-50 Counting Results	90
3.2.1	Counting Techniques	91
3.2.2	Summary of Measurements	92
4	Background Sources	96
4.1	Beta-decay Backgrounds	96
4.1.1	Forbidden Transitions	100
4.1.2	Cherenkov Backgrounds	100
4.2	Gamma-ray Backgrounds	101
4.2.1	Types of X-rays and Gamma-rays	101
4.2.2	Gamma-ray Sources	103

4.2.3	Cherenkov Signals	103
4.3	Alpha-decay and Surface Backgrounds	106
4.3.1	The Alpha-decay Process	106
4.3.2	Alpha Backgrounds	107
4.3.3	Surface Backgrounds	108
4.4	Neutron Backgrounds	119
4.4.1	Cosmogenic Neutrons	120
4.4.2	Fission Neutrons	130
4.4.3	(alpha,n) Neutrons	135
5	Neutron Veto Concept	157
5.1	Water Cherenkov Veto	157
5.1.1	Passive Shielding	158
5.1.2	The Cherenkov Effect	158
5.2	Liquid Scintillator Veto	160
5.2.1	The Scintillator	160
5.2.2	Wavelength Shifter	163
5.2.3	Boron-loading Agent	165
5.2.4	Prompt Coincidence Signals	166
5.2.5	Delayed Coincidence Signals	167
5.2.6	Muon and Cosmogenic Neutron Signals	170
6	Designing and Prototyping	172
6.1	Material Choices	172
6.1.1	WCV	172
6.1.2	LSV Scintillator	173
6.1.3	LSV Reflector	177
6.2	Full Prototype Detector	181

6.2.1	Apparatus	181
6.2.2	Modes of Operation	183
6.2.3	Small Cell Tests	184
6.2.4	Filled Vessel Tests	187
6.3	Neutron Beam Measurements	193
6.3.1	Liquid Scintillator Detector	193
6.3.2	Neutron Beam Setup	193
6.3.3	LiF Target	197
6.3.4	Electrons and Data Acquisition	199
6.3.5	Stability	202
6.3.6	Thermalization Signal	202
6.3.7	Modeling	206
7	Optical Simulations	208
7.1	Framework	208
7.2	Prototype Validation	209
7.2.1	Optical Model	210
7.2.2	Simulation Results	211
7.3	DarkSide-50 Performance Prediction	211
7.3.1	Position Dependence	214
7.3.2	Predictions for a Larger Detector	215
8	The DarkSide-50 Outer Detectors	216
8.1	Electronics and Data Acquisition System	216
8.1.1	Hardware	216
8.1.2	Zero-Suppression	217
8.1.3	Offset Calibration	218
8.2	Outer Detector DAQ and Low-Level Analysis	219

8.2.1	Pulse Reconstruction	221
8.2.2	Charge Integral Estimators	223
8.3	The Water Cherenkov Veto	226
8.3.1	Detector Design	227
8.3.2	Muon Signal	227
8.4	The Liquid Scintillator Veto	230
8.4.1	Detector Design	230
8.4.2	Run Phases	231
8.4.3	Carbon-14 Background	233
8.4.4	Measured Energy Spectra	237
8.4.5	Muon Signals	244
8.4.6	Stability	247
8.4.7	Dark Matter Search Cuts	248
9	Neutron Source Calibration	253
9.1	Calibration Insertion System	253
9.2	Americium-Beryllium Source Calibration	254
9.2.1	Neutron Spectrum	255
9.2.2	Neutron Capture Signal	257
9.3	Americium-Carbon Source Calibration	260
9.3.1	Source Design	260
9.3.2	Neutron Energy-Angle Distribution	265
9.3.3	Data Analysis and Results	268
10	Background Simulations and Neutron Vetoing Efficiency	281
10.1	Cosmogenic Neutron Backgrounds	281
10.2	(alpha,n) Neutron Backgrounds	282
10.2.1	TPC Event Reconstruction	282

10.2.2 LSV Event Reconstruction	285
10.2.3 Results	286
10.3 Fission Neutron Backgrounds	298
10.4 Gamma-ray Backgrounds	300
11 Conclusions	301
11.1 Impact of Nuclear Recoil Backgrounds	303
11.2 Additional Backgrounds	305
11.2.1 Cherenkov-Mixed Electron Recoils	305
11.2.2 Neutrino-Induced Nuclear Recoils	305
11.3 Larger Detectors	306
11.4 Final Remarks	306
A Neutron Yield Calculation and Measurement Comparisons	308
Bibliography	322

List of Tables

1.1	Possible Lorentz-invariant WIMP-nucleus interaction terms up to second order in spin and momentum transfer arising from exchange particles of spin 1 or below (<i>i.e.</i> in the non-relativistic limit) [1, 2]. \vec{S}_N and \vec{S}_χ are the spins of the nucleus and the WIMP, respectively, \vec{q} is the momentum transfer, and v^\perp is the magnitude of the component of the WIMP velocity perpendicular to the momentum transfer.	21
1.2	Nuclear form factor operators for different kinds of WIMP-nucleus interactions, based on multipole expansions of the nucleon distribution in the nucleus. Definitions for M_{JM} , Σ'_{JM} , Σ''_{JM} , Δ_{JM} , $\tilde{\Phi}'_{JM}$, and Φ''_{JM} are given in [1]. The indices J and M refer to the angular momentum modes of the multipole expansion.	22
1.3	A summary of the parameters used to draw the sensitivity curves in Figure 1.14.	42
2.1	Summary of DarkSide-50 TPC components.	55
2.2	Summary of DarkSide-50 TPC components.	56
3.1	List of x-rays produced by the electron capture decay of ^{40}K [3]. . . .	76
3.2	Counting results measured per unit mass. Measurements' sources are marked with superscripts: ^a Germanium Counting, ^b GDMS, ^c NAA, ^d ICP-MS, ^e <i>in situ</i> measurement.	93

3.3	Counting results measured per piece. Measurements' sources are marked with superscripts: ^a Germanium Counting, ^b GDMS, ^c NAA, ^d ICP-MS, ^e <i>in situ</i> measurement.	93
3.4	Counting measurements of ²¹⁰ Pb about the rest of the lower ²³⁸ U chain, as measured by the Germanium Counting technique.	94
3.5	Germanium Counting results measured per unit mass or per piece for ⁶⁰ Co and ⁴⁰ K contamination.	94
4.1	Selection rules for various β -decay transitions with nuclear spin changes of ΔS and parity changes $\Delta\pi$	100
4.2	Selection rules for γ -ray decays for which the nuclear angular momentum changes by ΔL and the parity changes by $\Delta\pi$	102
4.3	γ -ray lines in DarkSide-50. Energies are in keV. Intensities relative to the top of the decay chain and correlation probabilities (in parentheses) relative to γ -ray line intensity are in percent. Only γ -ray lines with intensities greater than 1% and energies over 100 keV are included. Isotopes marked with a * have many γ -ray, and only the three highest intensity lines are included. Correlated γ -rays are only listed if they have a probability greater than 1%. γ -ray lines that are part of a cascade with many lines are labeled with \mathcal{C} , and only correlations with listed lines are explicitly included.	104

4.4	γ -ray lines in DarkSide-50. Energies are in keV. Intensities relative to the top of the decay chain, and correlation probabilities (in parentheses) relative to γ -ray line intensity are in percent. Only γ -ray lines with intensities greater than 1% and energies over 100 keV are included. Isotopes marked with a * have many γ -ray, and only the three highest intensity lines are included. Correlated γ -rays are only listed if they have a probability greater than 1%. γ -ray lines that are part of a cascade with many lines are labeled with \mathcal{C} , and only correlations with listed lines are explicitly included.	105
4.5	Uranium and thorium contamination measurements in the DarkSide-50 underground liquid argon as measured by C. Stanford.	108
4.6	Muons	122
4.7	Cosmogenic β -delayed neutron emitters expected in DarkSide-50. . .	127
4.8	Rough estimates of the rate of delayed neutron precursors generated by muons within 30 cm of the TPC cryostat	129
4.9	Summary of the fission neutrons produced in the ^{232}Th , ^{238}U , and ^{235}U decay chains, as simulated by SOURCES4A [4]. $\langle\nu\rangle$ is the mean number of neutrons produced per fission and $\langle T\rangle$ is the mean neutron kinetic energy.	132
4.10	(α, n) yield calculations determined by the calculator presented in this work, the presented data compilations, and calculations performed using SOURCES4A. Entries marked with a dash were not possible to calculate due to limitations in the available data.	149
4.11	(α, n) yield calculations determined by the calculator presented in this work, the presented data compilations, and calculations performed using SOURCES4A. Entries marked with a dash were not possible to calculate due to limitations in the available data.	150

4.12	Isotopes missing from the SOURCES4A and data neutron yield calculations presented in Tables 4.10 and 4.11. Where only a chemical symbol is given, all isotopes of that element are missing from the data. For borosilicate glass, the contribution of lithium to the total yield was only available at low energy in SOURCES4A, and therefore was included for ^{210}Pb and the lower ^{238}U chain.	151
4.13	Total (α, n) neutron production rates per year as predicted by the calculator presented in this work, published data, and SOURCES4A calculations. Estimates are made for each component of the DarkSide-50 TPC and are arranged into 10 groups. (A) Cryostats, (B) Cryosat Flanges, (C) Extraction Grid, (D) Teflon Structures, (E) TPC Windows, (F) Liquid Argon, (G) Electric Field Cage, (H) PMTs, (I) PMTHardware, (J) Multilayer Insulation.	152
5.1	Dominant inelastic scatter signals that may be seen in the LSV in prompt coincidence with neutron scatters in the TPC. The excited nucleus, its ground state angular momentum and parity J_{π}^0 , their excited state energy levels E' and their angular momenta and parity J_{π}' , the excited states' lifetime $T_{1/2}$, the cross section for a 3 MeV neutron incident on each nucleus exciting the nucleus to each given state σ , the final state energy levels that each excited state can decay to E^f , the γ -ray energy released by this transition E_{γ} , and the branching ratio for the excited state to decay to each final state BR. Energy level structure information from [3, 5, 6], and cross sections from [7].	168

6.1	Summary of the Tyvek and Lumirror reflectances shown in Figs. 6.5 and 6.6. R_{250} and R_{500} are the reflectances measured at 250 nm and 500 nm, respectively. $\lambda_{50\%}$ and $\lambda_{90\%}$ are the wavelengths at which each reflector had 50% or 90% reflectance, respectively. Notably, the six layers of Tyvek 1070D was above 50% for the entire range of wavelengths measured while dry, and the second time it passes the 90% reflectance mark is reported. While it is dry, this sample is below 50% reflectance for the entire range of wavelengths measured.	180
6.2	Radionuclide measurements made of a 1.7 kg sample of Lumirror E6SR (188 μm thick). Measurements were made by M. Laubenstein using the Germanium Counting technique. Nuclides that are part of the same decay chain are grouped together.	181
6.3	Measured light yields of the bell jar in the scintillator cell configuration with a ^{54}Mn γ -ray source and Tyvek or Lumirror as the reflector, in PE/keV. For all measurements the bell jar was filled with air and the scintillator cell was placed opposite the PMT.	187
6.4	Beamline angles and distance between the center of the liquid scintillator detector and the coincidence detectors, and the corresponding $p(n, n')p$ scattering expected recoil energy for a 690 keV neutron. . . .	196
6.5	Summary of expected, simulated, and observed recoil energies in the liquid scintillator detector when there is a coincidence with each coincidence detector.	204
6.6	Summary of fits of scintillation quenching Models I and II explored here.	206
8.1	Summary of the zero-suppression parameters used for DarkSide-50 data acquisition.	219
8.2	Water contamination levels before and after purification with the Borexino water purification plant, as reported by [8]	228

8.3	Summary of Phase-I and Phase-II of LSV operation, including start and end dates of both phases, the type of argon used in the TPC for each phase, the scintillator composition, the observed ^{14}C decay rate, the thermal neutron capture time, the calculated ratio of thermal neutrons capturing on ^1H to those capturing on ^{10}B , and the acquisition window used relative to when the outer detector DAQ received the trigger signal.	232
8.4	Accelerator mass spectrometry measurements comparing the $^{14}\text{C}/^{13}\text{C}$ fraction for each measured sample to the modern fraction.	236
8.5	Dark matter selection cuts made based on the WCV and LSV signals during Phase-I and Phase-II, along with their acceptances. These cuts are explained in more detail in the text.	249
9.1	Composition of the $^{241}\text{Am}^{13}\text{C}$ source graphite pellet.	262
9.2	The mean and RMS neutron energy and total neutron production rate of the $^{241}\text{Am}^{13}\text{C}$ source with a $100\mu\text{Ci}$ ^{241}Am source and three different thicknesses of gold foil between the ^{241}Am source and graphite pellet, as simulated using GEANT4 and TALYS.	267
9.3	Best fit parameters for Model II and their χ^2/NDF	277
10.1	Recombination probability model parameters used in the model tuned by D. Franco and P. Agnes on electron recoil data.	283

10.2 Neutron background predictions from the dominant (α , n) sources in DarkSide-50 for one year. Columns show the total number of neutrons produced in one year, the number expected to pass all TPC cuts, the number expected to pass the TPC cuts and each of the three LSV cuts, and all TPC and all LSV cuts. Neutron sources included here are the PMT stems, PMT hardware, cryostat walls, multilayer insulation (MLI), the inner cryostat flange, and the outer cryostat flange. Uncertainties are discussed in the text.	296
--	-----

List of Figures

1.1	Galactic rotation curves of two different galaxies measured [9]. A three parameter fit was performed on the data: the dashed lines represent the contributions from visible components, the dotted lines from gas, and the dash-dotted line from the dark halo.	3
1.2	A cartoon showing the process of gravitational lensing as predicted by general relativity. θ is the observed angle between the image and the lens, β is the observed angle if the lens were massless (i.e. with no lensing), α is the reduced lensing angle, and $\hat{\alpha}$ is the deflection angle., d_{LS} is the distance between the lens and the source, d_{OL} is the distance between the observer and the lens, and d_{OS} is the distance between the observer and the source.	4
1.3	X-ray photograph of the Bullet Cluster taken by the Chandra X-ray Observatory [10].	5
1.4	Temperature variations over the entire sky as measured by Planck. The bright horizontal strip across the center is due to distortions from the Doppler shift caused by Earth's motion [11].	6
1.5	The power spectrum of the CMB as measured by Planck. The x-axis is logarithmic for $l < 50$ and linear above, and the y-axis scale is in $\frac{l(l+1)C_l}{2\pi}$ [11].	8

1.6	A Feynman diagram showing the Primakoff conversion process of an axion into a photon in the presence of an external magnetic field, as described by Equation 1.9	12
1.7	WIMP speed distribution based on Equation 1.47 for $v_0 = 230$ km/s, $v_{esc} = 600$ km/s, and (<i>green</i>) $v_E = 244$ km/s or (<i>blue</i>) $v_E = 0$ km/s. . .	28
1.8	Nuclear form factors computed using Equation 1.62 for (<i>blue</i>) argon, (<i>green</i>) xenon, (<i>purple</i>) germanium, and (<i>red</i>) helium.	31
1.9	Recoil energy spectrum for (<i>blue</i>) 10 GeV, (<i>green</i>) 100 GeV, (<i>purple</i>) 500 GeV, and (<i>orange</i>) 1000 GeV WIMPs on an (<i>top</i>) argon or (<i>bottom</i>) xenon target, assuming a cross section of 1 pb. Solid lines include the effects of the form factor, while dashed lines set $F^2(q) \equiv 1$	32
1.10	A Feynman diagram depicting a WIMP scattering off of a nucleus via some exchange particle. g_χ and g_N are the coupling constants for the WIMP and the nucleus with the exchange particle, respectively. . . .	33
1.11	The most recent exclusion curves resulting from several different direct detection experiments at the time of writing this document [12]. (<i>Solid purple</i>) WARP, (<i>dotted red</i>) DarkSide-50 (Atmospheric Argon), (<i>dashed red</i>) DarkSide-50 (Underground Argon), and (<i>solid red</i>) DarkSide-50 (combined) used liquid argon targets. (<i>Dash-dotted black</i>) ZEPPLIN-III, (<i>dotted black</i>) PandaX-I, (<i>dashed black</i>) XENON100, and (<i>solid black</i>) LUX use liquid xenon targets. (<i>Solid green</i>) CDMS used a germanium crystal target. (<i>Solid cyan</i>) PICO used a liquid C_3F_8 target.	35

1.12	The number of WIMP-like events needed to claim a WIMP observation to (<i>from top to bottom</i>) 1, 3, and 5σ based on the predicted background rate of the experiment. Solid lines show the number of WIMP-like events needed, including backgrounds, while dashed lines show the number of WIMP-scatters needed (after subtracting the expected background). Blue, green, and purple curves were made assuming uncertainty on the background model of 100%, 10%, and 1%, respectively.	39
1.13	The number of WIMP-like events needed to claim a WIMP observation to (<i>from top to bottom</i>) 1, 3, and 5σ based on the predicted background rate of the experiment, on a Log-Log scale. Solid lines show the number of WIMP-like events needed, including backgrounds, while dashed lines show the number of WIMP-scatters needed (after subtracting the expected background). Blue, green, and purple curves were made assuming uncertainty on the background model of 100%, 10%, and 1%, respectively.	40
1.14	(<i>Top</i>) Sensitivity curves drawn for an idealized liquid argon dark matter detector similar to DarkSide-50. The area above each curve is the area of parameter space where this detector would be able to detect a WIMP; each curve represents a different threshold number of WIMP-nucleus scattering events the detector would have to observe in order to claim a detection with 50% probability. The background-free case is shown in black, where 1 event is required. Beyond that we see the curves for (<i>blue</i>) 2, (<i>green</i>) 3, (<i>purple</i>) 4, (<i>orange</i>) 5, (<i>teal</i>) 10, (<i>red</i>) 15, and (<i>pink</i>) 20. (<i>Bottom</i>) The same plots, drawn for an identical experiment using liquid xenon instead of argon.	41

1.15	A family of sensitivity curves based on a toy model of the LUX detector with 10065 kg days of exposure and a nuclear recoil energy range of 3–25 keV.	43
2.1	A drawing of the DarkSide-50 detectors.	48
2.2	A cutaway view of the DarkSide-50 TPC.	49
2.3	(<i>Left</i>) The DarkSide-50 outer cryostat as it was lowered into the LSV. (<i>Right</i>) The inner cryostat, removed from the outer cryostat, which contains the TPC. Photographs by Y. Suvarov.	51
2.4	(<i>Left</i>) The inner cryostat top dome covered in embossed multilayer insulation with the outer cryostat dome above it. (<i>Right</i>) The inside of the outer cryostat with netted multilayer insulation along the walls and a Viton o-ring in the flange. Photographs by Y. Suvarov.	52
2.5	(<i>Left</i>) The cryostat opened up to reveal the TPC. (<i>Right</i>) Top-down view of the TPC. Photographs by Y. Suvarov.	52
2.6	(<i>Left</i>) The grid that maintains the extraction field in the TPC. (<i>Right</i>) TPC field cage rings and resistors. Photographs by Y. Suvarov.	53
2.7	The top array of PMTs in their Teflon support mount, with circuitry and cold amplifiers visible. Photographs by P. Meyers.	54
2.8	Diagram showing the process of liquid argon scintillation after a particle X has scatters on an argon nucleus. f is the fraction of energy that the recoiling argon nucleus loses to electrons; $(1 - f)$ is the fraction lost to other argon nuclear recoils. N_{ex} and N_{ion} are the number of excitons and the number of ions produced, respectively, and $\alpha = N_{\text{ex}}/N_{\text{ion}}$. r is the recombination rate for argon ions capturing an ionized electron. Electrons that fail to recombine escape, and can either be collected as an S2 signal or be absorbed by another atom (such as an impurity), or they may be recaptured by another ion a long time later.	59

2.9	Drawing of Hamamatsu R11065 PMT detecting a photon.	64
2.10	(<i>Black</i>) The SPE charge distribution readout from a Hamamatsu R11065 PMT. (<i>Red</i>) A Gaussian distribution plus an exponential is fit to this curve above the pedestal, taken from [13].	65
2.11	The Hamamatsu R11065 quantum efficiency as a function of incident photon wavelength, as reported by Hamamatsu [14].	66
3.1	Decay diagram of ^{39}Ar . Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q_{β^-} -value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^- decays, the given energy is the endpoint energy of the β^- [15].	70
3.2	Decay diagram of ^{85}Kr . Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q_{β^-} -value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^- decays, the given energy is the endpoint energy of the β^- [16].	71

3.3	Decay diagram of ^{60}Co . Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Blue arrows and text refer to internal de-excitation processes within a nucleus, with the relative intensity, γ -ray energy, and lifetime of each transition. Different shades of blue were used to help identify which text describes which arrow. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q_{β^-} -value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^- decays, the given energy is the endpoint energy of the β^- [17].	73
3.4	Decay diagram of ^{40}K . Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission (“EC” refers to electron capture), its energy, relative intensity, and half-life. Red lines and text show Q -value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^\pm decays, the given energy is the endpoint energy of the β^\pm [3]. *The decay of ^{40}Ca to ^{40}Ar has not been observed, but is believed to be possible due to the positive Q -value of this reaction and the fact that it does not violate any conserved quantities.	75

3.5	Decay diagram of ${}^3\text{H}$. Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q-value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β decays, the given energy is the endpoint energy of the β [18].	77
3.6	The ${}^3\text{H}$ β^- spectrum (<i>blue</i>) with the Fermi function correction for the nucleus's Coulomb potential and (<i>green</i>) without. The correction decreases the mean β^- energy by 0.28 keV.	77
3.7	Decay diagram of ${}^{14}\text{C}$ and ${}^{14}\text{O}$ to ${}^{14}\text{N}$. Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Blue arrows and text refer to internal de-excitation processes within a nucleus, with the relative intensity, γ -ray energy, and lifetime of each transition. Different shades of blue were used to help identify which text describes which arrow. Red lines and text show Q-value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^\pm decays, the given energy is the endpoint energy of the β^\pm [19]. The three $T = 1$ multiplet states are labeled.	79
3.8	The ${}^{14}\text{C}$ β spectrum (<i>blue</i>) with the shape factor correction predicted by [20], and (<i>green</i>) without this correction.	80

3.9	The ^{232}Th decay chain. Blue circles are each isotope in the chain above the isotope's half-life. Green arrows and text represent decay modes and their associated Q-values in keV. Red numbers show the branching ratio whenever the chain splits. Purple wavy lines and text represent γ -rays produced by the decay they are associated with, written as the γ -ray energy in keV followed by the intensity of that γ -ray in parentheses. The intensity of a γ -ray is the probability of that γ -ray appearing in a decay. Only γ -rays with intensities greater than 1% and energies greater than 100 keV are included. Only branches with a branching ratio above 1×10^{-6} are included [21].	82
3.10	Fission neutron production rate per mBq of ^{232}Th , as simulated by SOURCES4C [4].	83
3.11	The ^{238}U decay chain. Blue circles are each isotope in the chain above the isotope's half-life. Green arrows and text represent decay modes and their associated Q-values in keV. Red numbers show the branching ratio whenever the chain splits. Purple wavy lines and text represent γ -rays produced by the decay they are associated with, written as the γ -ray energy in keV followed by intensity of that γ -ray in parentheses. Only γ -rays with intensities greater than 1% and energies greater than 100 keV are included. Only branches with a branching ratio above 1×10^{-6} are included [22]. The black dashed line above ^{226}Ra shows where secular equilibrium is likely broken.	85
3.12	Fission neutron production rate per mBq of ^{238}U , as simulated by SOURCES4C [4].	86

3.13	Decay diagram of ^{210}Tl via the β^-n channel to ^{209}Pb . The much more likely branches that ^{210}Tl can decay to via normal β^- decay to ^{210}Pb are not shown. Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q-value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^\pm decays, the given energy is the endpoint kinetic energy of the β^\pm . Spin and parity of the excited states shown for ^{210}Pb are not shown because they were not listed in the ENSDF dataset. Similarly, the relative intensities of both of these branches were not specified, but their total intensity is 0.007% [23, 24].	87
3.14	The ^{235}U decay chain. Blue circles are each isotope in the chain above the isotope's half-life. Green arrows and text represent decay modes and their associated Q-values in keV. Red numbers show the branching ratio whenever the chain splits. Purple wavy lines and text represent γ -rays produced by the decay they are associated with, written as the γ -ray energy in keV followed by intensity of that γ -ray in parentheses. Only γ -rays with intensities greater than 1% and energies greater than 100 keV are included. Only branches with a branching ratio above 1×10^{-6} are included [25].	89
3.15	Fission neutron production rate per mBq of ^{235}U , as simulated by SOURCES4C [4].	90

4.1	Fermi functions showing the Coulomb effects on β spectra as a function of β kinetic energy for β^- emission in a (<i>blue</i>) ^3H -like nucleus and a (<i>green</i>) ^{14}C -like nucleus, and for β^+ emission in a (<i>purple</i>) ^3H -like nucleus and a (<i>orange</i>) ^{14}C -like nucleus. Dashed lines show the full Fermi function given in Equation 4.4, and solid lines show the approximation given in Equation 4.5.	99
4.2	A drawing of the surface background of interest to DarkSide-50, in which a ^{210}Po nucleus decays and sends a recoiling ^{206}Pb nucleus into the liquid argon and an α into the wavelength shifter (TPB for DarkSide-50), which may scintillate as well.	109
4.3	(<i>Left</i>) A schematic of the detector used to measure surface background scintillation (courtesy of C. Stanford). (<i>Right</i>) A photograph of the detector.	110
4.4	The setup used to deposit radon daughters on samples.	112
4.5	Measured visible energy spectrum of ^{214}Po α -decays in liquid argon. The high energy peak is from ^{214}Po α scintillation, while the lower energy peak is from ^{210}Pb scintillation. The inset plot shows visible energy vs. f_{90} for these events. At lower energies, we see a tail leading down to lower f_{90} , which we conclude is from decays of ^{214}Po embedded in the substrate. We therefore require that ^{210}Pb events be in the red box shown in this inset plot to appear in the mean energy spectrum.	113
4.6	Molecular diagrams of (<i>left</i>) tetraphenyl butadiene (TPB) and (<i>right</i>) p-terphenyl (pTP).	114
4.7	An illustration of the scintillation process of organic scintillators. . . .	115
4.8	(<i>Left</i>) Emission spectrum of TPB measured at different temperatures, taken from [26]. Vertical offsets between curves have been inserted for clarity. (<i>Right</i>) Emission spectrum of pTP, from [27].	117

4.9	Acrylic slide with pTP wavelength shifter evaporated onto it, glowing in a UV lamp.	118
4.10	α -decay scintillation spectra for ^{222}Rn daughters on LAr and either TPB or pTP. The inset plot shows the f_{90} distribution for the events where the α scintillated in the wavelength shifter and the recoiling ^{210}Pb nucleus scintillated in the LAr	118
4.11	Cosmogenic muon energy distribution calculated from Equation 4.7 at (<i>blue</i>) LNGS, (<i>green</i>) SNOLAB, and (<i>purple</i>) SURF, normalized to the laboratory's muon flux, summarized in Table 4.6.	121
4.12	(<i>Left</i>) The muon azimuthal angular distribution and (<i>right</i>) muon zenith angular distribution, as measured by (<i>blue</i>) Borexino and (<i>green</i>) MACRO, and (<i>red</i>) simulated by FLUKA with a limited zenith angle to match MACRO, and (<i>black</i>) with a full zenith angular distribution to match Borexino. Plot from [28].	122
4.13	(<i>Left</i>) Particle production rates of 280 GeV muons as a function of how far the muon has traveled through rock, normalized to a maximum value of 1 for each particle species. (<i>Right</i>) The flux of particles produced by muons in Hall C of LNGS per muon as a function of the particle's kinetic energy. Both plots were generated by a FLUKA simulation and taken directly from [28].	125
4.14	The cosmogenic neutron energy spectrum per meter of muon track length produced by a 285 GeV muon. From [29].	126
4.15	The distribution of the number of neutrons produced in ^{238}U spontaneous fission as reported by [30]. The data is well fit by a Gaussian with mean 2.05 ± 0.04 and standard deviation 1.04 ± 0.03	132

4.16	(<i>Top</i>) The total energy carried away by γ -rays as a function of neutron multiplicity for the spontaneous fission of ^{238}U , as described by Equation 4.13. (<i>Bottom</i>) The mean number of γ -rays produced in the fission fragment decay chains of ^{238}U as a function of the neutron multiplicity, as described by Equations 4.14 and 4.13.	134
4.17	A scatter plot comparing the (α, n) yields calculated by the tool presented in this present work versus benchmark calculations in [31, 32, 33]. The straight black line is $x = y$, showing where this calculator agrees with a benchmark perfectly. Solid dots represent measurements presented in Heaton, squares measurements in Stelson and McGowan, triangles pointing up measurements in Bair and Gomez del Campo, hollow circles in Roughton et al., and triangles pointing down calculations done by SOURCES4A. Different colors correspond to measurements done on different target isotopes; progressions of points the same shape and color represent yields varying over energy, which increase monotonically. A breakdown of the yield calculations for each target isotope versus energy as measured by each benchmark and this calculator is given in Appendix A.	142
4.18	The distribution of the deviation of neutron yields calculated in this work compared to the various benchmark measurements and calculations for a range of isotopes with Z ranging from 3 to 49 and incident α energies ranging from 5 to 11 MeV. (<i>Top</i>) A scatter plot of the deviation versus the neutron yield. (<i>Bottom</i>) A histogram of all of the deviations, with a mean of 29.6% and an RMS of 31.9%.	143
4.19	The fractional contribution to the total stopping power for each element in (<i>top</i>) borosilicate glass and (<i>bottom</i>) stainless steel.	146
4.20	(α, n) neutron energy spectrum from the cryostat body.	154

4.21	(α , n) neutron energy spectrum from the cryostat flanges.	154
4.22	(α , n) neutron energy spectrum from the multilayer insulation.	155
4.23	(α , n) neutron energy spectrum from the TPC PMTs.	155
4.24	(α , n) neutron energy spectrum from the TPC PMT hardware.	156
5.1	(<i>Left</i>) Total interaction cross section for neutrons on (<i>blue</i>) ^1H and (<i>green</i>) ^{16}O . Data from [7]. (<i>Right</i>) The neutron mean free path in water.	158
5.2	Illustration of the Cherenkov effect from a charged particle moving through some medium with a velocity greater than the speed of light in that medium.	159
5.3	Molecular diagrams of (<i>top left</i>) PC, (<i>top right</i>) TMB, and (<i>bottom</i>) PPO.	161
5.4	The (<i>black</i>) attenuation length, (<i>blue</i>) absorption spectrum, and (<i>red</i>) emission spectrum of PC, as measured by [34].	162
5.5	(<i>Top</i>) The emission spectrum of PPO, as measured by [35]. (<i>Bottom</i>) The absorbance of (<i>black, dashed</i>) PC+ 1.5 g/L of PPO, (<i>red, solid</i>) PC+ 2.5 g/L of PPO, and (<i>blue, dotted</i>) an equal mixture of PC and TMB with 2.5 g/L of PPO, measured by A. Ianni and Y. Suvorov. . .	164
5.6	Fast neutron capture reaction cross sections on ^{11}B for (<i>blue</i>) the $^{11}\text{B}(\text{n},\alpha)^8\text{Li}$ reaction, (<i>green</i>) the $^{11}\text{B}(\text{n},\text{t})^9\text{Be}$ reaction, and (<i>purple</i>) the $^{11}\text{B}(\text{n},\text{p})^{11}\text{Be}$ reaction.	166
6.1	The relative light yield of a PPO+linear alkylbenzene solution, for various PPO concentrations taken inside a 5 cm diameter by 5 cm tall cell, from [36].	174

6.2	The (<i>solid</i>) absorption and (<i>dashed</i>) emission spectra for (<i>black</i>) POPOP dissolved in cyclohexane [35], and (<i>red</i>) bis-MSB dissolved in PC [34].	175
6.3	The relative light yield of a linear alkylbenzene solution with 3 g/L PPO in a 5 cm diameter by 5 cm tall cell with varying concentrations of either (<i>left</i>) POPOP or (<i>right</i>) bis-MSB dissolved in it, from [36]. .	175
6.4	The degradation of the attenuation length of PC over time when exposed to a surface-to-volume ratio of 20 m^{-1} of electropolished stainless steel compared to PC left unexposed. In order to speed up the interactions, the sample in stainless steel was heated by 50°C for the first 300 hours and then by 100°C afterwards, from [13].	176
6.5	Reflectance measurements of Tyvek reflectors. The black curve shows the reflectance of Tyvek 4077D with one layer, while blue shows the same with 6 layers. Additional curves compare two different types of Tyvek after soaking in a PC+TMB cocktail for various lengths of time: (<i>red</i>) 4077D after 3 hours, (<i>orange</i>) 4077D after 3 weeks, (<i>green</i>) 1070D before soaking, and (<i>cyan</i>) 1070D after 3 weeks. (<i>Not shown</i>) Tyvek 1082D exhibited very similar behavior to 1070D.	178
6.6	Reflectance measurements of Lumirror 188 E6SR, after soaking in a PC+TMB cocktail for varying lengths of time: (<i>black</i>) before soaking, (<i>blue</i>) after 3 days, (<i>orange</i>) after 3 weeks, (<i>green</i>) after 4 months, and (<i>cyan</i>) after 10 months. (<i>Red</i>) The improvement of a dry sample while using 4 layers. The bump around 240 nm is due to the specular component of the reflectivity.	179

6.7	A small sample of Lumirror 188 E6SR that has been soaked in scintillator for nine months. Some degradation of the reflectance at the edges has occurred due to the creeping of scintillator through the sides. The reflectance of the central part of this sample remains unaffected, however (see Figure 6.6).	179
6.8	Cross section of bell-jar setup for each experiment. (<i>Left</i>) The cell setup shows a cell filled with scintillator in the otherwise dry bell jar. (<i>Right</i>) The filling setup shows the bell jar, which has been modified to be filled with scintillator.	182
6.9	Plot of a typical quantum efficiency distribution for an R11065 PMT, provided by Hamamatsu [14], shown alongside the reflectance spectrum for Lumirror and the emission spectra of PPO [35], POPOP [35], and bis-MSB [37], all scaled to arbitrary units.	182
6.10	The data acquisition electronics setup for the bell jar measurements.	183
6.11	A picture of the distillation column used to purify PC and TMB for prototype tests.	184
6.12	(<i>Left</i>) A photograph of the setup used to sparge and seal the scintillator cells. (<i>Right</i>) A diagram of this setup, labeling each of the parts.	185
6.13	A picture of three scintillator cells after filling and sealing them.	186

6.14	A sample ^{54}Mn spectrum measured in the bell jar, background subtracted and scaled to the number of photoelectrons by the SPE peak. (<i>Red</i>) A smoothed simulated spectrum has been fit to the data around the Compton edge; <i>blue</i> the simulated spectrum before smoothing. The large narrow spike around 390 PE is the 835 keV full energy peak, while the broader spike just below 300 PE is the Compton edge. The part of the spectrum beyond the Compton edge is due to multiple scattering. The fit was performed specifically around the Compton edge in order to avoid effects from ionization quenching.	189
6.15	The light yield (in photoelectrons per keV) measured in the bell jar over 42 days, neglecting the exponential tail in the single photoelectron distribution, decreasing at a rate of $0.52 \pm 0.02\%$ per week.	190
6.16	An example ADC waveform of a neutron capturing on ^{10}B in the bell jar. The thermalization signal can be seen around $0 \mu\text{s}$, and the neutron capture signal is seen around $1.5 \mu\text{s}$. The blue curve on the waveform shows the total integral of the curve from the start of the acquisition window to a given time.	190
6.17	The background-subtracted prompt energy spectrum from the AmBe source. The peak at high energy is the Compton edge of the 4.4 MeV γ -ray produced by the $^9\text{Be}(\alpha, n)^{12}\text{C}$ reaction with a branching ratio of $\sim 64\%$	191
6.18	The background-subtracted delayed energy spectrum from the AmBe source. The peak at $\sim 30 \text{ keV}_{ee}$ represents the energy deposit of the ^7Li and α alone, and the broad peak around 400 keV_{ee} represents the Compton spectrum of the γ -ray plus the ^7Li and the α	192

6.19	The canister holding the scintillator cell for measurements at the University of Notre Dame neutron beam. (<i>Cyan</i>) The scintillator cell, (<i>yellow</i>) the PMT, (<i>gray</i>) the canister holding everything in place, (<i>red</i>) high voltage and signal cables, (<i>white</i>) Teflon support cup, (<i>purple</i>) support structure and springs holding everything in place.	194
6.20	(<i>Top</i>) A picture of the detectors arranged around the neutron beam. (<i>Bottom</i>) A diagram of the experimental setup.	195
6.21	Simulation of 2.44 MeV protons propagating in the LiF target. (<i>Left</i>) Proton energy versus depth. (<i>Right</i>) Proton angle relative to the beam-line versus depth.	197
6.22	Energy-angle relationship for outgoing neutrons in the ${}^7\text{Li}(p, n){}^7\text{Be}$ reaction with a 2.44 MeV proton.	198
6.23	${}^7\text{Li}(p, n){}^7\text{Be}$ cross section data from Burke et al. [38]. In both of these plots, the color represents the cross section in mb/sr. (<i>Left</i>) Neutron angle with respect to the proton versus the proton energy versus reaction cross section. (<i>Right</i>) Neutron kinetic energy versus proton energy versus cross section, made by transforming the left plot as described in the text.	199
6.24	Neutron energy and angular distribution from the ${}^7\text{Li}(p, n){}^7\text{Be}$ reaction of a 2.44 MeV proton on a 0.52 mg/cm ² LiF target.	200
6.25	Schematic diagram of the data acquisition system used for the neutron beam measurements.	201
6.26	Spectra in the liquid scintillator detector for events with coincidences in coincidence detectors 1–6.	203

6.27	An example proton recoil spectrum with a simulated spectrum fit to it. The example shown here is for coincidence detector 2. (<i>Black</i>) the measured spectrum. (<i>Red</i>) The fit simulated spectrum. The departure of the fit spectrum from the measured spectrum is due to carbon recoils, which are quenched more heavily than was accounted for in this simulation; the simulated spectrum was fit to the measured spectrum above where we see the carbon recoil contribution to the spectrum.	205
6.28	(<i>Black dots</i>) Quenching factor measurements in the liquid scintillator detector. (<i>Red triangles</i>) Measurements reported by Hong et al. [39] for a different organic liquid scintillator for comparison. (<i>Blue and green curves</i>) Two different quenching models (see text) fit to the measurements presented here.	206
7.1	A labeled drawing of the DarkSide-50 LSV as used for the GEANT4 simulations.	212
7.2	A schematic drawing of the Hamamatsu R5912 LRI PMTs used in the DarkSide-50 LSV, from [40].	213
7.3	Hamamatsu R5912 LRI PMT quantum efficiency curve as reported by Hamamatsu in [41].	213
7.4	The light yield versus position map generated by H. Qian. The x-axis is the square of the radius, the y-axis is the z-coordinate (parallel to the rotational axis of symmetry of the cryostat), and the color represents the light yield, normalized to give ^{60}Co scintillation events a light yield of 1.	214
7.5	A drawing of the LSV with a 3.8 tonne LArTPC, used for exploring how the optics of the LSV scale with a larger detector.	215

8.1	Illustration of the zero-suppression algorithm applied to the outer detector data. Made by L. Pagani.	218
8.2	A diagram showing the analysis flow of the DarkArtOD data reconstruction software. Red nodes are the raw data-level, green nodes are pulse-level, purple nodes are sum waveform-level, blue nodes are charge estimator-level, and yellow nodes are output-level.	220
8.3	An example of a waveform in the LSV with its corresponding cluster, illustrating the Top-Down Cluster Finder algorithm. The variable in the horizontal axis is the time in ns with respect to the trigger.	225
8.4	The inside of the DarkSide-50 WCV. The inner surfaces are all covered in Tyvek, as is the LSV, seen in the middle, standing on four stilts. Taken by Y. Suvarov.	227
8.5	Total charge collected in the WCV for each event in the ~ 70 day underground argon campaign [42]. (<i>Blue</i>) All events, including those with cosmogenic signals and those without. (<i>Red</i>) Events with coincidental high energy signals in the LSV, predominantly due to cosmogenic signals in both detectors.	228
8.6	An example waveform of a cosmogenic signal in the WCV. The inset plot is zoomed in around the time that the muon passed through the detector.	229
8.7	Two photographs of the inside of the DarkSide-50 LSV, taken by Y. Suvarov. (<i>Left</i>) The TPC cryostat suspended inside the LSV, held in place by connective pipes. (<i>Right</i>) The Lumirror reflector and PMTs mounted on the walls of the detector.	230
8.8	Block diagram of the fluid handling systems for the DarkSide-50 LSV, from [40].	232

8.9	(<i>Black</i>) The low end of the photoelectron spectrum seen in the LSV in Phase-I, with ^{14}C -rich TMB, at 50% concentration, compared to (<i>red</i>) the Phase-II low- ^{14}C TMB at $\sim 5\%$ concentration. The dramatic decrease in ^{14}C background is evident.	234
8.10	Photoelectron spectrum of the LSV clusters during Phase-II: (<i>black</i>) for all clusters and (<i>red</i>) for clusters starting within a $\pm 25\text{ ns}$ prompt coincidence with the TPC. Region A is the low energy signal that consists largely of afterpulses, Region B is the ^{60}Co peak, and Region C is the ^{208}Tl peak.	238
8.11	LSV cluster start time during the ~ 70 day underground argon campaign, showing the time distribution relative to the LSV trigger time of scintillation events in the LSV. The large spike around $-6\ \mu\text{s}$ corresponds to prompt coincidences between the LSV and the TPC. The blue and green curves highlight high and low energy clusters, respectively, showing that the different origins of the features in the cluster start time profile.	239
8.12	(<i>Top</i>) Photoelectron spectrum of the ^{14}C in the LSV in Phase-II, evaluated with the clustering method. The black curve is the measured low energy spectrum after applying cuts to remove afterpulses, and the red curve is the theoretical ^{14}C spectrum, convolved with a Gaussian response function, fit over the range of 35–130 PE. (<i>Bottom</i>) The ratio of the fit spectrum to the measured spectrum over the fit energy range.	240

8.13	(<i>Top</i>) The (<i>black</i>) photoelectron spectrum of the ^{60}Co in the LSV in prompt coincidence with a signal in the TPC compared to (<i>red</i>) a fit of a simulated energy distribution convolved with a Gaussian response function. The fit, performed around the two peaks from the 1.17 and 1.33 MeV γ -rays, gives a light yield of 0.592 ± 0.011 PE/keV. (<i>Bottom</i>) The (<i>black</i>) photoelectron spectrum of the ^{208}Tl in the LSV in prompt coincidence with a signal in the TPC, compared to (<i>red</i>) a fit to a Gaussian plus a falling exponential fit to it. The 2.6 MeV γ -ray peak is seen at 1440 ± 5 PE.	243
8.14	An example waveform of a cosmogenic signal in the LSV. The inset plot is zoomed in around the time that the muon passed through the detector.	245
8.15	(<i>Top</i>) Total charge collected in the LSV for each event in the ~ 70 day underground argon campaign [42]. (<i>Blue</i>) All events, including those with cosmogenic signals and those without. (<i>Red</i>) Events with coincidental high energy signals in the WCV, predominantly due to cosmogenic signals in both detectors. (<i>Bottom</i>) Total charge collected in the LSV and in the WCV per event. (<i>Red</i>) Cosmogenic signals in the LSV and WCV. (<i>Blue</i>) Front-end oscillations in LSV electronics. (<i>Green</i>) Some front-end oscillations in WCV electronics and some cosmogenic signals. (<i>Black</i>) Standard LSV and WCV signals.	246

8.16	The waveform seen in the LSV for an event seen during the 70 day underground argon campaign published in [12]. This event produced a single nuclear recoil in the WIMP energy range in the TPC. This signal was vetoed by the LSV due to the large prompt signal seen around - $6.5\mu\text{s}$ and in the left inset plot, as well as by the two capture signals: one at $7.8\mu\text{s}$ and one (seen in the right inset plot) at $69.9\mu\text{s}$. This signal was also vetoed because it produced light with a high concentration in one TPC PMT, indicating that it may have been a fission event in the TPC fused silica in which at least two neutrons were produced. . . .	250
9.1	(<i>Left</i>) A diagram of CALIS lowered into the LSV from the port in the clean room, from [40]. (<i>Right</i>) Photograph of the inside of the LSV with CALIS deployed inside of it, from [40].	254
9.2	A photograph of the $^{241}\text{AmBe}$ source, opened up. Photograph taken by Y. Suvarov.	255
9.3	Diagram illustrating the $^{241}\text{AmBe}$ (α , n) reaction. An α emitted by an ^{241}Am nucleus interacts with a ground state ^9Be nucleus to form a compound ^{13}C nucleus. This ^{13}C nucleus then decays to the ground state, first excited state, or second excited state of ^{12}C by emitting a neutron. Depending on which state the ^{12}C ends up in, it may release up to two γ -rays as it decays to the ground state. Black lines represent nuclear states. Numbers to the left of each black line represent the J_π value of that state; numbers to the right are the excitation energy of the state in keV. Branching ratios were tuned by Borexino, from I. Machulin.	256

9.4	Neutron energy spectrum from the $^{241}\text{AmBe}$ source, from [43]. Spectra are normalized so that the integral of each spectrum is the probability of the given decay mode. (<i>Blue</i>) The spectrum for the $^{241}\text{AmBe}$ reaction resulting directly in the ground state of ^{12}C . (<i>Green</i>) The spectrum for the $^{241}\text{AmBe}$ reaction going to the first excited state of ^{12}C , which emits a 4.439 MeV γ -ray. (<i>Purple</i>) The spectrum for the $^{241}\text{AmBe}$ reaction going to the second excited state of ^{12}C , which emits 3.215 and 4.439 MeV γ -rays.	256
9.5	The neutron capture signal measured during the $^{241}\text{AmBe}$ calibration campaign, showing the signal produced by neutrons capturing on ^{10}B . The signal produced by the $\alpha+^7\text{Li}$ in the capture channel going directly to the ground state of ^7Li can be seen at 29.2 ± 0.2 PE. The signal produced by the $\alpha+^7\text{Li}+\gamma$ (478 keV) produced by the capture channel that goes to the excited state of ^7Li can be seen at 264.6 ± 0.1 PE. . .	257
9.6	The time distribution for neutron captures measured during an $^{241}\text{AmBe}$ calibration campaign. (<i>Black</i>) The low energy and (<i>blue</i>) high energy peaks, corresponding to neutron captures on ^{10}B leading to the ^7Li ground state and excited state, respectively. (<i>Red</i>) An exponential plus a constant were fit to both of these curves (shown here only for the low energy peak curve), giving a decay time of $22.1\ \mu\text{s}$, consistent with the expected capture time for thermal neutrons during Phase-II.	258
9.7	(<i>Top</i>) A schematic diagram of the $^{241}\text{Am}^{13}\text{C}$ source holder, showing the lead holder and its lid, the ^{241}Am source in its container, the ^{13}C pellet wrapped in a $1\ \mu\text{m}$ thick sheet of gold foil, and the polyethylene spacer. (<i>Bottom</i>) A photograph of the $^{241}\text{Am}^{13}\text{C}$ source, covered in gold foil with the lid removed. Photograph taken by Y. Suvarov. . . .	261

9.8	Diagram illustrating the $^{241}\text{Am}^{13}\text{C}$ (α , n) reaction. An α emitted by an ^{241}Am nucleus interacts with a ground state ^{13}C nucleus to form a compound ^{17}O nucleus. This ^{17}O nucleus then decays to the ground state or first excited state of ^{16}O by emitting a neutron. If the incident α has less than 5.048 MeV of kinetic energy, kinematic constraints require that this decay go exclusively to the ground state of ^{16}O . If the decay goes to the first excited state, a e^+e^- pair is emitted as the nucleus decays to the ground state; the positron will then annihilate with an electron and produce two 511 keV γ -rays. Black lines represent nuclear states. Numbers to the left of the black lines represent the J_π value of that state; numbers to the right are the excitation energy of the state. Data from [44, 45].	263
9.9	GEANT4 simulations of the energy deposited in the ^{13}C pellet by ^{241}Am α s propagating through the source, through (<i>blue</i>) $0\ \mu\text{m}$ of foil, (<i>green</i>) $1\ \mu\text{m}$ of foil, or (<i>purple</i>) $2\ \mu\text{m}$ of foil, and into the pellet. . . .	264
9.10	Neutron energy spectra coming from the $^{241}\text{Am}^{13}\text{C}$ source, as simulated for source configurations with (<i>blue</i>) no gold foil between the ^{241}Am source and the graphite pellet, (<i>green</i>) $1\ \mu\text{m}$ of gold foil between them, and (<i>purple</i>) $2\ \mu\text{m}$ of foil between them.	267
9.11	The energy-angle distribution of neutrons emitted by the $^{241}\text{Am}^{13}\text{C}$ source with $1\ \mu\text{m}$ of gold foil between the ^{241}Am source and the graphite pellet target. The angle of the neutron θ_n is given relative to the line normal to the gold-graphite interface. The colors of each bin correspond to the probability of a neutron being generated with the energy and angle given by that bin per α -decay in the ^{241}Am source.	268

9.12	The reconstructed quenched recoil energy vs. f_{90} for single scatter recoils observed in the TPC during the $^{241}\text{Am}^{13}\text{C}$ calibration campaign. Events with high $f_{90} \gtrsim 0.84$ are primarily events with Cherenkov light in the Teflon or fused silica. The upper band (approximately $0.5 \lesssim f_{90} \lesssim 0.84$) is due to nuclear recoils from $^{241}\text{Am}^{13}\text{C}$ neutrons scattering in the LAr, while the lower band with $0.1 \lesssim f_{90} \lesssim 0.5$ is due to electron recoils due to backgrounds in the TPC, and the lowest band with $f_{90} \lesssim 0.1$ is due to S2 triggers and events near the extraction grid.	270
9.13	Recoil energy spectrum in the TPC during the $^{241}\text{Am}^{13}\text{C}$ calibration campaign between 20–500 keV, after correcting for Lindhard-Birks quenching, in coincidence with a neutron thermalization and capture signal.	272
9.14	LSV cluster start time vs. charge for clusters in prompt coincidence with a single nuclear recoil in the TPC passing all cuts mentioned in the text. The <i>early prompt</i> region can be seen left of -12.8 ns, while the late prompt region can be seen to the right.	273
9.15	The prompt LSV spectrum showing LSV clusters in prompt coincidence with nuclear recoils in the TPC. (<i>Black</i>) Shows all prompt signals, (<i>blue</i>) shows the spectrum for early prompt events, and (<i>green</i>) shows the spectrum for late prompt events.	274

9.16	$^{241}\text{Am}^{13}\text{C}$ prompt coincidence spectrum in the LSV. (<i>Black</i>) Data taken during the calibration campaign. This spectrum is dominated by the neutron thermalization signal, which primarily consists of proton recoils; the kink around 700 PE is from the γ -ray produced by the first excited state of ^{40}Ar , and much of the tail at higher energies results from γ -rays produced by other inelastic scatters. (<i>Blue</i>) The best fit Monte Carlo simulation modeling the liquid scintillator quenching with Model II with $kB = 0$ and $C = 3.5 \times 10^{-5} (\text{cm/MeV})^2$. This fit has $\chi^2/\text{NDF}=1.7$	276
9.17	$^{241}\text{Am}^{13}\text{C}$ delayed coincidence spectrum in the LSV, with cuts to remove afterpulses from the prompt coincidence signal. (<i>Black</i>) Data taken during the calibration campaign, primarily resulting from the neutron capture signal. The peak at ~ 20 PE is from the signal produced by the $\alpha + ^7\text{Li}$ produced by the neutron capture reaction on ^{10}B ; the peak ~ 240 PE is from the $\alpha + ^7\text{Li} + \gamma$ -ray (478 keV) produced by the neutron capture on ^{10}B . (<i>Blue</i>) The best fit Monte Carlo simulation modeling the liquid scintillator quenching with Model II with $kB = 0.022$ and $C = 5 \times 10^{-6} (\text{cm/MeV})^2$. This fit has $\chi^2/\text{NDF}=1.5$. .	277
9.18	χ^2/NDF versus kB and C values in Quenching Model II for (<i>top</i>) the prompt $^{241}\text{Am}^{13}\text{C}$ signal, (<i>middle</i>) the delayed $^{241}\text{Am}^{13}\text{C}$ signal, and (<i>bottom</i>) both signals combined. These three fits have χ^2/NDF of 1.7, 1.5, and 4.8, respectively.	278
9.19	Quenching factors for (<i>black</i>) electrons, (<i>blue</i>) hydrogen nuclei, (<i>light green</i>) helium nuclei, (<i>purple</i>) lithium nuclei, (<i>orange</i>) beryllium nuclei, (<i>teal</i>) boron nuclei, (<i>red</i>) carbon nuclei, (<i>pink</i>) nitrogen nuclei, and (<i>dark green</i>) oxygen nuclei depositing all of their energy in the LSV. .	280

10.1	Fraction of S1 photoelectrons that come from nuclear recoils for neutrons simulated between 0–10 MeV in the TPC PMT stems. For this analysis, nuclear recoils were defined as those with an S1 nuclear fraction above 0.8.	284
10.2	Probability of neutrons from the PMT stems surviving cumulative TPC cuts as a function of neutron energy. Each curve represents the probability of neutrons surviving each given cut as well as the cuts above it. (<i>Black</i>) All neutrons simulated at each energy bin (normalized to 1 for each bin), (<i>blue</i>) neutrons depositing any nonzero amount of energy in the active LAr, (<i>green</i>) those neutrons that produce nuclear recoils and pass the pulse shape discrimination cut, (<i>purple</i>) those neutrons that also pass the multiple scatter cut, (<i>orange</i>) those neutrons that pass fiducial z cuts, and (<i>teal</i>) those neutrons that also leave signals in the WIMP energy window.	286
10.3	Probability of neutrons from the inner cryostat flange surviving cumulative TPC cuts as a function of neutron energy. Each curve represents the probability of neutrons surviving each given cut as well as the cuts above it. (<i>Black</i>) All neutrons simulated at each energy bin (normalized to 1 for each bin), (<i>blue</i>) neutrons depositing any nonzero amount of energy in the active LAr, (<i>green</i>) those neutrons that produce nuclear recoils and pass the pulse shape discrimination cut, (<i>purple</i>) those neutrons that also pass the multiple scatter cut, (<i>orange</i>) those neutrons that pass fiducial z cuts, and (<i>teal</i>) those neutrons that also leave signals in the WIMP energy window.	287

10.4	Probability of neutrons from the outer cryostat flange surviving cumulative TPC cuts as a function of neutron energy. Each curve represents the probability of neutrons surviving each given cut as well as the cuts above it. (<i>Black</i>) All neutrons simulated at each energy bin (normalized to 1 for each bin), (<i>blue</i>) neutrons depositing any nonzero amount of energy in the active LAr, (<i>green</i>) those neutrons that produce nuclear recoils and pass the pulse shape discrimination cut, (<i>purple</i>) those neutrons that also pass the multiple scatter cut, (<i>orange</i>) those neutrons that pass fiducial z cuts, and (<i>teal</i>) those neutrons that also leave signals in the WIMP energy window.	288
10.5	Probability of neutrons from the cryostat walls surviving cumulative TPC cuts as a function of neutron energy. Each curve represents the probability of neutrons surviving each given cut as well as the cuts above it. (<i>Black</i>) All neutrons simulated at each energy bin (normalized to 1 for each bin), (<i>blue</i>) neutrons depositing any nonzero amount of energy in the active LAr, (<i>green</i>) those neutrons that produce nuclear recoils and pass the pulse shape discrimination cut, (<i>purple</i>) those neutrons that also pass the multiple scatter cut, (<i>orange</i>) those neutrons that pass fiducial z cuts, and (<i>teal</i>) those neutrons that also leave signals in the WIMP energy window.	289
10.6	Probability of neutrons from the PMT stems surviving LSV cuts as a function of neutron energy. (<i>Black</i>) Neutrons passing all TPC cuts, (<i>blue</i>) neutrons passing all TPC cuts and the LSV prompt cut, (<i>green</i>) neutrons passing all TPC cuts and the LSV preprompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all three LSV cuts.	290

10.7	Probability of neutrons from the inner cryostat flange surviving LSV cuts as a function of neutron energy. (<i>Black</i>) Neutrons passing all TPC cuts, (<i>blue</i>) neutrons passing all TPC cuts and the LSV prompt cut, (<i>green</i>) neutrons passing all TPC cuts and the LSV preprompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all three LSV cuts. . . .	291
10.8	Probability of neutrons from the outer cryostat flange surviving LSV cuts as a function of neutron energy. (<i>Black</i>) Neutrons passing all TPC cuts, (<i>blue</i>) neutrons passing all TPC cuts and the LSV prompt cut, (<i>green</i>) neutrons passing all TPC cuts and the LSV preprompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all three LSV cuts. . . .	292
10.9	Probability of neutrons from the cryostat walls surviving LSV cuts as a function of neutron energy. (<i>Black</i>) Neutrons passing all TPC cuts, (<i>blue</i>) neutrons passing all TPC cuts and the LSV prompt cut, (<i>green</i>) neutrons passing all TPC cuts and the LSV preprompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all three LSV cuts.	293
10.10	Initial energy of neutrons generated in the PMT stems for (<i>black</i>) all neutrons, (<i>teal</i>) neutrons passing all TPC cuts, (<i>green</i>) neutrons passing all TPC cuts and the LSV preprompt cut, (<i>blue</i>) neutrons passing all TPC cuts and the LSV prompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.	293

10.11	Initial energy of neutrons generated in the PMT hardware for (<i>black</i>) all neutrons, (<i>teal</i>) neutrons passing all TPC cuts, (<i>green</i>) neutrons passing all TPC cuts and the LSV preprompt cut, (<i>blue</i>) neutrons passing all TPC cuts and the LSV prompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.	294
10.12	Initial energy of neutrons generated in the inner cryostat flange for (<i>black</i>) all neutrons, (<i>teal</i>) neutrons passing all TPC cuts, (<i>green</i>) neutrons passing all TPC cuts and the LSV preprompt cut, (<i>blue</i>) neutrons passing all TPC cuts and the LSV prompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.	295
10.13	Initial energy of neutrons generated in the outer cryostat flange for (<i>black</i>) all neutrons, (<i>teal</i>) neutrons passing all TPC cuts, (<i>green</i>) neutrons passing all TPC cuts and the LSV preprompt cut, (<i>blue</i>) neutrons passing all TPC cuts and the LSV prompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.	296
10.14	Initial energy of neutrons generated in the cryostat walls for (<i>black</i>) all neutrons, (<i>teal</i>) neutrons passing all TPC cuts, (<i>green</i>) neutrons passing all TPC cuts and the LSV preprompt cut, (<i>blue</i>) neutrons passing all TPC cuts and the LSV prompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.	297

10.15	Initial energy of neutrons generated in the multilayer insulation for (<i>black</i>) all neutrons, (<i>teal</i>) neutrons passing all TPC cuts, (<i>green</i>) neu- trons passing all TPC cuts and the LSV preprompt cut, (<i>blue</i>) neu- trons passing all TPC cuts and the LSV prompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.	298
10.16	Initial energy of fission neutrons generated in the PMT stems for (<i>black</i>) all neutrons, (<i>teal</i>) neutrons passing all TPC cuts, (<i>green</i>) neu- trons passing all TPC cuts and the LSV preprompt cut, (<i>blue</i>) neu- trons passing all TPC cuts and the LSV prompt cut, (<i>purple</i>) neutrons passing all TPC cuts and the LSV delayed cut, and (<i>orange</i>) neutrons passing all TPC cuts and all LSV cuts.	299
A.1	(α, n) yields calculated by the present work as compared to a series of benchmarks for lithium, beryllium, boron, and carbon targets. See text for a description of the reference sources. (<i>Blue solid dots</i>) Cal- culations made by the present work, (<i>green squares</i>) measurements reported by Heaton et al, (<i>empty cyan circles</i>) measurements made by Bair and Gomez del Campo, (<i>purple triangles</i>) calculations made by the SOURCES4A code.	310
A.2	(α, n) yields calculated by the present work as compared to a se- ries of benchmarks for nitrogen, oxygen, fluorine, and sodium tar- gets. See text for a description of the reference sources. (<i>Blue solid dots</i>) Calculations made by the present work, (<i>green squares</i>) measure- ments reported by Heaton et al, (<i>purple triangles</i>) calculations made by the SOURCES4A code, (<i>empty light purple squares</i>) measurements by Roughton et al. Regions highlighted in red show where the Coulomb barrier for a given nucleus falls.	311

A.3	(α, n) yields calculated by the present work as compared to a series of benchmarks for magnesium, aluminum, silicon, and iron targets. See text for a description of the reference sources. (<i>Blue solid dots</i>) Calculations made by the present work, (<i>green squares</i>) measurements reported by Heaton et al, (<i>empty cyan circles</i>) measurements made by Bair and Gomez del Campo, (<i>purple triangles</i>) calculations made by the SOURCES4A code, (<i>empty light purple squares</i>) measurements by Roughton et al, (<i>upside-down orange triangles</i>) measurements reported by Stelson and McGowan. Regions highlighted in red show where the Coulomb barrier for a given nucleus falls.	312
A.4	(α, n) yields calculated by the present work as compared to a series of benchmarks for cobalt, nickel, copper, and zinc targets. See text for a description of the reference sources. (<i>Blue solid dots</i>) Calculations made by the present work, (<i>green squares</i>) measurements reported by Heaton et al, (<i>empty cyan circles</i>) measurements made by Bair and Gomez del Campo, (<i>purple triangles</i>) calculations made by the SOURCES4A code, (<i>empty light purple squares</i>) measurements by Roughton et al, (<i>upside-down orange triangles</i>) measurements reported by Stelson and McGowan. Regions highlighted in red show where the Coulomb barrier for a given nucleus falls.	313
A.5	(α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick boron target, as calculated both by (<i>green</i>) SOURCES4A and (<i>blue</i>) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.	314

A.6	(α , n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick carbon target, as calculated both by (<i>green</i>) SOURCES4A and (<i>blue</i>) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.	315
A.7	(α , n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick oxygen target, as calculated both by (<i>green</i>) SOURCES4A and (<i>blue</i>) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.	316
A.8	(α , n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick fluorine target, as calculated both by (<i>green</i>) SOURCES4A and (<i>blue</i>) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.	317

A.9	(α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick silicon target, as calculated both by (<i>green</i>) SOURCES4A and (<i>blue</i>) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.	318
A.10	(α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick iron target, as calculated both by (<i>green</i>) SOURCES4A and (<i>blue</i>) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.	319
A.11	(α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick chromium target, as calculated both by (<i>green</i>) SOURCES4A and (<i>blue</i>) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.	320

A.12 (α , n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick nickel target, as calculated both by (<i>green</i>) SOURCES4A and (<i>blue</i>) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.	321
---	-----

Chapter 1

Introduction

In recent years, the field of particle physics and our understanding of the standard model has taken massive strides forward, as exemplified by the discovery of the Higgs boson by the ATLAS and CMS collaborations at the Large Hadron Collider [46, 47]. However, despite this recent progress, there exists a preponderance of evidence, as discussed in Section 1.1, suggesting that this standard model describes less than 20% of the total mass of the universe, and less than 5% of the total energy density, accounting for dark energy [48].

The identity of this missing matter, called dark matter, remains one of the biggest mysteries in modern physics. The first step to understanding this dark matter is to detect its interactions with the known particles described by the Standard Model of particle physics. Studying these interactions will allow us to probe its properties and compare it to theoretical models, many of which are discussed in Section 1.2. One such model describes a class of theoretical particles called Weakly Interacting Massive Particles (WIMPs), which will be the primary focus of the detection efforts discussed in this dissertation. Various experiments trying to detect WIMPs are discussed in Section 1.3.

Due to their expectedly very low interaction cross sections, dark matter detectors require extremely high sensitivity. For WIMP detection experiments, this means that detectors must be able to discriminate between backgrounds and potential signals with a very high degree of efficiency. While the various backgrounds will be discussed more thoroughly in Chapter 4, the most problematic background for many experiments comes from neutrons scattering on the target nuclei. The focus of this dissertation will be on a highly efficient neutron veto system designed to tag neutron events, especially in the context of the DarkSide-50 detector [49, 50, 42]. The goal of this veto system is to allow a dark matter experiment to run free of neutron backgrounds and be able to perform highly sensitive dark matter searches.

1.1 Evidence for Dark Matter

We currently lack a concrete understanding of what dark matter is from a particle physics perspective. However, many astrophysical observations over the last century

have indicated that dark matter exists, is very abundant, and is made of non-baryonic particles beyond the Standard Model of particle physics. While direct evidence of dark matter interacting with other forms of matter is lacking, its effects can be observed through the gravitational pull it exerts. In particular, three important pieces of evidence come from studies of galactic rotation curves, gravitational microlensing measurements, and measurements of the cosmic microwave background.

1.1.1 Galactic Rotation Curves

Measurements of galactic rotation curves of many galaxies—the rotational velocity of stars in the galaxy as a function of their distance from the center of the galaxy—provided some of the earliest evidence for the existence of dark matter. This discovery was first made by F. Zwicky in 1933 [51], and was then seen more robustly in many galaxies observed by V. Rubin in the 1970s [52, 53]. Later, K. Freeman [54] and J. Ostriker and J. Peebles [55] showed that dark matter could explain Zwicky’s and Rubin’s observations.

Observations of spiral galaxies indicates that they typically consist of a dense central bulge surrounding a supermassive black hole, with much lower density spiral arms branching off and extending outward. The vast majority of the mass of the galaxy is in this central bulge, and so a simple galactic model can assume that all of the mass is uniformly distributed in this volume.

Doing so, we can expect a rotation curve approximately of the form

$$v(r) = \begin{cases} \sqrt{\frac{4}{3}\pi G \rho r} & \text{if } r < R \\ \sqrt{\frac{4}{3}\pi G \rho R^3 / r} & \text{if } r > R \end{cases} \quad (1.1)$$

where r is the distance from the center of the galaxy, R is the radius of the central bulge, and ρ is the density of the bulge.

As a result, we would expect to see the rotational velocity increase linearly near the center of the galaxy, and then decrease like $r^{-1/2}$ at larger radii.

Figure 1.1 show how visible matter, gas, and dark matter may all separately contribute to the shape of these rotation curves. However, as can be seen in this figure, the measured rotation curves do not agree with the prediction in Equation 1.1. We see a roughly linear increase in velocity at small r , as expected, but then the rotation curves approximately level off at large radii, rather than decrease. Extending Equation 1.1 to account for the mass in the spiral arms cannot explain this discrepancy. As we can see from the fit components in this figure, the visible and gaseous components alone cannot explain the flatness of the curves at large radii. Instead, we must either assume the existence of a dark halo or a modified theory of gravity (MOND) to explain the shapes of these curves.




Figure 1.1: Galactic rotation curves of two different galaxies measured [9]. A three parameter fit was performed on the data: the dashed lines represent the contributions from visible components, the dotted lines from gas, and the dash-dotted line from the dark halo.

1.1.2 Gravitational Microlensing

General relativity tells us that massive objects warp the local metric, so that a beam of light passing by such an object would have its trajectory appear deflected to a distant observer.

Figure 1.2 depicts an observer viewing a source being lensed by a massive body that lies closer to the observer than the source. Instead of seeing the source in its actual position, the observer sees the image deflected by some reduced lensing angle α . The reduced lensing angle is related to the deflection angle $\hat{\alpha}$ by the relation $\alpha = \frac{d_{LS}}{d_S} \hat{\alpha}$. Since α is just the difference between θ and β , we obtain the lens equation

$$\beta = \theta - \frac{d_{LS}}{d_S} \hat{\alpha} \quad (1.2)$$

The deflection angle depends on the metric assumed and the the gravitational potential ϕ that is deflecting the light, and can be found by $\hat{\alpha} = 2 \int \vec{\nabla}_{\perp} \phi ds$. For the simple case of a point mass M , we find that $\hat{\alpha} = \frac{4GM}{d_L \theta}$. We can therefore find that the observed angle θ is related to the actual angle β by the relation

$$\theta_{\pm} = \frac{1}{2} \left(\beta \pm \sqrt{\beta^2 + 4\theta_E^2} \right) \quad (1.3)$$

When the source and lens are collinear with the observer (i.e. $\beta = 0$), Equation 1.3 becomes degenerate, and the observer sees the image distorted into a ring known as an Einstein ring. θ_E in this equation is the Einstein angle, defined to be the angular separation of the lens and image in the Einstein ring. The Einstein angle can be




Figure 1.2: A cartoon showing the process of gravitational lensing as predicted by general relativity. θ is the observed angle between the image and the lens, β is the observed angle if the lens were massless (i.e. with no lensing), α is the reduced lensing angle, and $\hat{\alpha}$ is the deflection angle., d_{LS} is the distance between the lens and the source, d_{OL} is the distance between the observer and the lens, and d_{OS} is the distance between the observer and the source.

calculated by

$$\theta_E = \sqrt{\frac{4GMd_{LS}}{d_L d_S}} \quad (1.4)$$

Notice that in the non-collinear case ($\beta \neq 0$), Equation 1.3 gives two angles. For non-point mass distributions in the lens, a third angle can generally also be observed. Therefore, by observing the angular separation of the images from a gravitationally lensed source or measuring the angular diameter of an Einstein ring, the mass and position of the lens can be determined.

With sufficient resolution, the observation of a small Einstein ring with an invisible target can indicate the presence of a non-luminous celestial body with a mass as low as 10^{-7} solar masses [56]. Lensing effects on this scale are known as micro-lensing. Multiple experiments have tried to use micro-lensing to search for such bodies to see if the dark halo can be explained by celestial bodies that are baryonic but otherwise invisible because they do not give off light. These bodies are commonly referred to as Massive Compact Halo Objects (MaCHOs), and such experiments have placed an upper bound on their potential contribution to the dark halo of 1.7×10^{-7} [56, 57].




Figure 1.3: X-ray photograph of the Bullet Cluster taken by the Chandra X-ray Observatory [10].

Additionally, observations of the Bullet Cluster (see Figure 1.3) provide further evidence that dark matter cannot be accounted for by ordinary baryonic matter. The Bullet Cluster consists of a pair of galactic clusters that have collided with each other. Since the mass of these galactic clusters is dominated by interstellar gas and stars, optical techniques such as x-ray photography can be used to track the center of mass of the baryonic matter in both galaxies.

Measurements of gravitational lensing caused by these galaxies can provide an independent measure of the center of mass of all matter (dark and baryonic alike) in each galaxy. Remarkably, measurements of the center of mass using optical techniques and measurements using gravitational lensing techniques significantly differ from each other, implying that the centers of mass of baryonic matter differ from the total centers of mass by over 8σ [58]. This discrepancy is a very strong indication for the presence of a far more weakly interacting form of matter than baryonic matter, with a total mass that far exceeds the baryonic mass.

1.1.3 Cosmic Microwave Background

Shortly after the big bang, the universe was in such a hot state that photons could not travel far before interacting with other particles. This led to a period where the universe was largely opaque, and much of the universe was held in thermal equilibrium. As the universe continued to expand, regions of the universe came out of causal contact with each other, and the universe began to cool down as the average energy density decreased.

As the universe expanded, different types of particles cooled down at different rates, depending on how their energy densities dilute over space. Specifically, the energy densities of non-relativistic forms of matter decreased like $a(t)^{-3}$, where $a(t)$




Figure 1.4: Temperature variations over the entire sky as measured by Planck. The bright horizontal strip across the center is due to distortions from the Doppler shift caused by Earth’s motion [11].

is the scale factor representing the size of the universe. For photons and relativistic particles, for which the energy is given by $E = \frac{hc}{\lambda}$ and $E = pc = \frac{hc}{\lambda_{dB}}$ (where λ_{dB} is the particle’s de Broglie wavelength), respectively, the energy density drops off like $a(t)^{-4}$, due to the additional affects of red shifting.

As the energy density of each species decreased, various forms of matter began to “freeze out” as the available energy density dropped below the chemical potential needed to further create and destroy that form of matter. As this happened, the particles that were freezing out came out of thermal equilibrium with the rest of the universe.

Since distant parts of the universe were out of causal contact at this time, the rapid expansion of the universe allowed small quantum fluctuations to grow, until over- and under-densities formed throughout the universe, guided by when each type of particle happened to freeze out in different regions.

As the universe continued to cool, the energy of photons continued to decrease until they became too low to interfere with nuclear and atomic processes, allowing more complicated structures such as nuclei and atoms to form. The period after nuclei formed but before atoms became stable is known as the epoch of recombination. After this epoch ended, the universe became transparent, and photons were able to freely travel ever since, with the exception of the epoch of re-ionization, which introduced additional anisotropies as larger-scale structures that could absorb light began to form and were still densely packed. These relic photons formed what is known as the Cosmic Microwave Background (CMB), and can still be observed today. A picture of the temperature map of the CMB as produced by Planck is shown in Figure 1.4.

While the CMB is mostly isotropic, showing that the universe is largely homogeneous, small anisotropies are driven almost entirely by the temperature fluctuations in the early universe that were caused by the under- and over-densities in different regions as particles began to freeze out at different times. The intensity and size of these fluctuations depend entirely on the particle physics that was happening at this time, making the CMB a very powerful probe into fundamental properties of different species of particles.

These temperature variations can be characterized by the power spectrum of the CMB. The power spectrum is simply a measure of the amount each point deviates from the mean temperature, decomposed into spherical harmonics,

$$\Theta_{l,m} = \int Y_{l,m}^*(\theta, \phi) \frac{T(\theta, \phi) - \langle T \rangle}{\langle T \rangle} d\Omega \quad (1.5)$$

where $T(\theta, \phi)$ is the temperature at a given point in the sky, and l and m describe an angular size of $\frac{\pi}{l}$ and the orientation, respectively, of a given spherical harmonic. Since the universe is isotropic the parameter m is unimportant.

The anisotropy of the power spectrum is often written in terms of the two point correlation function of various modes of $\Theta_{l,m}$, given by the parameter C_l , where $\langle \Theta_{l,m}, \Theta(l', m') \rangle = \delta_{l,m} \delta_{l',m'} C_{l,m}$ and $C_l = \sum_m C_{l,m}$ is the magnitude of the l^{th} spherical harmonic.

The amplitude of the power spectrum at various values of l gives a quantitative measure of the anisotropies in the CMB, which allows us to determine the abundance of different kinds of particles in the universe after they froze out, as well as many other properties.

Since dark matter does not interact directly (*i.e.* on the tree-level) with photons, baryonic matter couples to photons substantially more strongly than dark matter does. This means that while fluctuations in baryonic density in the over- and under-dense regions of the early universe drive anisotropies in the CMB, fluctuations in the presence of dark matter have a much weaker effect. This allows us to use measurements of these anisotropies to determine how much of the matter in the universe is baryonic.

Figure 1.5 shows the power spectrum of the anisotropies in the CMB as measured by the Planck Collaboration [11]. A clear explanation of this spectrum in context of the COBE experiment is given in [59]. The peaks shown in Figure 1.5 are generally referred to as acoustic peaks because they relate to the density of sound waves of a given wave number when the universe became transparent to them. The acoustic oscillations are caused by inward pressure from the pull of baryons against each other, and outward pressure from the radiation exerted by photons and relativistic particles.

The first acoustic peak in the power spectrum is seen around $l \sim 200$ and corresponds to compression of the plasma. The position of this peak, as well as the subsequent peaks, depends on the curvature of the universe, since the l value of the peaks are set by the angular distance traveled by photons to their surface of last scattering. Furthermore, the greater the mass of all of the baryonic matter in the universe, the more the plasma of the early universe will compress in these acoustic




Figure 1.5: The power spectrum of the CMB as measured by Planck. The x-axis is logarithmic for $l < 50$ and linear above, and the y-axis scale is in $\frac{l(l+1)C_l}{2\pi}$ [11].

oscillations and the larger we expect the first peak and subsequent odd-numbered peaks to be.

The second acoustic peak is simply a harmonic of the first peak, but corresponds to rarefaction of the plasma. These peaks correspond to troughs in the acoustic waves; however, the multipole expansion of the CMB temperature is only sensitive to the absolute value of the acoustic waves, and so we see both of these as peaks in the power spectrum. The amplitude of this peak and subsequent even-numbered peaks depends on how much the plasma will expand in the acoustic oscillations, and so does not depend on the total baryonic mass of the universe.

Since increasing the baryonic mass of the universe causes the first peak to increase in amplitude but does not affect the second peak, we can infer the amount of baryonic matter in the universe by the ratio of these two peaks. Furthermore, since the total amount of non-relativistic mass increases the inertia of the plasma, the frequency of the harmonics of the acoustic waves decreases with the mass. This means that higher frequency oscillations (which have smaller values of l) will tend to be suppressed, causing subsequent peaks in the power spectrum to decrease in amplitude.

The power spectrum of the CMB therefore depends separately on the total mass of all non-relativistic matter in the universe and on the amount of baryonic mass, these two terms can be separately evaluated. The green line in Figure 1.5 is a fit of the Λ CDM cosmological model, which allows for a cosmological constant and cold (non-relativistic) dark matter. The quality of this fit indicates that a model of the universe with 4.9% dark matter, 26.8% baryonic matter, and 68.3% dark energy is highly consistent with observations [11].

This consistency is perhaps one of the most compelling forms of evidence for the existence of some non-baryonic form of dark matter that lies outside our current Standard Model of particle physics.

1.2 Dark Matter Candidates

While we know very little about dark matter, there are a number of possible extensions to the standard model that could provide dark matter candidates. These models predict different rest masses and modes of interaction with the Standard Model. These differences allow comparisons across various direct searches to constrain the set of models that might explain what dark matter is. In addition to direct searches, indirect searches look for cosmic dark matter particles annihilating with each other into Standard Model particles, while those done at colliders like the Large Hadron Collider look for Standard Model particles annihilating with each other to produce dark matter. By studying dark matter interactions through these interactions, we can learn more about what dark matter is or isn't as we constrain the space of what model or models might be consistent with observations.

One important categorization for different models is whether it is hot (relativistic), warm, or cold (non-relativistic). Since the total mass of dark matter is more than five times greater than that of baryonic matter, the formation of galactic structure in the early universe was heavily guided by the motion of dark matter. In particular, cold dark matter particles will more easily end up trapped in each other's gravitational potentials, meaning that they will form more well-defined large scale structures. As the dark matter particles become hotter and hotter, the increased kinetic energy means that dark matter particles that start out farther from each other will mostly remain gravitationally unbound, and so the large scale structure gets progressively more smeared. Such simulations have shown that hot dark matter can constitute no more than 10% of the total mass of dark matter in the galaxy in order for the structure formation in these simulations to be consistent with the observed universe [60].

While there are many theories that may explain dark matter, we discuss only a few of them here, limiting the discussion to models predicting warm or cold dark matter consistent with observations.

1.2.1 Axions

The Strong CP Problem Much of the Standard Model of particle physics has been built on a series of symmetries that have been observed in nature. These symmetries can generally be expressed in terms of a group structure, and the properties of these groups can be used to formally describe particles and their interactions. For example, a gauge field in particle physics is a field that remains invariant under a continuous series of local transformations. A formal mathematical treatment of these groups, including the allowed algebraic operations that can be performed on them, can be used to completely describe electromagnetic interactions when combined with information about the electromagnetic charge of fundamental particles.

Three important discrete symmetries that are generally used to give Standard Model interactions better context are the CPT symmetries. These are a set of three discrete transformations:

- **Charge conjugation** : Reversing the sign of all charges involved in the interactions (i.e. $q \rightarrow -q$), which typically maps particles to their antiparticles

- **Parity reversal** : Changing the sign on all spatial terms in the described interactions (i.e. $x_i \rightarrow -x_i$), like playing everything out in a world that has been mirrored in all spatial dimensions
- **Time reversal** : Changing the sign on all time terms in the described interactions (i.e. $t \rightarrow -t$), like running the entire interaction in reverse

Eigenvectors of these transformations have eigenvalues of 1 or -1 and are said to be even or odd states, respectively, depending on whether they remain completely unchanged under this transformation or pick up a factor of -1. It is generally expected that all Standard Model interactions should be even under the total CPT transformation, however some interactions may be allowed to be odd under a pair of these transformations, as long as it is even under the third. Since the eigenvalue of an interaction under one transformation can be determined by the the eigenvalues under the other two transformations, people often just refer to CP symmetry to discuss whether or not an interaction is even under all three transformations (making it a CP-even interaction) or if it is odd under exactly two of these transformations (making it a CP-odd interaction). Fields whose terms are all the same sign under CP transformations are said to be CP-invariant, while those that have CP-even and CP-odd interactions are CP-violating.

Another, more complicated gauge field in the Standard Model is the gluon field, describing strong interactions in quantum chromodynamics (QCD). The portion of an effective Standard Model Lagrangian describing QCD interactions can be written as

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4}G_{\mu\nu}^a G^{a\mu\nu} + \sum_r \bar{q}_r^a (i\not{D}_a^b - m_r\delta_a^b)q_{rb} + \frac{\bar{\theta}}{32\pi^2}G_{\mu\nu}^a \tilde{G}^{a\mu\nu} \quad (1.6)$$

where $G_{\mu\nu}^a$ is the gluon gauge tensor field, q_r^a is a quark of charge r , flavor a , and mass m_r , and \not{D} is the gauge-invariant differential operator. The tilde over the G field denotes that that term is meant to be the dual field.

The first term in Equation 1.6 describes the interactions between gluons, while the second term describes the interactions of quarks with gluons and each other. Both of these terms describe QCD interactions that have been predicted and verified by high energy experiments, and both of these terms are CP-even. The third term, however, is often referred to as the θ term, and describes CP violating QCD interactions. This term is often included in the Standard Model Lagrangian because it obeys all of the symmetries that are typically expected of QCD. Since all mathematically allowed terms describing electromagnetic and weak interactions have been observed in nature, we would naively expect the same to hold for strong interactions. However, thus far no experiment has managed to find any CP violation in the strong sector, indicating that θ is vanishingly small.

One common avenue considered in looking for strong CP violation is the neutron electric dipole moment. Experiments so far have already measured a nonzero magnetic dipole moment in the neutron. Since magnetic dipole moments are odd under T but even under P, while electric dipole moments are odd under P but even under T, the presence of a nonzero electric dipole moment would indicate strong CP violation,

as predicted by the θ term. However, several experiments have tried to measure the neutron electron dipole moment and so far the most precise measurements have failed to find any, placing an upper bound on θ of [61]

$$|\theta| < 3.0 \times 10^{-10} \quad (1.7)$$

This value is incredibly small, and this vanishingly small value of the θ term often seems unnatural. This is generally referred to as the strong CP problem, and seems to indicate that something more might be going on to suppress this term so much.

Peccei-Quinn Symmetry One potential solution to the strong CP problem was proposed by Robert Peccei and Helen Quinn and is referred to as Peccei-Quinn symmetry. This solution describes a global chiral $U(1)$ symmetry acting separately on left- and right-handed quarks. Since the Higgs mechanism couples quarks of different chirality, Peccei-Quinn symmetry requires at least two Higgs doublets that both transform differently under this symmetry in order for the total Lagrangian to remain invariant. As a result of this relation to the Higgs mechanism, a complete description of Peccei-Quinn theory depends on the model of the Higgs mechanism being considered.

It is possible to write a set of Peccei-Quinn symmetry transformations that have the net effect of changing the value of θ . This means that this symmetry allows us to continuously transform θ in such a way that all QCD theories would be equivalent to a theory in which the CP-violating θ term is 0. This explains why no neutron electric dipole moment has been observed, and solves the strong CP problem.

The Axion The full Standard Model Lagrangian does not explicitly include quark masses. Instead, effective mass terms precipitate out of the spontaneous breaking of electroweak symmetry via the Higgs mechanism. However, it can easily be shown that these effective mass terms explicitly break Peccei-Quinn symmetry when considering the evolution of chiral projections of quark fields under Peccei-Quinn transformations. This means that spontaneous breaking of electroweak symmetry also breaks Peccei-Quinn symmetry.

The breaking of Peccei-Quinn symmetry can be countered by promoting the $U(1)$ Peccei-Quinn transformation term into a field a with mass m_a . Doing so, and modifying the Standard Model Lagrangian accordingly, we find symmetry breaking term arising from this new field exactly cancels the previous symmetry breaking.

This field is commonly referred to as the axion field and contributes the following term to the Lagrangian

$$\mathcal{L}_{\text{axion}} = -\frac{1}{4} \left(g_{a\gamma\gamma} a F_{\mu\nu} \tilde{F}^{\mu\nu} + g_{agg} a G_{\mu\nu}^i \tilde{G}^{i\mu\nu} \right) \quad (1.8)$$

where $F_{\mu\nu}$ is the electromagnetic field strength tensor, $g_{a\gamma\gamma}$ is the coupling of the axion to a pair of photons, and g_{agg} is the axion coupling to a pair of gluons. The product of $F_{\mu\nu}$ with its dual field can be written as the dot product of the electric

and magnetic fields, giving rise to the equivalent QED interaction term

$$\mathcal{L}_{a+\text{QED}} = g_{a\gamma\gamma} a \vec{E} \cdot \vec{B} \quad (1.9)$$

where \vec{E} is the electric field and \vec{B} is the magnetic field.




Figure 1.6: A Feynman diagram showing the Primakoff conversion process of an axion into a photon in the presence of an external magnetic field, as described by Equation 1.9

The interaction described by Equation 1.9 allows for three primary axion-photon interactions: axion decay ($a \rightarrow 2\gamma$), photon-to-axion coalescence ($2\gamma \rightarrow a$), and Primakoff conversion ($a \leftrightarrow \gamma$), which mixes photons and axions in strong electromagnetic fields, such as those found near nuclei or in stars. A Feynman diagram depicting the process of Primakoff conversion is shown in Figure 1.6.

These interactions make the axion a very unique dark matter candidate, as it can couple directly to photons in a strong electromagnetic field. Additionally, the mass of the axion is proportional to its coupling constant, and therefore its relic kinetic energy. This relationship means that more massive axions will be hotter than lighter axions, and so in order for the axion to be a cold dark matter candidate, we expect its mass to be on the μeV – meV scale.

Experiments searching for axions typically make use of the Primakoff conversion process to either search for stellar or cosmological axions by waiting for them to pass through resonant cavities that will convert these axions into photons which they can detect, or by trying to produce laboratory axions with a pair of resonant cavities, one having a laser and the other having photon detectors. Experiments of this latter kind are often called Light Through Wall experiments and work by trying to convert photons in one resonant cavity into an axion, which will pass through the wall into the next cavity, where it may be detected.

1.2.2 Sterile Neutrinos

There are currently three known flavors of neutrino: ν_e , ν_μ , and ν_τ , which couple to electrons, muons, and taus, respectively. These neutrinos have an incredibly light mass; measurements done by Planck indicate an upper bound on their total mass of 0.85 eV [62]. There are also three mass eigenstates that each of these neutrinos is some superposition of. Because the mass and flavor eigenstates are not simultaneously diagonalizable, these three flavors oscillate into each other. Additionally, all observed neutrino interactions have involved left-handed neutrinos, while all observed anti-neutrino interactions have involved right-handed neutrinos.

A sterile neutrino could be one of two things: a fourth flavor of neutrino that does not couple to any other known leptons, except possibly indirectly by oscillating into another flavor, or a right-handed neutrino (or left-handed anti-neutrino) that is invisible due to its failure to couple with leptons like its left-handed counterpart.

In order for the sterile neutrino to be consistent with Planck’s measurement of the total neutrino mass, its mass must either be less than 0.85 eV, or it must have a large enough mass to be non-relativistic. Planck’s measurement specifically limits the total mass of relativistic species in the early universe. Since lighter particles are expected to have higher velocities when in thermal equilibrium, a sufficiently heavy sterile neutrino may be non-relativistic and therefore not in conflict with Planck’s measurements. Alternatively, a sterile neutrino with a mass below 0.85 eV could still be a warm or cold dark matter candidate if it is produced non-thermally by some mechanism.

One commonly predicted mechanism for sterile neutrino production in the early universe involves non-sterile neutrinos oscillating into them at temperatures ~ 100 MeV. The resulting density of sterile neutrinos produced by this mechanism is close to what one would expect from a thermal relic. Numerical studies have shown that this density varies with the mass of the sterile neutrino and its mixing angle with non-sterile neutrinos; this relationship can be numerically approximated by [63]

$$\Omega_{\nu_s} \approx 0.2 \frac{\sin^2 2\theta}{10^{-8}} \left(\frac{m_s}{3 \text{ keV}} \right)^{1.8} \quad (1.10)$$

where Ω_{ν_s} is the sterile neutrino density, θ is the mixing angle, and m_s is the sterile neutrino mass.

Sterile neutrinos produced through this mechanism will typically be warm, and so their mass may be constrained by the level of small-scale structure formation we see in the universe, which may be suppressed by dark matter with too low of a mass and therefore too high of a speed. Astrophysical observations currently limit the mass of sterile neutrinos produced through this mechanism to $m_s \gtrsim 10$ keV and $\sin^2 2\theta \lesssim 10^{-9}$ [63]. However, as we can see in Equation 1.10, sterile neutrinos with lower masses cannot explain all of the dark matter we observe today. Therefore, if sterile neutrinos are near this lower mass bound, there must be at least one other form of dark matter; however, sufficiently massive sterile neutrinos may explain the modern dark matter density without violating small-scale structure observations.

A massive right-handed sterile neutrino can potentially explain the seeming unnaturalness of the observed neutrino masses, which are several orders of magnitude lighter than all other Standard Model particles. Such a sterile neutrino could explain the masses through the seesaw mechanism, which treats each neutrino as a doublet $(\nu_i^L \ \nu_i^R)$, where ν_i^L is a Standard Model left-handed neutrino of flavor i , and ν_i^R is a corresponding right-handed partner that is singlet under weak isospin symmetry (and therefore sterile). This allows us to add mass term to the Standard Model Lagrangian

of the form

$$\begin{aligned}\mathcal{L}_\nu &= \frac{1}{2} \begin{pmatrix} \nu_i^L & \nu_i^R \end{pmatrix} \begin{pmatrix} B' & M \\ M & B \end{pmatrix} \begin{pmatrix} \nu_i^L \\ \nu_i^R \end{pmatrix} \\ &= \frac{1}{2} B' \nu_i^{L,\alpha} \nu_{i,\alpha}^L + \frac{1}{2} B \nu_i^{R,\alpha} \nu_{i,\alpha}^R + M \nu_i^{R,\alpha} \nu_{i,\alpha}^L\end{aligned}$$

Since ν_i^L interacts with the electroweak force, renormalizability forces B' to be zero. However, since ν_i^R is uncharged under all Standard Model gauge symmetries, B is a free parameter that can have any value. The last term is forbidden under electroweak gauge symmetry; however, an effective nonzero value for M can be generated by the spontaneous breakdown of electroweak gauge symmetry under the Higgs mechanism. This mechanism will naturally generate a value for M on the scale of the vacuum expectation value for the Higgs field, around ~ 200 GeV.

The eigenvalues of the neutrino mass matrix in the above equation are given by $\lambda_\pm = \frac{1}{2} (B \pm \sqrt{B^2 + 4M^2})$. In the case where $B \gg M$, these simplify to $\lambda_+ \approx B$ and $\lambda_- \approx -\frac{M^2}{B}$. This will generate two masses, one very large and one very small, whose masses tend to vary inversely of each other, as B increases. This relationship is typically called the seesaw mechanism.

Depending on whether the neutrino is a Dirac or Majorana particle and on specific details of how the seesaw mechanism generates neutrino masses, the sterile neutrino is typically predicted to have a mass in the 100s of MeV scale or above ~ 100 GeV to ~ 1 TeV [64, 65, 66]. This mass range would allow the sterile neutrino to be a viable cold or warm dark matter candidate.

1.2.3 Asymmetric Dark Matter

Another, more observationally-driven explanation of dark matter is the asymmetric dark matter model. This model is based on the observation that the current density of dark matter is approximately five times that of Standard Model matter, and posits that dark matter particles have antiparticles.

Somehow Standard Model particles in the early universe came to outnumber the antiparticles, creating the matter-dominated universe we see today. This small excess of matter over antimatter can be written as

$$\eta(B) = \frac{n_B - n_{\bar{B}}}{s} \simeq 10^{-10} \quad (1.11)$$

where n_B and $n_{\bar{B}}$ are the number densities for baryons and anti-baryons, respectively, and s is the entropy density. While the mechanism for this excess is not well understood, the asymmetric dark matter model postulates that dark matter came to dominate anti-dark matter in the early universe through a similar mechanism.

Assumptions that the number density of dark matter is currently comparable to that of baryonic matter therefore predict a dark matter mass on the scale of 1–15 GeV. Many more detailed models exist that typically assume that the matter-antimatter asymmetry among Standard Model matter drives the asymmetry for dark

matter, vice-versa, that both are independently asymmetric, or that a linear combination of Standard Model and dark matter and anti-matter is conserved. These models often include more details about a dark sector that may include many dark matter-dark matter interactions that do not interact directly with the Standard Model. Depending on the details of the theories, asymmetric dark matter theories typically predict masses anywhere from the 10 MeV— 100 GeV scale. Many of these models may also constitute viable WIMP candidates [67].

1.2.4 Weakly Interacting Massive Particles

The WIMP is possibly the most generic of the dark matter candidates, as there are many independent theories that predict viable WIMP candidates, such as sterile neutrinos, as discussed in Section 1.2.2, and asymmetric dark matter, as discussed in Section 1.2.3. A WIMP is any particle that only directly interacts with matter via the weak and gravitational forces and has a large mass, typically on the scale of 1 GeV and above. The large mass is required for the particles to be cold dark matter candidates, as required by structure formation. While there is no reason to require that dark matter be weakly interacting, this is a reasonable assumption because most Standard Model particles interact with the weak force.

The WIMP Miracle

WIMP interaction cross sections on the weak scale predict a relic density for cold WIMP dark matter consistent with the observed dark matter density. In the early universe, the annihilation rate per particle for a particle such as a WIMP is given by

$$\Gamma_- = \int dv \sigma \frac{d\Phi}{dv} = \langle \sigma \Phi \rangle = n \langle \sigma v \rangle \quad (1.12)$$

where σ is the particle-antiparticle annihilation cross section, Φ is the flux of antiparticles, n is the antiparticle number density, and v is the speed of the antiparticles. Since Γ_- is the annihilation rate per dark matter particle, the total annihilation rate is given by

$$\frac{dn}{dt} = -\langle \sigma v \rangle n^2 \quad (1.13)$$

Similarly, other particles may annihilate to produce dark matter. This process is described by

$$\Gamma_+ = \sum_i \int dv \sigma \frac{d\Phi_i}{dv} = \sum_i \langle \sigma \Phi_i \rangle = \sum_i n_i \langle \sigma v \rangle e^{-m_{\chi}/T} \quad (1.14)$$

where T is the temperature of the universe at that time, and particles denoted by the subscript i are those that can annihilate to form a dark matter particle. The exponential factor at the end suppresses the dark matter production rate according to the Boltzmann distribution since dark matter can only be produced when the total energy of the collisions that form it is more than twice the dark matter particles' mass.

At high temperatures $T \gg m_{i,\chi}$, dark matter particles should be able to annihilate into other particles at the same rate that those particles can annihilate into dark matter. This means the number of dark matter particles approaches some equilibrium n_{eq} governed by the Boltzmann equation

$$\frac{1}{a^3} \frac{d(na^3)}{dt} = \langle \sigma v \rangle (n_{eq}^2 - n^2) \quad (1.15)$$

$$\frac{dn}{dt} = -3nH + \langle \sigma v \rangle (n_{eq}^2 - n^2) \quad (1.16)$$

where $a(t)$ is the scale factor of the universe and H is the Hubble parameter given by $H = \frac{1}{a} \frac{da}{dt}$. The equilibrium number density decreases as the universe cools, and undergoes a phase transition as the WIMPs go from being relativistic to nonrelativistic. At relativistic temperatures, a WIMP of mass M_χ and momentum \vec{p} has energy $E = p$. Integrating the Boltzmann distribution over all p , we get

$$\begin{aligned} n_{eq}(T \gg M_\chi) &= g \int \frac{d^3 \vec{p}}{(2\pi)^3} e^{-p/T} \\ &= \frac{2g}{(2\pi)^2} \int dp p^2 e^{-p/T} \\ &= g \frac{T^3}{\pi^2} \end{aligned}$$

Similarly, in the nonrelativistic case we have $E = \frac{p^2}{2M_\chi} + M_\chi$. This gives us the equilibrium relationship

$$\begin{aligned} n_{eq}(T \ll M_\chi) &= g \int \frac{d^3 \vec{p}}{(2\pi)^3} e^{-p^2/2M_\chi T - M_\chi/T} \\ &= \frac{2g}{(2\pi)^2} e^{-M_\chi/T} \int dp p^2 e^{-p^2/2M_\chi T} \\ &= g \left(\frac{M_\chi T}{2\pi} \right)^{3/2} \end{aligned}$$

where g is the number of degrees of freedom (2 for a spin 1/2 particle, 3 for a spin 1 particle). Combining these two equations, we get

$$n_{eq} = \begin{cases} g \frac{T^3}{\pi^2} & ; T \gg M_\chi \\ g \left(\frac{M_\chi T}{2\pi} \right)^{3/2} e^{-M_\chi/T} & ; T \ll M_\chi \end{cases} \quad (1.17)$$

The entropy density of the early universe can be calculated by

$$s = \frac{2\pi^2 g_*}{45} T^3 \quad (1.18)$$

where g_* is the number of relativistic degrees of freedom [68]. This is a useful parameter for understanding the thermodynamics of the early universe, and it is worth noting that entropy per co-moving volume s is conserved as the universe expands. That is to say

$$\begin{aligned}\frac{1}{a^3} \frac{d(sa^3)}{dt} &= 0 \\ \frac{ds}{dt} &= -3Hs\end{aligned}$$

It is also convenient to define the variables $Y \equiv \frac{n}{s}$ and $Y_{eq} \equiv \frac{n_{eq}}{s}$. Equation 1.16 can then be written in terms of Y and s as

$$s \frac{dY}{dt} = -\langle \sigma v \rangle s^2 (Y^2 - Y_{eq}^2) \quad (1.19)$$

We can additionally define $x = \frac{M_\chi}{T}$ to write

$$\frac{dY}{dx} = -\frac{\langle \sigma v \rangle s}{Hx} (Y^2 - Y_{eq}^2) \quad (1.20)$$

The Friedmann equation lets us relate the Hubble parameter H to the energy density of the universe ρ :

$$H^2 = \frac{8\pi G}{3} \rho = \frac{8\pi}{3M_p^2} \rho \quad (1.21)$$

where $G \approx 6.7 \times 10^{-39} \text{ GeV}^{-2}$ is the gravitational constant and $M_p \approx 1.22 \times 10^{19} \text{ GeV}$ is the Planck mass.

For heavy dark matter, like we expect for WIMPs, the self-annihilation cross section can be expanded in powers of v^2 as

$$\langle \sigma v \rangle = a + b \langle v^2 \rangle + \mathcal{O}(\langle v^4 \rangle) \approx a + 6 \frac{b}{x} \quad (1.22)$$

for constants a and b . It also helps to define $\Delta = Y - Y_{eq}$. Differentiating with respect to x then gives

$$\frac{d\Delta}{dx} = -\frac{dY_{eq}}{dx} - f(x)\Delta(2Y_{eq} + \Delta) \quad (1.23)$$

where

$$f(x) = \sqrt{\frac{\pi g_*}{45}} M_\chi M_p (a + 6b/x) x^{-2} \quad (1.24)$$

As is evident from Equation 1.15, as the universe expands and cools down, the co-moving number density of dark matter particles will approach some equilibrium, as the density becomes too low for dark matter particles to efficiently be produced or annihilate with each other. This process is called freezing out, and is generally characterized by the freeze-out temperature T_f . Similarly, we can define $x_f = M_\chi/T_f$. Doing so, we can solve Equation 1.23 in the limits where $x \ll x_f$ and $x \gg x_f$. This

gives

$$\Delta = -\frac{\frac{dY_{eq}}{dx}}{2f(x)Y_{eq}} \quad \text{for } x \ll x_f \quad (1.25)$$

$$\frac{d\Delta}{dx} = -f(x)\Delta^2 \quad \text{for } x \gg x_f \quad (1.26)$$

Dividing both sides of Equation 1.26 by Δ and integrating from the freeze-out time to infinity (assuming $\Delta_{x_f} \gg \Delta_\infty$, *i.e.* that $Y \rightarrow Y_{eq}$), we get

$$\frac{1}{\Delta} = \sqrt{\frac{\pi g_*}{45}} M_p M_\chi \frac{a + 3b/x_f}{x_f} \quad (1.27)$$

Equation 1.17 shows us that the equilibrium number density drops off as the temperature decreases. Therefore, as $x \rightarrow \infty$, we expect $Y_{eq} \rightarrow 0$. Returning to our definition of Δ , we find that

$$Y^\infty = \sqrt{\frac{45}{\pi g_*}} \frac{1}{M_p M_\chi} \frac{x_f}{a + 3b/x_f} \quad (1.28)$$

The freeze-out temperature x_f can be numerically solved through iterative solutions to

$$x_f = \ln \left(c(c+2) \sqrt{\frac{45}{8}} \frac{g}{2\pi^3} \frac{M_\chi M_p (a + 6b/x_f)}{\sqrt{g_* x_f}} \right) \quad (1.29)$$

where c is a constant of order unity. Numerical solutions to this equation show that $x_f \approx 20$ varies little with M_χ [69].

The present day WIMP relic density can then be calculated by $\rho_\chi = M_\chi n_\chi = M_\chi s_0 Y^\infty$, where $s_0 = 2889.2 \text{ cm}^{-3}$ is the modern day entropy density assuming three Dirac neutrinos [68].

If we make the simplifying assumption that $\frac{3b}{x_f} \ll \langle \sigma v \rangle$ (this assumption may or may not be justified, though the worst case scenario if it is wrong is a factor of 2 error, which is adequate for an order of magnitude estimate), we can write the present day relic density relative to the critical density $\Omega_\chi h^2$ as

$$\Omega_\chi h^2 \approx \sqrt{\frac{45}{\pi g_*}} \frac{s_0 x_f}{M_p} \frac{1}{\langle \sigma v \rangle} \quad (1.30)$$

where g_* is evaluated at the freeze-out temperature and $h = H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

Finally, plugging in values for these numbers gives the relic WIMP density to be

$$\Omega_\chi h^2 \approx \frac{10^{-36} \text{ cm}^2}{\langle \sigma v \rangle} \quad (1.31)$$

Astoundingly, this gives us WIMP self-annihilation cross sections on the weak interaction scale of the order $\times 10^{-36}$, consistent with observations. This is more remarkable

because Equation 1.30 does not depend on the mass of the WIMP, and makes very few assumptions about the annihilation process. This *WIMP miracle* shows that a weakly interacting massive particle naturally gives a relic abundance consistent with measurements, with a minimal set of assumptions, and is one of the most compelling arguments for the naturalness of the WIMP paradigm.

These weak-scale interactions hint at a WIMP mass roughly on the weak scale as well, hinting at the GeV to TeV mass range, close to the ~ 100 GeV mass scale of the the weak force carrier particles.

Interaction Terms

The physics of WIMP-nuclei interactions quickly gets very deep, involving a menagerie of interaction terms one could add add to the Standard Model Lagrangian. Before diving into this depth, we will start with the simpler, more digestible case of two-body scattering mediated by a massive scalar field. An example of such an interaction would come from WIMPs scattering via Higgs-exchange.

These interactions are generally described by the Yukawa potential, which is given by,

$$V_Y(\vec{r}) = -\frac{g_N g_\chi}{4\pi} \frac{1}{r} e^{-m_E r} \quad (1.32)$$

where g_N and g_χ are coupling constant between the mediator and a nucleus or a WIMP, respectively, \vec{r} is the distance between the two bodies, and m_E is the mass of the exchange particle.

Solving Schrödinger's equation for the scattering wave function generally gives,

$$\psi = \phi(\vec{r}) + \int d^3\vec{r}' G(\vec{r}, \vec{r}') U(\vec{r}') \psi(\vec{r}') \quad (1.33)$$

where $U(\vec{r}) = \frac{2M_\chi}{\hbar^2}$. Since we are primarily concerned with low energy recoils of non-relativistic particles, we can assume plane-wave scattering where $\phi(\vec{r})$ is a linear combination of $\phi_k(\vec{r}) = e^{i\vec{k}\cdot\vec{r}}$ and $G = -\frac{1}{4\pi} \frac{e^{ik|\vec{r}-\vec{r}'|}}{|\vec{r}-\vec{r}'|}$.

In the far-field limit where we assume the particles have traveled a long distance since scattering (*i.e.* $r' \ll r$), we can simplify G to $G = -\frac{1}{4\pi} \frac{e^{ikr}}{r} e^{-i\vec{k}'\cdot\vec{r}'}$. This lets us rewrite the wave function as,

$$\psi_k = e^{i\vec{k}\cdot\vec{r}} - \frac{1}{4\pi} \int d^3\vec{r}' \frac{e^{ikr'}}{r'} U(\vec{r}') e^{-i\vec{k}'\cdot\vec{r}'} \psi_{k'}(\vec{r}') \quad (1.34)$$

Finally, the matrix element for this scattering interaction can be calculated by,

$$\begin{aligned} f(\theta, \phi) &\simeq -\frac{1}{4\pi} \langle \phi_{k'} | U | \psi_k \rangle \\ &\simeq -\frac{1}{4\pi} \int d^3\vec{r}' e^{-i\vec{k}'\cdot\vec{r}'} \frac{e^{ik'r'}}{r'} U(\vec{r}') \psi_k(\vec{r}') \end{aligned}$$

This equation can be solved iteratively for ψ_k . In the first Born approximation, which is typically valid at the low momentum transfers we are interested in, we can assume that the incoming and outgoing wave are approximately the same—that is, $\psi_k = \phi_k = e^{i\vec{k}\cdot\vec{r}}$.

Doing so, we find that, integrating $f(\theta, \phi)$ over all angles,

$$\begin{aligned} f &\simeq -\frac{1}{4\pi} \int d^3\vec{r}' \frac{e^{ik'\cdot\vec{r}'}}{r'} U(\vec{r}') \\ &\simeq -\frac{1}{4\pi} \int_0^{2\pi} d\phi \int_{-1}^1 d(\cos\theta) \int_0^\infty dr r e^{ik'r'} U(r, \theta, \phi) \end{aligned}$$

The total cross section for scattering with some momentum transfer $\vec{q} = \hbar(\vec{k} - \vec{k}')$ can finally be calculated from the matrix element using

$$\sigma(q) = |f(q)|^2$$

For the Yukawa potential, this then gives us the cross section,

$$\sigma(q) = \frac{4g^4}{\hbar^4} \left(\frac{1}{((q/\hbar)^2 + m_E^2)^2} \right) \mu_N^2 \quad (1.35)$$

where μ_N is the reduced mass of the WIMP-nucleus system.

In the limit of zero momentum transfer (*i.e.* when the momentum transfer is small compared to the mass of the exchange particle), this becomes,

$$\sigma_0 = \left(\frac{4g_N^2 g_\chi^2}{\hbar^4 m_E^4} \right) \mu_N^2 \quad (1.36)$$

Since direct detection experiments are typically looking at low momentum transfers for which $q \sim 10\text{--}100\text{ MeV}$, an interaction with a heavy mediator $m_E \gg 100\text{ MeV}$ can be approximated well by neglecting the cross section dependence on q . This is typically the case for weak interactions where the W^\pm, Z, H have masses $\sim 100\text{ GeV}$.

This calculation neglected a description of the nucleus's response to the momentum transfer. This effect is typically described by the nuclear form factor, which can dictate whether the WIMP scatters coherently off of a single nucleus or off of an individual nucleon. The nuclear form factor tends to introduce additional q -dependence, especially for larger nuclei.

The space of potential interaction terms is defined by the types of mechanisms through which WIMPs and nuclei may interact with each other, and is explored in detail by A. L. Fitzpatrick et al. [1] as well as M. I. Gresham and K. M. Zurek [2]. To describe these interactions, they consider the effective field theory describing two-body elastic scattering,

$$\mathcal{L}_{eff} = \mathcal{O}_\chi \chi^+ \chi^- N^+ N^- \quad (1.37)$$

where \mathcal{L}_{eff} is the effective Lagrangian for the interaction and \mathcal{O}_χ is an operator that depends on the underlying physics.

Table 1.1: Possible Lorentz-invariant WIMP-nucleus interaction terms up to second order in spin and momentum transfer arising from exchange particles of spin 1 or below (*i.e.* in the non-relativistic limit) [1, 2]. \vec{S}_N and \vec{S}_χ are the spins of the nucleus and the WIMP, respectively, \vec{q} is the momentum transfer, and v^\perp is the magnitude of the component of the WIMP velocity perpendicular to the momentum transfer.

	Term	Example interaction mechanism
\mathcal{O}_1	1	scalar (<i>e.g.</i> Higgs)
\mathcal{O}_2	$(v^\perp)^2$	relativistic pseudoscalar (<i>e.g.</i> π^0)
\mathcal{O}_3	$i\vec{S}_N \cdot (\vec{q} \times \vec{v}^\perp)$	electric dipole coupling (<i>e.g.</i> dark photon)
\mathcal{O}_4	$\vec{S}_\chi \cdot \vec{S}_N$	axial-vector coupling (<i>e.g.</i> W^\pm, Z)
\mathcal{O}_5	$i\vec{S}_\chi \cdot (\vec{q} \times \vec{v}^\perp)$	magnetic dipole coupling (<i>e.g.</i> dark photon)
\mathcal{O}_6	$(\vec{S}_\chi \cdot \vec{q})(\vec{S}_N \cdot \vec{q})$	magnetic dipole coupling (<i>e.g.</i> dark photon)
\mathcal{O}_7	$\vec{S}_N \cdot \vec{v}^\perp$	axial-vector coupling (<i>e.g.</i> W^\pm, Z)
\mathcal{O}_8	$\vec{S}_\chi \cdot \vec{v}^\perp$	pseudoscalar anapole coupling (<i>e.g.</i> π^0)
\mathcal{O}_9	$i\vec{S}_\chi \cdot (\vec{S}_N \times \vec{q})$	pseudoscalar anapole coupling (<i>e.g.</i> π^0)
\mathcal{O}_{10}	$i\vec{S}_N \cdot \vec{q}$	axial-vector coupling (<i>e.g.</i> W^\pm, Z)
\mathcal{O}_{11}	$i\vec{S}_\chi \cdot \vec{q}$	electric dipole coupling (<i>e.g.</i> dark photon)

Due to the constraints imposed by the symmetries of the Standard Model, there is a finite set of possible interaction terms that can mathematically be written out. Of these terms, there are 11 possibilities that are at most second order in spin and momentum exchange and can therefore produce non-negligible cross sections in the nonrelativistic limit. These interaction terms and an example mediator that may produce such terms are summarized in Table 1.1. In addition to the operators listed in this table, all operators $q^{2n}\mathcal{O}_\chi$ also obey all relevant symmetries.

Summing all possible WIMP-nucleon interaction terms over all nucleons in a nucleus allows us to write the total WIMP-nucleus interaction term as,

$$\mathcal{L}_{eff} = F_{\mathcal{O}}(q^2)\mathcal{O} \quad (1.38)$$

where \mathcal{O} is a linear combination of \mathcal{O}_χ and $F_{\mathcal{O}}(q^2)$ is the nuclear form factor, operating on the wave function of the target nucleus. All possible WIMP-nucleus interactions can be formulated as a sum over these operators acting on the incident WIMP and target nucleus.

The possible nuclear form factors depend on what properties of the nucleus the WIMP may couple to, and are summarized in Table 1.2. Since the transverse electric operator does not appear in many UV complete models [2], it has been left out of this table.

The M_{JM} term in Table 1.2 is the Bessel spherical harmonic expansion of the spatial nucleon distribution in the nucleus. At low momentum transfers, this is dominated by the M_{00} mode, which is approximately equivalent to a Fourier transform of the nucleon distribution. Assuming isospin symmetry in the nucleus, the nucleon distribution is generally well-approximated by the charge distribution, which is what

Table 1.2: Nuclear form factor operators for different kinds of WIMP-nucleus interactions, based on multipole expansions of the nucleon distribution in the nucleus. Definitions for M_{JM} , Σ'_{JM} , Σ''_{JM} , Δ_{JM} , $\tilde{\Phi}'_{JM}$, and Φ''_{JM} are given in [1]. The indices J and M refer to the angular momentum modes of the multipole expansion.

Response Name	Interaction Type	Term	Low Energy Limit
Charge	Spin-independent	M_{JM}	$\frac{1}{\sqrt{4\pi}}$
Axial Longitudinal	Spin-dependent (l)	Σ''_{JM}	$\frac{1}{2\sqrt{3\pi}}\sigma_{1M}$
Axial Transverse Electric	Spin-dependent (t)	Σ'_{JM}	$\frac{1}{\sqrt{6\pi}}\sigma_{1M}$
Transverse Magnetic	Angular momentum spin-dependent	$\frac{q}{M_N}\Delta_{JM}$	$-\frac{q}{2M_N\sqrt{6\pi}}\ell_{1M}$
Longitudinal	Angular momentum and spin-dependent	$\frac{q}{M_N}\Phi''_{JM}$	$-\frac{q}{3M_N\sqrt{4\pi}}\vec{\sigma} \cdot \vec{\ell}$

we use in Section 1.3.1. This form factor is the momentum transfer-dependent term that was missing from cross section calculated using the Born approximation in Equation 1.35, and is applicable to Higgs-mediated interactions, for example.

The standard spin-independent analysis typically assumes the M_{JM} form factor, while the standard spin-dependent analysis is a linear combination of $\Sigma'_{JM} + \Sigma''_{JM}$. The Δ_{JM} and Φ''_{JM} terms are not usually part of the standard WIMP analyses, though there is not a compelling reason why they should be ignored. These two terms introduce a coupling between a WIMP and the nucleons' angular momenta, which are not allowed by the other terms. Such interactions may allow nuclei with no net spin to couple to the spin of a WIMP.

The full definition of each of the responses in Table 1.2 is given in [1]; in summary, each term is a combination of various Bessel spherical harmonics, vector spherical harmonics, differential and spin operators.

The matrix element \mathcal{M} can be computed from these interaction terms and form factors, as described in [1]. These matrix elements can be used to compute the scattering cross section through the relation given by,

$$\sigma = \frac{\mu^2}{\pi} \langle |\mathcal{M}|^2 \rangle \quad (1.39)$$

Depending on the nature of these interactions, B. A. Dobrescu and I. Mocioiu show that natural interaction choices may result in cross sections that scale like v^2 , $v^4, q^{\pm 2}$, and $q^{\pm 4}$ [70]. For example, neutrinos, which also interact weakly, have an elastic scattering cross section that scales like v^2 .

As will be discussed in Section 1.3.1, the signal we expect to see from a given WIMP has two model-dependent factors. The first is the WIMP-nucleus scattering cross section (which depends on the nuclear operator and interaction terms), as discussed above, as well as the velocity and density distribution of the WIMP halo. This latter factor depends strongly on the nature of WIMP-WIMP interactions.

In broad strokes, if WIMP-WIMP interactions are rare or on a scale comparable to WIMP-nucleus interactions, we would expect the WIMP halo to be mostly spherical with randomly oriented momenta. In this case, we may expect WIMPs to have an average velocity of zero with respect to the galaxy—that is, WIMPs may end up co-rotating with the galaxy. However, as the strength of the WIMP-WIMP interactions increases, the WIMP halo may virialize and pick up a net angular momentum. The magnitude of this angular momentum depends on the strength of the self-interactions. This may be further affected if WIMPs have excited states to which they can inelastically collide.

The effects of these interactions under certain theories are explored by Munoz et al. in [71]. One particular case may arise if WIMPs interact with each other on a scale comparable to the strong force. Such an interaction may cause them to co-rotate with the sun, giving them a very small velocity with respect to the earth.

WIMP Models

There exists a vast array of models that predict WIMPs that could lend insight into how they would interact with the Standard Model. While a discussion of all of these models goes far beyond the scope of this dissertation, a more complete discussion can be found in [68]. Our goal here is, instead, to cite a couple of models which tend to get a fair amount of discussion, and to hopefully provide a glimpse into the depth of the field. Here we will discuss vanilla WIMPs—the simplest form of weakly interacting dark matter that we could add to the Standard Model, requiring the fewest additions—supersymmetry, and Kaluza-Klein theory.

Vanilla WIMPs The simplest form of WIMP would be one, not necessarily motivated by any external theory, that requires the least amount of addition to the Standard Model. This would describe a WIMP that mostly fits within the Standard Model as it is today, but would only interact with Standard Model matter through neutral-current weak coupling. This could include vector interactions via Z-exchange and scalar interactions via Higgs-exchange on the tree level, and W^\pm -exchange on the loop level (as predicted by Minimal Dark Matter models) [72].

Most models predict that Z-exchange would give cross sections roughly on the order of $\sigma_{\text{SI}} \approx \alpha^2 m_N^2 / M_Z^4 \approx 10^{-38} \text{ cm}^2$, where σ_{SI} is the spin-independent scattering cross section, m_N is the mass of the target nucleus, M_Z is the mass of the Z boson, and α is the weak coupling constant. Similarly, Higgs-exchange for a 100 GeV WIMP would give cross sections on the order of $\sigma_{\text{SI}} \approx \alpha^2 m_N^4 / M_\chi^2 M_Z^4 \approx 10^{-43} \text{ cm}^2$, where M_χ is the mass of the WIMP, assuming that the dark matter-Higgs coupling is comparable to the weak coupling. If WIMPs instead couple at the loop level with W^\pm bosons, we would expect a scattering cross section on the order of $\sigma_{\text{SI}} \approx \alpha^4 m_N^4 / (4\pi)^2 M_\chi^2 M_W^6 \approx 10^{-46} \text{ cm}^2$ [72]. Since we do not know the strengths of the WIMP-boson coupling constants, we typically assume they are of order unity in the above cited cross section estimates. We get some wiggle room in these predicted cross sections by allowing the coupling constants to vary by one or two orders of magnitude. However, if the

coupling constant is too low, one might begin to question the naturalness of their values, which may be an indication of deeper underlying physics.

As will be discussed in Section 1.3, direct detection experiments have largely ruled out Z-mediated interactions for WIMPs in the ~ 100 GeV mass scale, assuming that there does not exist more exotic interactions that may suppress these interactions. The next generation of tonne-scale dark matter experiments are expected to reach the sensitivity where they might be able to detect or rule out Higgs- and W^\pm -mediated interactions for this mass range.

If Higgs-mediated interactions are not seen by these experiments, it may mean that some exotic mechanism suppresses this interaction, that WIMPs have a lower mass for which these interactions have not yet been excluded, or that WIMPs get their mass through some means other than the Standard Model Higgs field. One possible mechanism for this mass generation would be a “dark Higgs” particle that gives WIMPs mass but does not interact with Standard Model particles, which will make WIMPs much harder to detect.

Supersymmetric WIMPS The existence of WIMPs is predicted by many separate extensions to the Standard Model. Supersymmetric theories, in particular, all require the existence of WIMPs that would make feasible dark matter candidates. Supersymmetric theories are typically motivated by a number of problems in particle physics, such as the hierarchy problem (why is gravity so weak?), the fine-tuning problem (small variations in Standard Model parameters can cause the Lagrangian to diverge; however, superpartners may allow these divergences to cancel out), and unification of the fundamental forces.

Supersymmetry is a correspondence between fermions and bosons that states that every Standard Model particle has a superpartner (for each Standard Model boson, there is a super fermion, and for every Standard Model fermion, there is a super boson). These superpartners typically are expected to have masses between 100 GeV and 1 TeV [73].

Supersymmetry adds an additional discrete symmetry to the Standard Model called R-parity. All Standard Model particles have R-parity of 1, while superpartners have R-parity of -1, and all interactions between these particles must conserve this value. This means that a Standard Model particle can decay into two superparticles, and two superparticles can interact to create a Standard Model particle, but a single superparticle cannot decay into a Standard Model particle. Since conservation of energy tells us that particles can only decay into other particles of equal or smaller mass, if we have a collection of superparticles, R-parity and conservation of energy mean that the lightest such particle must be entirely stable. This particle is typically called the Lightest Supersymmetric Particle (LSP). Since these particles would only interact with Standard Model particles through gravitational and weak interactions and would have mass scales in the WIMP range, LSPs make very good WIMP candidates.

While the details of the LSP vary across different supersymmetric theories, the most commonly predicted LSP is the neutralino—an admixture of the bino (super-

partner to the photon), the neutral wino (superpartner to the Z boson), and the Higgsinos (superpartners to the Higgs boson). Other common candidates include the sneutrino (superpartner to the neutrino) and the gravitino (superpartner to the graviton).

Kaluza-Klein WIMPs Another possible WIMP candidate comes from Kaluza-Klein theory. Kaluza-Klein theory is a theory of small extra dimensions and attempts to explain the hierarchy problem and unify the fundamental forces by allowing fields to propagate in these dimensions.

One result of these extra dimensions is that fields may pick up extra Fourier modes while propagating through them, creating additional energy levels. Since we do not see interactions in these extra dimensions, we would observe particles with higher Kaluza-Klein Fourier modes as having an increased mass, given by

$$m^2 = \sum_i \frac{n_i^2}{R_i^2} + m_0^2 \quad (1.40)$$

where R_i is the size of the i th extra dimension, n_i is the excitation level of the Fourier mode in the i th extra dimension, m_0 is the mass of the particle in the $n = 0$ mode, and m is the excited state mass.

Since all of these extra dimensions should be very small, the excited states grow in mass very quickly. Additionally, if we postulate that motion in these extra dimensions is translationally invariant, then the population of each extra dimension must be a conserved quantity, and the lowest mode becomes stable. These lightest Kaluza-Klein particles (LKPs) would be weakly interacting and very massive, and are therefore good candidates for WIMPs. Since we expect higher modes in each dimension to eventually decay to the lowest mode, it is also possible for these LKPs to be sufficiently abundant to account for the dark matter we observe. Most Kaluza-Klein models predict the LKP to be an excitation of the photon, and it has been shown the LKPs with a mass between 0.5 and 1.2 TeV would be created in sufficient abundance to satisfy our observations [74].

1.3 WIMP Detection Experiments

The crux of WIMP dark matter models is that they mostly include some way in which dark matter can interact with Standard Model matter (though there are, unfortunately, non-WIMP dark matter models that do not include these interactions). Astrophysical observations tell us that the density of the dark halo—the dark matter cloud that surrounds the galaxy—is $\sim 0.4 \text{ GeV}/\text{cm}^3$ [75].

The astrophysical phase-space distribution of WIMPs depends on the WIMP model and how they interact with each other and with Standard Model matter, as well as how they may have been produced. Since different particle and astrophysical models affect how we would expect WIMPs to interact with a detector, a set of simple (though not necessarily physical) assumptions are often used in standard analyses.

A common assumption is that the dark halo around the Milky Way galaxy is locally co-rotating with the galaxy with some speed. It is therefore typical to assume that the average WIMP speed will match the galactic rotation speed at a given distance from the center of the galaxy, roughly 230 km/s. This produces a WIMP flux on the order of $\frac{100 \text{ GeV}}{m_\chi} \times 10^5 \text{ cm}^{-2}\text{s}^{-1}$, where m_χ is the WIMP mass. With such a high flux, it is therefore possible to observe a weak interaction-scale collision within a few years of a sufficiently large and sufficiently sensitive detector riding earth through this wind of WIMPs.

1.3.1 Searching Parameter Space

This detection scheme centers around being able to detect extremely rare events, as WIMPs collide with the detector medium and produce a visible signal. Dark matter detection is therefore an exercise in low background and low statistics. Most experiments strive to have as low a background as possible, so that a small handful of events passing all of their cuts could be enough to make a claim. Nonetheless, most experiments so far have only succeeded in pushing lower and lower limits for where WIMPs could live in parameter space.

The standard parameter space used to discuss WIMP searches plots WIMP mass against WIMP-nucleon cross section. If a limit has been placed, an exclusion curve is drawn in parameter space, excluding the entire region above the curve. If a WIMP is detected, the experiment instead plots a point with a region in parameter space, showing where the detected WIMPs might fall.

An important point to discuss before saying any more about these plots is that the details that go into determining parameters depends very heavily on both the astrophysical and particle models used to describe WIMP physics. Different models affect different experiments differently, and so the entire picture of the WIMP parameter space is very heavily dependent on the models assumed. The standard WIMP model makes very basic assumptions about how WIMPs interact with nuclei and their configuration and velocity distribution in the dark halo. These assumptions are not very well motivated beyond their simplicity; however, the important purpose that they serve is that they provide a common language for experiments to use when discussing their results, and they act as a standard basis on which to compare various experiments. In fact, experiments with seemingly contradictory results are among the most useful. While it is possible that the contradictions may be formed by errors in one experiment or another, many of these contradictions are model dependent, and so experiments that disagree with each other allow us to narrow down the number of potential astrophysical and particle models that may describe WIMPs. For this reason, is actually very important that many different types of experiments and many different target nuclei be used to try to gain a better understanding of WIMPs.

That said, we will make these simple assumptions for now and follow the discussion of [75] to understand how experiments place such limits.

The rate at which a detector should detect WIMPs depends on the flux of WIMPs passing through the detector, the probability of a WIMP interacting with the target, and probability of the detector seeing the interaction. Assuming 100% detector effi-

ciency for interactions that produce detectable signals, the simplest case expression for the rate of dark matter interactions per unit detector mass is given by

$$R = \frac{N_0}{A} \eta_\chi \langle \sigma_0 v \rangle \quad (1.41)$$

where $\frac{N_0}{A}$ is the number of target nuclei per unit total target mass, σ_0 is the WIMP-nucleus scattering cross section, η_χ is the number density of WIMPs, and $\langle v \rangle$ is the average WIMP velocity.

While, as discussed in Section 1.2.4, many WIMP models predict that the WIMP-nucleus cross section depends on the WIMP velocity, the simplest model, which is usually assumed for comparing direct detection experiments, is that the only operator in the WIMP effective field theory is \mathcal{O}_1 . It should be noted that while this assumption is far from guaranteed to be correct, it is not altogether an unlikely model. Models that have WIMPs scattering with nucleons via Higgs-exchange—the models of WIMP interaction that requires the least modification to the Standard Model, and therefore are attractive candidates—follow this assumption. Conveniently, under this assumption the cross section becomes totally independent of the WIMP velocity, allowing us to rewrite Equation 1.41 as

$$R = \frac{N_0}{A} \eta_\chi \sigma_0 \langle v \rangle \quad (1.42)$$

We will assume this model throughout this calculation. However, it is important to bear in mind that different models may yield very different sensitivity and exclusion curves, and an experiment that are more sensitive than another under one model may be less under another.

While details of the WIMP phase space distribution depend on the model of WIMP—for example, some models predict that the WIMP halo surrounding the galaxy would collapse into a disk while many others predict the halo to remain mostly spherical (though different models still have the distributions of WIMPs within the sphere)—the simplest assumption is that WIMPs follow a Maxwellian velocity distribution. In particular, we assume to first order that the WIMP phase space distribution can be described by

$$dn = k \eta_\chi f(\vec{v}, \vec{v}_E) d^3v \quad (1.43)$$

where k is a normalization constant, and $f(\vec{v}, \vec{v}_E)$ is the WIMP velocity distribution relative to the earth, which depends on both the WIMP velocity, v , and the velocity of the earth relative to the galaxy, $v_E \approx 244 \pm 15$ km/s (depending on the time of the year). Assuming a Maxwellian distribution, this velocity distribution is then given by

$$f(v, v_E) = e^{-(\vec{v} + \vec{v}_E)^2 / v_0^2} \quad (1.44)$$

$$= e^{-(v^2 + v_E^2 + 2vv_E \cos \theta) / v_0^2} \quad (1.45)$$

where v_0 is the average WIMP speed relative to the earth. Assuming that the WIMPs are moving with roughly the same speed as the local galactic rotation speed—another

model-dependent assumption— v_0 is generally taken to be ~ 230 km/s. If we assume the direction of the WIMPs is isotropically distributed, we can average over all angles in Equation 1.45 to get

$$\bar{f}(v, v_E) = \int_0^{2\pi} d\phi \int_{-1}^1 d(\cos \theta) v^2 f(v, v_E) \quad (1.46)$$

$$= \frac{\pi v_0^2 v}{v_E} \left(e^{-(v-v_E)^2/v_0^2} - e^{-(v+v_E)^2/v_0^2} \right) \quad (1.47)$$




Figure 1.7: WIMP speed distribution based on Equation 1.47 for $v_0 = 230$ km/s, $v_{esc} = 600$ km/s, and (green) $v_E = 244$ km/s or (blue) $v_E = 0$ km/s.

Figure 1.7 shows the WIMP speed distribution described by Equation 1.47. The mean of this distribution is ~ 350 km/s when we use $v_E = 244$ km/s, and 260 km/s in the limit when $v_E = 0$.

Since we expect the WIMPs we hope to detect to be gravitationally bound in the galaxy, we expect $|\vec{v} + \vec{v}_E| < v_{esc}$, where v_{esc} is the escape velocity in the Milky Way, assumed to be 600 km/s. We can integrate Equation 1.43, assuming radial symmetry, with the constraint that

$$\int dn = k\eta_\chi \int_0^{2\pi} d\phi \int_{-1}^1 d(\cos \theta) \int_0^{v_{esc}} f(\vec{v}, v_E) v^2 dv \equiv \eta_\chi \quad (1.48)$$

gives

$$k = \left(\frac{1}{\pi v_0^2} \right)^{3/2} \left[\text{erf} \left(\frac{v_{esc}}{v_0} \right) - \frac{2}{\pi^{1/2}} \frac{v_{esc}}{v_0} e^{-v_{esc}^2/v_0^2} \right]^{-1} \quad (1.49)$$

We can simplify Equation 1.42 by first considering the expected rate when $v_{earth} = 0$ and $v_{esc} = \infty$, and considering the effects of these variables' actual values

as a separate factor limiting the rate. We can therefore define R_0 to be the event per unit mass in this case by calculating the mean velocity, using Equation 1.44. Doing so gives

$$\int_0^\infty v k' f(\vec{v}, \vec{v}_E) d^3v = \frac{2}{\sqrt{\pi}} v_0 \quad (1.50)$$

where k' is k in the limit where $v_{esc} \rightarrow \infty$. In order to perform this integral in the limit where $v_E \rightarrow 0$, we must notice that $\frac{1}{v_E} \left(e^{-(v-v_E)^2/v_0^2} - e^{-(v+v_E)^2/v_0^2} \right)$ is equal to the derivative of e^{-v^2/v_0^2} with respect to v . Doing so, the integral reduces to $2\pi v_0$, and multiplying by k' gives the answer. It should be noted that $\frac{2}{\sqrt{\pi}} v_0 = 260 \text{ km/s}$, equal to the mean we found in Figure 1.7.

Using this mean velocity in Equation 1.42 and setting the WIMP number density η_χ to ρ_χ/M_χ , we get

$$R_0 = \frac{2}{\sqrt{\pi}} \frac{N_0}{A} \sigma_0 \frac{\rho_\chi}{M_\chi} v_0 \quad (1.51)$$

Finally, we can normalize the parameters in Equation 1.51 to their typical values and multiply out the physical constants, noting that $M_T \approx 0.932A$, to get

$$R_0 = \frac{503}{M_\chi M_T} \left(\frac{\sigma_0}{1 \text{ pb}} \right) \left(\frac{\rho_\chi}{0.4 \text{ GeV cm}^{-3}} \right) \left(\frac{v_0}{230 \text{ km/s}} \right) \text{ kg}^{-1} \text{ day}^{-1} \quad (1.52)$$

where M_χ is the WIMP mass and M_T is the target nucleus mass (both in units of GeV); ρ_χ is the local WIMP density, generally assumed to be in the range 0.3–0.7 GeV cm⁻³ with a generally accepted value of 0.4 GeV cm⁻³.

This simplification lets us rewrite Equation 1.42 as

$$dR = R_0 k \frac{\pi^{1/2}}{2v_0} v f(\vec{v}, \vec{v}_E) d^3\vec{v} \quad (1.53)$$

Since we want to understand how the rate changes as a function of the WIMP's recoil energy E_r on a target nucleus, we consider that the recoil energy from a classical elastic two body scatter is given by

$$E_r = Er(1 - \cos \theta)/2 \quad (1.54)$$

where $E = \frac{1}{2} M_\chi v^2$ is the kinetic energy of a WIMP with velocity v , θ is the center-of-mass scattering angle, and r is a kinematic factor describing the maximum fractional energy transfer that can occur, given by

$$r = \frac{4M_\chi M_T}{(M_\chi + M_T)^2} \quad (1.55)$$

If we assume that WIMP scatters are isotropic (*i.e.* uniform in $\cos \theta$), then we can write that $\frac{dR}{d\cos \theta} = \frac{1}{2} dR$ (so that multiplying both sides by $d\cos \theta$ and integrating from -1 to 1 gives consistent results on both sides). Equation 1.54 tells us that

$dE_r = -\frac{E_r}{2}d(\cos \theta)$. Solving for $d(\cos \theta)$ and combining these two equations, we get

$$\frac{dR}{dE_r} = \int_{E_{min}}^{E_{max}} \frac{1}{E_r} dR(E) \quad (1.56)$$

Since $E \propto v^2$, we can rewrite this in terms of E_0 , the average energy (corresponding to a WIMP of velocity v_0), and get

$$\frac{dR}{dE_r} = \frac{1}{E_0 r} \int_{v_{min}}^{v_{esc}} \frac{v_0^2}{v^2} dR(v) \quad (1.57)$$

$$= \frac{\pi^{1/2} v_0}{2} k \frac{R_0}{E_0 r} \int_{v_{min}}^{v_{esc}} \frac{1}{v} f(\vec{v}, \vec{v}_E) d^3 \vec{v} \quad (1.58)$$

where $v_{min} = \left(\frac{E_r}{E_0 r}\right)^{\frac{1}{2}} v_0$ is the minimum velocity a WIMP could have to produce a recoil with energy E_r .

Integrating over $f(\vec{v}, \vec{v}_E)$ gives

$$\frac{dR(v_E, v_{esc})}{dE_r} = k \frac{R_0}{E_0 r} \left(\frac{\pi^{1/2}}{4} \frac{v_0}{v_E} \left[\text{erf} \left(\frac{v_{min} + v_E}{v_0} \right) - \text{erf} \left(\frac{v_{min} - v_E}{v_0} \right) \right] - e^{-v_{esc}^2/v_0^2} \right) \quad (1.59)$$

Equation 1.52 has a term proportional to σ_0 . This is the WIMP-nucleus scattering cross section, and depends on the target nucleus as well as the momentum transfer between the WIMP and the nucleus. In order to compare results across experiments, it is useful to instead calculate the WIMP-nucleon scattering cross sections for coherent nuclear scatters (*i.e.* when the WIMP appears to recoil off of the nucleus as a whole, rather than off of an individual nucleon). To do so, we relate the WIMP-nucleus cross section σ_0 to the WIMP-nucleon cross section σ_n with

$$\sigma_0 = F(q)^2 \left(Z_T \frac{f_p}{f_n} + (A_T - Z_T) \right)^2 \left(\frac{\mu}{\mu_n} \right)^2 \sigma_n \quad (1.60)$$

where $q = \sqrt{2M_T E_R}$ is the momentum transfer, $F(q)$ is the nuclear form factor, A_T and Z_T are the target nucleus's mass number and charge, respectively, μ is the reduced mass of the WIMP-nucleus system, μ_n is the reduce mass of the WIMP-nucleon system, and f_n and f_p are the WIMP-neutron and WIMP-proton coupling constants, respectively. The factor of $(\mu/\mu_n)^2$ accounts for the different kinematic phase-space factors for the WIMP-nucleon and WIMP-nucleus systems. The factor of $\left(Z_T \frac{f_p}{f_n} + (A_T - Z_T) \right)^2$ accounts for WIMPs potentially coupling differently to protons and neutrons. It is typically assumed that WIMP-nucleon scattering obeys isospin symmetry and that $f_p/f_n = 1$ —though this assumption is model-dependent—setting this term equal to A_T^2 . Lastly, the form factor $F(q)$ encompasses the cross section's dependence on the amount of energy transferred

Following the prescription we established in Section 1.2.4, we will assume that the form factor for our nuclei is the Fourier transform of the nucleon distribution.




Figure 1.8: Nuclear form factors computed using Equation 1.62 for (blue) argon, (green) xenon, (purple) germanium, and (red) helium.

Decent accuracy can be achieved by assuming a nucleon distribution following the Fermi distribution, assuming that the nucleon density follows the charge density

$$\rho(r) = \rho_0 [1 + e^{(r-c)/a}]^{-1} \quad (1.61)$$

where $\rho(r)$ is the mass distribution and ρ_0 , a , and c are parameters describing the nuclear structure. This equation describes a spherical density distribution that drops off exponentially from some maximum to some minimum value. We can simplify this distribution to something that can be solved analytically by assuming a uniform spherical distribution and treating the component that decreases exponentially as some attenuation across some thin skin layer. This model is the Helm distribution. Calculating the Fourier transform, we get

$$F(q) = 3 \frac{j_1(qr_n)}{qr_n} e^{-(qs)^2/2} \quad (1.62)$$

where $j_1(x)$ is the first order Bessel function, $r_n \approx 1.14A_T^{1/3}$ fm is the radius of the target nucleus, $s \approx 0.9$ fm is the nuclear skin thickness. Figure 1.8 shows how the nuclear form factor varies with E_R for a few different nuclei.

Figure 1.9 shows the recoil energy spectrum for WIMPs scattering on argon and xenon targets for different WIMP masses. A key feature to note from this spectrum is that the energy spectrum drops off quickly with recoil energy. This is especially so for lower mass WIMPs; however, even for 1000 GeV WIMPs, the decrease in rate with recoil energy cannot be neglected. This highlights the importance of the low energy threshold in a WIMP search. Unfortunately, lower energies tend to also have higher background rates, making WIMP detection an experimentally challenging exercise.




Figure 1.9: Recoil energy spectrum for (*blue*) 10 GeV, (*green*) 100 GeV, (*purple*) 500 GeV, and (*orange*) 1000 GeV WIMPs on an (*top*) argon or (*bottom*) xenon target, assuming a cross section of 1 pb. Solid lines include the effects of the form factor, while dashed lines set $F^2(q) \equiv 1$.

It is also notable that these spectra all drop off mostly exponentially, with no obvious features one could look for above the background.

Figure 1.9 shows that for an argon target, the recoil energy spectrum is primarily suppressed at higher energies for higher mass WIMPs due to the kinematics of the reaction—that is to say that the maximum energy transfer in a WIMP-nucleus scattering event is limited by the argon nuclear mass. While xenon’s larger mass means that it is kinematically able to experience higher energy transfers, its recoil energy spectrum is instead primarily suppressed by the nuclear form factor. The end result

is that an argon target is likely to see higher energy recoils than a xenon target is. This means that the lower energy threshold is more crucial for a xenon experiment than it is for an argon experiment. While being able to detect low energy recoils can dramatically increase the sensitivity of both types of detectors, 1000 GeV WIMPs drop off by an order of magnitude in rate as the recoil energy goes from 0 to ~ 40 keV, while argon detectors see such a drop going from 0 to ~ 120 keV.

Since most backgrounds tend to give low energy signals—a 1 MeV neutron scattering once in the detector will deposit a maximum of 95 keV in argon and 30 keV in xenon—and achieving a low energy threshold necessary for achieving a high sensitivity requires extremely low levels of background.



Figure 1.10: A Feynman diagram depicting a WIMP scattering off of a nucleus via some exchange particle. g_χ and g_N are the coupling constants for the WIMP and the nucleus with the exchange particle, respectively.

The WIMP-nucleon cross section depends on the assumed particle physics of the interaction and depends heavily on the exact WIMP model being considered. In general, these interactions will occur through processes like the one shown in Figure 1.10 and will have cross section of the form

$$\sigma_0 \propto \left(\frac{g_\chi^2 g_N^2}{M_E^4} \right) \mu^2 \quad (1.63)$$

where g_χ and g_N are the coupling constants of the WIMP and target nucleus, respectively, to the exchange particle moderating the interaction, and M_E is the exchange particle's mass. Since g_N is proportional to A_T for coherent spin-independent nuclear scattering, we see that $\sigma_0 \propto A_T^2$, and therefore that the zero momentum transfer cross section and event rate per unit mass are both proportional to A^2 . This shows that larger nuclei typically have larger cross sections.

However, since the nuclear radius is proportional to $A_T^{1/3}$, at low momentum transfer (as one would expect from a light WIMP) the nuclear form factor is suppressed more by larger nuclei than by smaller ones, as energy transfers less efficiently between the two bodies. Additionally, the more massive a nucleus, the more momentum the WIMP will need to transfer to it to obtain a given recoil energy. These effects are more easily seen in Figure 1.8. Additionally, the kinematic factor r is highest when the target mass equals the WIMP mass, and drops off as their difference grows. For

these reasons, lighter nuclei are often considered better targets for detecting light WIMPs, while heavier nuclei are preferable for detecting heavier ones.

We can understand the mass-dependence of an experiment's sensitivity by considering the full expansion of Equation 1.59. In particular, the cross section has a term proportional to A_T^2 and a term proportional to μ_T^2 . R_0 , in Equation 1.52, has a term proportional to $1/(M_\chi M_T)$, which, when multiplied by the factor of $1/r$, is proportional to $1/\mu_T^2$. These factors mean that the sensitivity of an experiment is proportional to A_T^2 . This number can change significantly once one takes into account the range of recoil energies that an experiment is sensitive to. Since the form factors drop off very differently for different targets, as shown in Figure 1.8, large factors may increase or decrease how one target performs with respect to another depending on their energy thresholds. Since form factors can be approximated as exponential to first order, we can roughly examine the effect that the different form factors play if we assume two experiments use the recoil energy range.

In the limit where the form factor is approximately exponential, $F^2(q) = e^{-2qr_n}$ and $F^2(E_R) \simeq e^{-CA_T^{5/6} E_R^{1/2}}$, where $C \sim 3.113$ is a constant. The definite integral of this form factor is

$$\int dE_R e^{-CA_T^{5/6} E_R^{1/2}} = -\frac{2e^{-CA_T^{5/6} E_R^{1/2}} (CA_T^{5/6} E_R^{1/2} + 1)}{(CA_T^{5/6})^2} \quad (1.64)$$

when $CA_T^{5/6} E_R^{1/2} \gg 1$, the first term in the numerator dominates, and the integral is proportional to $\frac{e^{-CA_T^{5/6} E_R^{1/2}}}{A_T^{5/6}}$. When $CA_T^{5/6} E_R^{1/2} \ll 1$, the second term dominates and the integral becomes proportional to $\frac{1}{A_T^{5/6}}$. The effects of the form factor on the recoil energy spectrum can be seen by comparing the solid and dashed curves in Figure 1.9.

In a WIMP detection experiment, the total number of events passing all cuts is the observed quantity. If there are zero background events during an exposure of T kg days, the expected number of WIMP events can be calculated from

$$N = T \times \int_{E_{min}}^{E_{max}} \frac{dR}{dE_r} dE_r \quad (1.65)$$

When an experiment excludes some region of parameter space, it usually does so to some confidence level C_0 . If the detector is expected to see n background events during its exposure time, the observed rate can be compared to the expected number of events as a function of the WIMP mass and WIMP-nucleon coherent scattering cross section. If WIMP events are seen, one can then calculate the mass and cross section most consistent with the observed rate. If no events are observed, one can instead find, as a function of WIMP mass, the value for the cross section for which one would have expected to see at least one WIMP with the desired confidence.

1.3.2 Current State of the Search

Figure 1.11 summarizes the exclusion curves produced by several different experiments reporting null results in the WIMP mass range of $10\text{--}10^4\text{ GeV}$. These curves show limits placed by several different experiments that tell us, for a given WIMP mass, what maximum WIMP-nucleon cross section is consistent with the experiments' null results.




Figure 1.11: The most recent exclusion curves resulting from several different direct detection experiments at the time of writing this document [12]. (*Solid purple*) WARP, (*dotted red*) DarkSide-50 (Atmospheric Argon), (*dashed red*) DarkSide-50 (Underground Argon), and (*solid red*) DarkSide-50 (combined) used liquid argon targets. (*Dash-dotted black*) ZEPPLIN-III, (*dotted black*) PandaX-I, (*dashed black*) XENON100, and (*solid black*) LUX use liquid xenon targets. (*Solid green*) CDMS used a germanium crystal target. (*Solid cyan*) PICO used a liquid C_3F_8 target.

Already, in the $\sim 100\text{ GeV}$ mass region hinted at by the WIMP miracle, we can see many models have already been cleanly ruled out. Recent experiments like LUX have also been encroaching into the territory where we might expected to find Higgs-mediated interactions, though they have not yet been ruled out entirely yet.

The curve labeled “DarkSide-50 (combined)” in Figure 1.11 shows the latest results published by the DarkSide collaboration, combining results from a campaign with atmospheric argon [50] and a run with underground argon [12], which has significantly lower levels of the radioisotope ^{39}Ar . The DarkSide-50 experiment will be the focus of this dissertation, and will be discussed in much greater detail.

Again, it is important to remember here that these plots are very heavily model dependent, and so a WIMP that interacts through some mechanism other than Higgs exchange may very well have a cross section that is not independent of the WIMP

velocity or momentum transfer, and these differences can reshape much of parameter space. For example, models that predict $\sigma_0 \propto v^2$ tend to favor argon nuclei over xenon nuclei, since this dependence means that higher velocity WIMPs are more likely to scatter than lower velocity WIMPs, and argon’s superior form factor at high momenta will allow the experiment to achieve high sensitivity, while xenon’s less favorable form factor will make it less sensitive to these WIMPs.

1.3.3 The Importance of Low Background

The goal of a dark matter experiment is, of course, to detect dark matter. As discussed earlier, dark matter is expected to interact with Standard Model extremely rarely. Depending on the mass and cross section that dark matter has, these interactions may only take place a few times over the course of running an experiment for several years, if at all. If an experiment hopes to maximize its sensitivity to dark matter, it is therefore important to be able to claim a discovery with as few WIMP signals as possible. It is therefore crucial that a dark matter experiment strives to minimize its backgrounds—events that may perfectly mimic a WIMP signal and therefore falsely look like one. Even a low but non-zero background rate can easily overwhelm the event rate expected from a WIMP with a very low scattering cross section, and any experiment trying to claim a discovery will need to see significantly more WIMP scatters. This greater number of requisite WIMP scatters means that the detector becomes less sensitive to dark matter.

Before going into the details of how backgrounds affect sensitivity in more detail, it is first important to draw a distinction between an experiment’s sensitivity to dark matter (*i.e.* the region of parameter space in which dark matter may exist that for which the experiment would be able to claim a detection), and the limit curves discussed in Section 1.3.2. When a collaboration draws a limit curve and claims a null result, they say they do not believe they saw any WIMPs scatter in their detector. They may have seen WIMP-like events, but not enough to be statistically significantly greater than their expected background rate. In this case, they can perform a likelihood analysis or a simple background subtraction to place a limit on the mass and cross section that WIMPs may have to be consistent with the number of background events they observed up to some confidence level, typically around 90%. Alternatively, experiments such as DarkSide-50 may see zero backgrounds, and exclude a region of parameter space based on the observation that they observe no signals consistent with what was expected for WIMPs up to some similar confidence level. However, while these exclusion curves are useful for telling us which dark matter models are not possible and for guiding future searches, they do not constitute a discovery. In order to make a discovery, it is important for experiments to have a high sensitivity so that they can make a claim with as few events as possible. Since discoveries typically require a confidence of 5σ while an exclusion curve uses a 90% confidence level, it is actually statistically much easier to exclude a region of parameter space than it is to make a detection, so one would generally expect to see an experiment’s projected exclusion curve after running for some period to extend to deeper cross sections than they are capable of making a discovery at.

The statistics of dark matter detection can be framed in terms of hypothesis testing. In this case, an experiment builds a background model and predicts a number of background events with some uncertainty. This model is the null hypothesis, and in order to claim a discovery, one needs to make a measurement that rejects the null hypothesis with some degree of confidence. Within the particle physics community, it is common to require that the measurement differ from the null hypothesis by 5σ in order to claim a discovery. Assuming that the model has uncertainty described by a Gaussian distribution around the expected number of events with uncertainty σ , this corresponds to a probability of $\sim 6 \times 10^{-7}$ that the observed signal could be an artifact of statistical fluctuations and uncertainty in the model. In the field of low-statistics counting, it isn't necessarily fair to assume that the uncertainty follows a Gaussian distribution, and an observation significantly under the null hypothesis certainly does not constitute a discovery of anything—most likely it just means that the model was wrong. We are therefore only interested in the upper tail of this distribution. So for a Gaussian distribution, 5σ above the mean corresponds to a probability of $\sim 3 \times 10^{-7}$.

While a disagreement between the data and the null hypothesis by 3σ does not constitute a discovery, it is common to consider this level of disagreement as “evidence” of a signal. Following our above discussion, a signal 3σ above the null hypothesis has a probability of $\sim 1 \times 10^{-3}$. Similarly, a disagreement of 1σ , which we might call “uninteresting”, has a probability of ~ 0.16 .

The relevant question then becomes a matter of determining the statistical significance of an observation of n WIMP-like signals when your background model predicted that you would see μ_B with uncertainty σ_B . For simplicity, we will assume that the uncertainty in the model follows a Gaussian distribution described by μ_B and σ_B with a cutoff at 0 (since you cannot observe a negative number of signals). Roughly speaking, fluctuations of this mean tell you the probability that your background model “should” have predicted a different number of background events. In addition to these fluctuations, however, one must also consider Poisson fluctuations in the number of backgrounds observed. In a kind world, we would be able to make a Gaussian approximation of the Poisson distribution and add the Poisson variance and Gaussian variance from the model uncertainty in quadrature. However, this is not the case; at the small number of events relevant to dark matter searches, we must consider the full probability distribution on the number of events observed.

We therefore consider the probability distribution

$$P(n) = \int_0^\infty \left[\frac{1}{\sqrt{2\pi\sigma_B^2}} e^{-\frac{(x-\mu_B)^2}{2\sigma_B^2}} \right] \times \left[\frac{x^n e^{-x}}{n!} \right] dx \quad (1.66)$$

where we require n , the number of events observed, to be an integer. The left-hand term of the integral is a Gaussian distribution centered around μ_B with standard deviation σ_B , telling us the probability that our model “should” have predicted x backgrounds. The right-hand term of the integral tells us the probability of seeing n events given our expectation of x . The product of the terms is therefore the probability that the model should have predicted x background events and that a model with

x background events sees n ; integrating over all x tells us the probability of observing n events.

For a given requirement on our p -value (*e.g.* 3×10^{-7} , corresponding to 5σ), we can then calculate the probability of seeing n or more background events given our model by integrating $\int_{n_0}^{\infty} P(n) dn$. Finding the lowest value of n_0 for which this integral is below the requisite p -value gives the number of WIMP-like events needed.

Figures 1.12 and 1.13 show the number of WIMP-like events an experiment would need to see in order to reject the null hypothesis to 1σ , 3σ , and 5σ as a function of the number of expected backgrounds when the background model predicts the background with 100%, 10%, and 1% uncertainty. As is evident from the plot, there is fairly little difference between 10% uncertainty and 1% uncertainty because when the uncertainty in the model gets small enough, the Poisson fluctuations begin to dominate the statistics.

When the model predicts 0 background events, a single WIMP-like event can be statistically significant enough to claim a detection. If the experiment expects 0.01 or 0.1 background events, it instead needs 3 or 5 WIMP-like events to claim a discovery. However, if an experiment expects to see even one background event, the number of WIMP-like events needed for a discovery goes up to 10–15, depending on the uncertainty in the background model. This means that in order to claim a detection, a dark matter experiment expecting one background will need ~ 2 – 3 times more WIMP interactions to claim a detection than a background free experiment would need. Holding all other things equal, this constitutes a factor of 2–3 difference in WIMP sensitivity.

These calculations underscore two important points. One is that in order to be able to claim a detection for WIMPs with as low a cross section as possible, an experiment must strive to minimize their number of expected backgrounds, since the requisite number of WIMP-nucleus interactions grows quickly with the expected background rate. The other point illustrated here is that in addition to reducing the number of expected background events, it is also important to reduce the uncertainty on this prediction. One way that can be accomplished, as we will discuss later in this document, is by building the experiment with active shielding that can provide *in situ* measurements of the background rate and confirm the background model while also vetoing background events.

The effect that backgrounds have on the sensitivity of a dark matter experiment is best shown by considering an example case. Since the focus of this document is on DarkSide-50, we will choose an idealized experiment that vaguely resembles DarkSide-50. Following the prescription outlined in Section 1.3.1, we can calculate the number of WIMP-nucleus recoils we would expect to see over some period of time by integrating Equation 1.59 over the recoil energy range where the experiment is looking for dark matter for some WIMP mass and some WIMP-nucleon scattering cross section. Figure 1.12 shows us the number of true WIMP events we would need to detect in order to claim a discovery. For each WIMP mass of interest, we can therefore search the space of cross sections to calculate the cross section needed to get the desired minimum number of events.

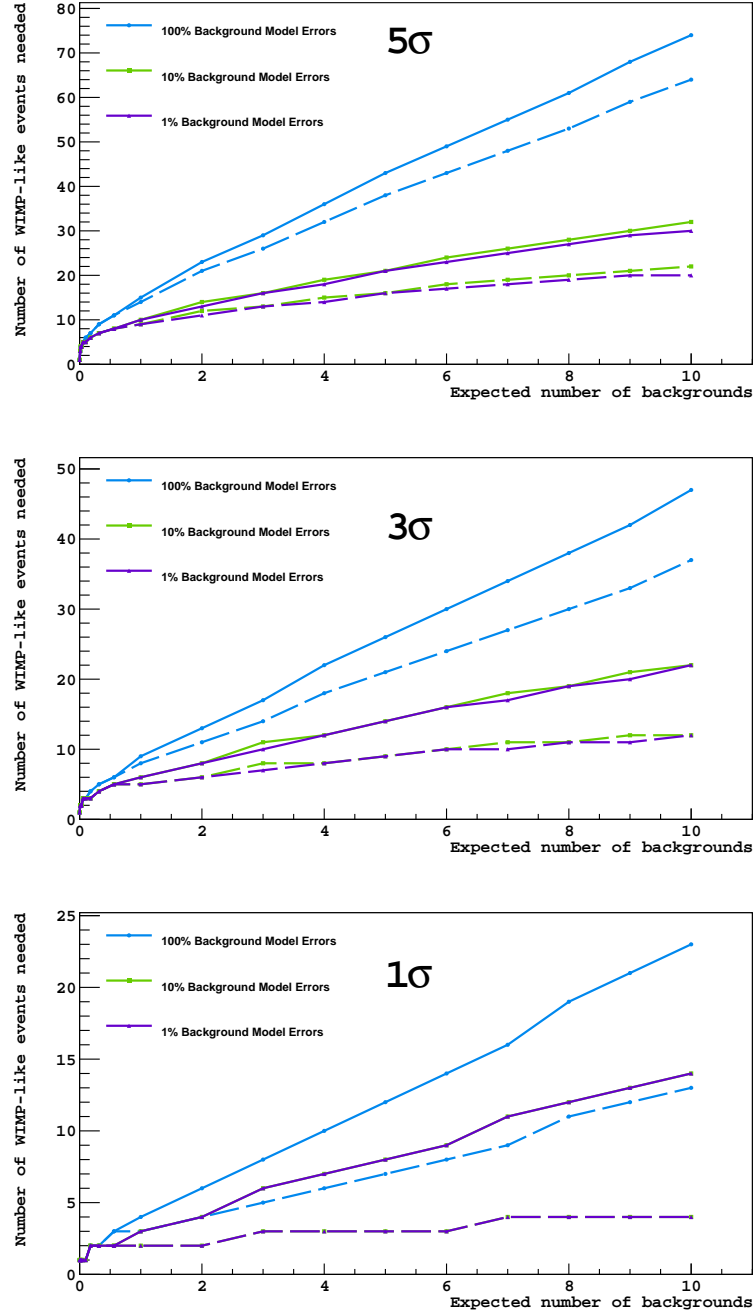


Figure 1.12: The number of WIMP-like events needed to claim a WIMP observation to (from top to bottom) 1 , 3 , and 5σ based on the predicted background rate of the experiment. Solid lines show the number of WIMP-like events needed, including backgrounds, while dashed lines show the number of WIMP-scatters needed (after subtracting the expected background). Blue, green, and purple curves were made assuming uncertainty on the background model of 100%, 10%, and 1%, respectively.

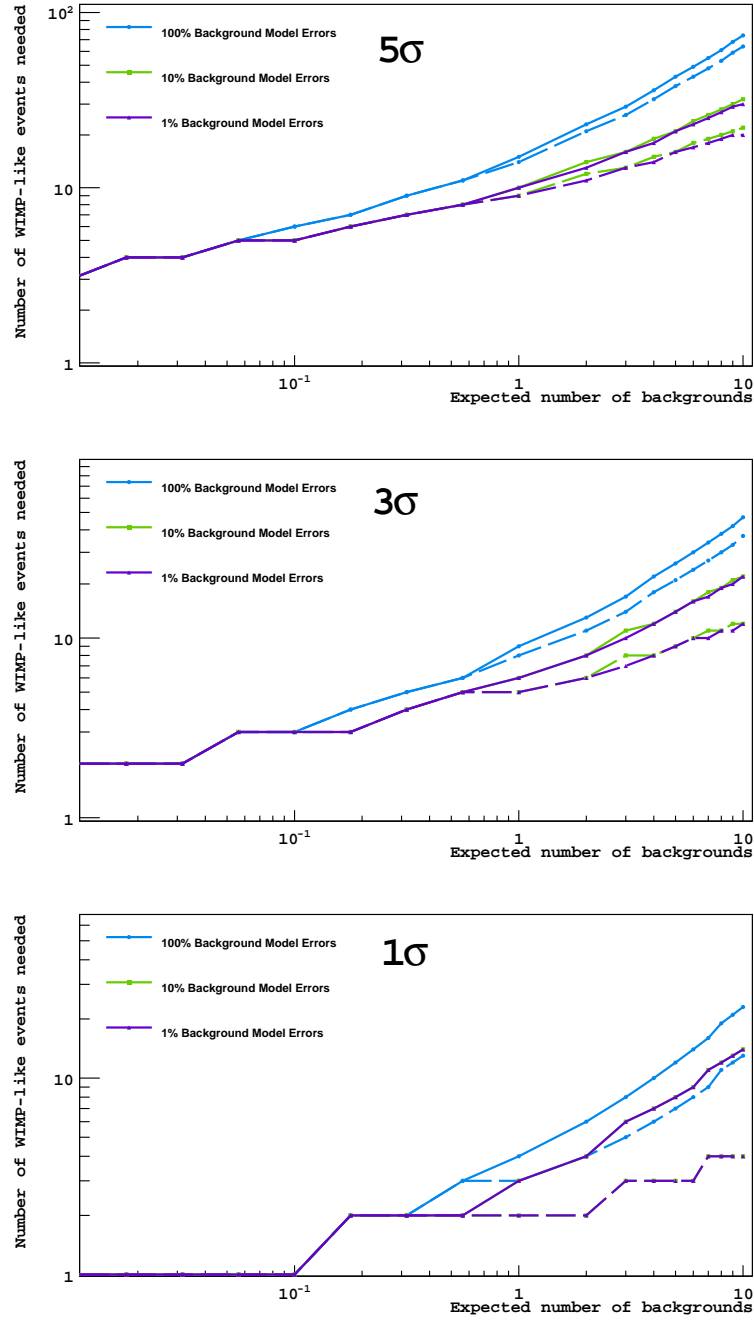


Figure 1.13: The number of WIMP-like events needed to claim a WIMP observation to (from top to bottom) 1 , 3 , and 5σ based on the predicted background rate of the experiment, on a Log-Log scale. Solid lines show the number of WIMP-like events needed, including backgrounds, while dashed lines show the number of WIMP-scatters needed (after subtracting the expected background). Blue, green, and purple curves were made assuming uncertainty on the background model of 100%, 10%, and 1%, respectively.




Figure 1.14: (*Top*) Sensitivity curves drawn for an idealized liquid argon dark matter detector similar to DarkSide-50. The area above each curve is the area of parameter space where this detector would be able to detect a WIMP; each curve represents a different threshold number of WIMP-nucleus scattering events the detector would have to observe in order to claim a detection with 50% probability. The background-free case is shown in black, where 1 event is required. Beyond that we see the curves for (blue) 2, (green) 3, (purple) 4, (orange) 5, (teal) 10, (red) 15, and (pink) 20. (*Bottom*) The same plots, drawn for an identical experiment using liquid xenon instead of argon.

Table 1.3: A summary of the parameters used to draw the sensitivity curves in Figure 1.14.

Parameter	Definition	Value	
v_0	Local galactic rotation speed	(230 ± 20) km/s	[76, 77, 78]
v_{esc}	Galactic escape speed	600^{+25}_{-20} km/s	[79]
v_{earth}	Speed of earth relative to DM halo	244 ± 15 km/s	[75]
ρ_χ	local WIMP density	$0.4^{+0.2}_{-0.1}$ GeV/cm ³	[80]
A_T	Target nucleus mass number	40	
E_{min}	Lower bound on recoil energy window	45 keV	[42]
E_{max}	Upper bound on recoil energy window	201 keV	[42]
T	Experimental exposure	54786 kg days	
C_0	Confidence for drawing sensitivity curve	50%	

Since our idealized experiment is similar to DarkSide-50, we will assume that the detector has a fiducial volume consisting of 50 kg of liquid argon, and that the experiment searches for three years, giving a total exposure of 54686 kg days. We will also assume that the detector is looking for WIMPs recoiling off of nuclei in the energy range of 45–201 keV, as was used in [12]. This detector is idealized in that we will assume it detects every nuclear recoil that occurs within it, so we will not account for additional cuts that may cause the experiment to accidentally miss a WIMP recoiling off of a nucleus.

Finally, following the prescription of [75] and using the parameters given in Table 1.3 to calculate the WIMP-nucleon coherent scattering cross section, we get the plots shown in Figure 1.14, which show this experiment’s sensitivity to WIMPs with various masses and cross sections if we require varying numbers of WIMP recoils. Each curve in this plot shows the cross section for which a WIMP of a given mass would have a 50% chance of being detected to 5σ given some threshold on the number of WIMPs needed to do so. As expected, the more recoils needed to make a detection, the less sensitive the experiment is. The loss of sensitivity at low WIMP masses comes from the kinematics of WIMP-nucleus scattering; the lighter the nucleus, the less likely it is to have enough kinetic energy to produce a recoil over the detector’s threshold energy. The loss of sensitivity at higher WIMP masses is due to a reduction of the number density of WIMPs as we assume a constant mass density.

Figure 1.12 shows us that with 100% uncertainty in the background model, expecting even one background means that we will need to observe 14 WIMPs scattering in our detector to make a claim. As can be seen in Figure 1.14, this requirement constitutes a sensitivity loss of a factor of ~ 21 compared to the case where exactly




Figure 1.15: A family of sensitivity curves based on a toy model of the LUX detector with 10065 kg days of exposure and a nuclear recoil energy range of 3–25 keV.

0 background events are expected. Even in the more optimistic cases where we have 10% or 1% uncertainty in the background model, a single expected background will mean that at least 9 WIMPs must scatter in the detector in order for the experiment to claim a discovery—reducing the sensitivity by a factor of ~ 14 .

For comparison, an experiment with 0.1 expected background may expect to be 2-3 times more sensitive than an experiment with 1 expected background, depending on the uncertainty in the background prediction.

Liquid xenon detectors offer an interesting point of comparison, here [81, 82]. From Figure 1.14, it can be seen that, all other things equal, a background-free liquid xenon experiment would be more sensitive than a background-free liquid argon experiment. However, a background-free liquid argon would be expected to outperform a liquid xenon experiment with one expected background. Figure 1.14 shows what would happen if the toy DarkSide-50 model were to use liquid xenon instead of argon. For each threshold, a factor of 1.8 increase in sensitivity is seen. This is equivalent to a scaling of $A_T^{0.49}$, smaller than the $\sim A_T^{7/6}$ scaling predicted from the mass factors alone. This is likely a result of both experiments using the same recoil energy ranges, meaning that the different form factors had a significant effect on the relative scaling. In particular, the form factor for xenon drops off much more quickly than it does for argon, meaning that the high recoil energy threshold used for these plots compromised the xenon sensitivity more so than in the argon case.

Realistically, a liquid xenon experiment can likely run with a lower energy threshold than a liquid argon experiment. Pulse-shape discrimination, the powerful tool used by liquid argon experiments to remain background free, requires a higher energy threshold than liquid xenon experiments typically use. Figure 1.15 shows a similar

toy model of the LUX detector, with $(85.3 \text{ days}) \times (118 \text{ kg}) = 10065 \text{ kg days}$ of exposure of a liquid xenon detector with a nuclear recoil energy range of 3–25 keV [83]. LUX reports a maximum sensitivity for 33 GeV WIMPs to be $7.6 \times 10^{-46} \text{ cm}^2$, close to the sensitivity predicted in Figure 1.15 with a target threshold of 5 WIMP recoils.

It is worth noting that [83] actually reports an expected background rate of $2.6 \pm 0.2_{\text{stat}} \pm 0.4_{\text{sys}}$ background events during their exposure (consistent with the $3.6 \pm 0.3_{\text{stat}}$ observed). According to Figure 1.12, a discovery with 2.6 expected backgrounds will require between 11–13 WIMP events. We would therefore expect a maximum sensitivity closer to $9.6 \times 10^{-46} - 1.4 \times 10^{-45} \text{ cm}^2$ for this exposure. However, as stated above, this prediction is based off of a simple background subtraction. LUX published an exclusion curve with a maximum upper limit of $7.6 \times 10^{-46} \text{ cm}^2$ at 33 GeV by performing a profile likelihood analysis. This exclusion curve means that they were able to rule out WIMPs with a cross section above this level (roughly falling between the 5 and 10 WIMP scatter requirement curves in Figure 1.15); however, their sensitivity to actually making a detection could have $6.9 \times 10^{-47} \text{ cm}^2$ had they been background free.

Even in the case where enough WIMP-nucleus scattering events are observed to claim a detection, it is worth remembering that WIMP detection is a practice in the statistics of low numbers, and that the ultimate goal is to understand what dark matter is. If an experiment has a nonzero number of background events, it will not necessarily be possible to tell which WIMP-like events are actually from WIMPs, and which one are from backgrounds. This means that any model testing that might require input on the details of these scattering events will be significantly hampered by the uncertainty of which events actually contain relevant information.

When we try to compare models to data, only a limited number of measurements are available for comparison. These differences will typically arise from measurements made of the WIMP-nucleus cross section measured with different nuclei, which may have different spins and form factors, as discussed in Section 1.2.4. If enough interactions are measured, the nuclear recoil energy spectrum for these interactions with each tested target nucleus may give us further information that we can compare to these models. However, since WIMPs have not yet been discovered, the probability of any experiment in the near future seeing a high enough rate of WIMP scatters to have a high-statistics nuclear recoil energy distribution that can be compared to models is pretty low. Instead, it will presumably come down to an exercise of performing likelihood fits of different models to a small amount of data. The low statistics in the energy spectrum will already be contending with uncertainties in the astrophysical model. If an experiment also has backgrounds mixed in with their WIMP signals, it will likely become impossible to distinguish one WIMP model from another based on that information. Comparing event rates between different targets might be a more robust way to test models, though Poisson statistics are likely to still be a big obstacle, and it will be very difficult to distinguish between varying event rates across experiments due to backgrounds versus different linear combinations of interaction operators. It is therefore paramount that a dark matter experiment have as low background as possible if we hope to use it to test a particle model of dark matter.

For these reasons, detecting and understanding dark matter both require detectors with extraordinarily low backgrounds and an accurate model of the backgrounds present in the experiment. As we saw in Figure 1.12, the sensitivity of an experiment to dark matter decreases as the uncertainty in the background model grows. While experiments typically use various counting techniques, some of which will be discussed in Chapter 3.2, to measure the radioisotopes present in many of their components, it is impractical to measure the activity of each component, down to every last screw and wire. Therefore, any Monte Carlo analysis of the expected background from such *ex situ* measurements will be inherently limited by the uncertainty in the backgrounds coming from the components that were not measured. One solution to this is to measure the backgrounds *in situ*, which can be accomplished by studying the γ -ray spectra observed in the detector, as will be discussed briefly in Chapter 3.2.

The signal produced by WIMP-nuclei interactions can generally be further verified by observations of annual modulation in the signal strength and the average direction that the WIMPs are coming from. These additional characteristics can provide strong evidence that an observed signal is, indeed, dark matter rather than an unforeseen background.

The key to detecting dark matter and finally understanding what it is therefore lies in the ability to design an experiment that will be able to efficiently reject background events, to build an experiment that is as clean as possible, and to model the experiment's backgrounds well enough to be confident that it will be background free.

1.4 Outline

Having motivated and explained the problems that this dissertation will try to solve in the quest of detecting and understanding dark matter, the rest of this document will focus on how the goals of modeling and vetoing backgrounds can be achieved. The primary focus of this dissertation will be on the DarkSide-50 WIMP detector, which will be described in more detail later in this document, though many of these details can be extrapolated to the bigger picture around developing the technology needed to model and eliminate backgrounds. The background that will be the primary focus here comes from neutrons, which, as will be discussed later, can most tenaciously mimic a WIMP in the detector.

Chapter 2 will go into detail of the design and operation of the LAr TPC of DarkSide-50. This discussion will provide context for the rest of the document, as it describes the part of the detector that is sensitive to WIMP recoils. We will discuss the design and composition of the detector, which is essential for construction a background model, and we will also discuss the physics underlying how the detector works.

Chapter 3 gives an overview of the various radioisotopes that are present in a dark matter experiment, where they come from, and what sort of signals they may produce. It then goes on to discussing the techniques with which we have measured the abundances of these isotopes in the various components of DarkSide-50, includ-

ing a summary of the concentration of various radioisotopes in most parts of the DarkSide-50 TPC.

Chapter 4 will review the various backgrounds that are relevant to DarkSide-50, as well as many other low-background detectors. We will discuss the physics underlying the processes that produce these backgrounds, with the goal of modeling their signals and understanding how we might eliminate them. We will also discuss processes that produce neutrons—especially (α, n) neutrons, which are the neutrons most likely to produce a WIMP-like signal in the detector. This discussion will include the introduction of a new tool developed here for calculating (α, n) yields.

Chapter 5 outlines the ideas behind a neutron veto system that is designed for efficiently eliminating neutron backgrounds in a low-background experiment. This design is ultimately what was used by the DarkSide-50 experiment, and so this implementation will be discussed in the most specific detail. In addition to discussing the signals that we expect to see in each detector, the physical processes that produce these signals, and how we can use these signals to veto neutron backgrounds, we discuss details of the optics and photonics and the ways in which they might affect the signals we observe.

Chapter 6 describes the process of prototyping and designing the DarkSide-50 neutron vetoes to ensure that they will be able to perform well enough to meet the experiment’s requirements. This process includes both research and development to address various design questions and building a prototype test stand to validate the final design before implementing it on a larger scale. We also discuss measurements taken to build a scintillator response model for predicting how the scintillator will respond to various types of particles depositing energy in it.

Chapter 7 discusses the optical simulations that were tuned on the prototype detector and then scaled to the full detector design to ensure adequate performance. These simulations were also used for better understanding the neutron veto’s optical response to various signals to help us understand signal seen in the actual neutron veto, as well as for building a model of how the detector works to help translate between data and physics simulations later on.

Chapter 8 goes into much more detail specific to the DarkSide-50 neutron veto system. It starts by discussing the low-level data acquisition and analysis framework used for measuring and evaluating signals from the neutron vetoes. We will then discuss the Water Cherenkov Veto (WCV) used as a muon veto in the detector, including the signals that we see from it and some discussion of its performance. Similar discussion will then describe the Liquid Scintillator Veto (LSV) used as a neutron veto, along with a detailed discussion of its performance and characterization.

Chapter 9 takes the discussion of the characterization of the neutron veto to the next level, where we discuss the calibration campaigns performed by actively inserting different neutron sources into the detector. This will include discussion of an early calibration campaign we performed using an $^{241}\text{AmBe}$ neutron source, as well as details of the development and analysis of an $^{241}\text{Am}^{13}\text{C}$ neutron source that more precisely fits the experiment’s needs. These calibration campaigns allowed us to further understand the response of the neutron veto to signals, which provided vital

input for modeling the behavior of the veto and understanding its neutron rejection efficiency.

Chapter 10 uses this model of the neutron veto's response function as well as a number of other tools to determine the flux of various backgrounds expected in DarkSide-50 and then, ultimately, to simulate the various backgrounds and determine the veto system's efficiency at removing these backgrounds, enabling DarkSide-50 to run background-free for three years. Most of this chapter will primarily be focused on (α, n) neutron backgrounds, which are typically among the hardest to veto. This chapter presents the heart of the background model and shows that this design can be used to drastically improve a low-background detector's sensitivity.

Chapter 11, finally, concludes the discussion of the background model and neutron veto system and draws upon the data taken by DarkSide-50 so far to show that we have achieved background-free conditions. We will also briefly discuss future directions and background source not elaborated upon in this document.

Lastly, though it was said in the acknowledgements, I wanted to reiterate that much of the work presented in this document was done in conjunction with a large set of collaborators, and I would not want to ignore their contributions. While I try to focus this document on the work that I performed myself, none of this would have been possible without the many people I worked with for each of these steps.

Chapter 2

The DarkSide-50 Detector

DarkSide-50 is a WIMP detector located at Laboratori Nazionali del Gran Sasso in central Italy. A drawing of the detector is shown in Figure 2.1. The experiment consists of three nested detectors.

The outermost detector is a stainless steel cylinder with a diameter of 11 m and a height of 10 m, referred to either as the muon veto or the Water Cherenkov Veto (WCV). This vessel is filled with ~ 1000 tonnes of water and is outfitted with an array of photomultiplier tubes (PMTs). This detector provides passive shielding to the inner detectors, and it can detect Cherenkov light produced by muons and their secondary particles, providing a veto signal for muons and any neutrons they may produce. This detector is discussed in more detail in Chapter 5.




Figure 2.1: A drawing of the DarkSide-50 detectors.




Figure 2.2: A cutaway view of the DarkSide-50 TPC.

Inside the WCV is the neutron veto, or Liquid Scintillator Veto (LSV). This detector is a 4 m diameter stainless steel sphere filled with ~ 30 tonnes of boron-loaded liquid scintillator and outfitted with an array of PMTs. The primary purposes of this detector are to provide prompt and delayed coincidence veto signals to help DarkSide-50 reject backgrounds—especially neutron backgrounds—and to provide *in situ* measurements of the backgrounds present in DarkSide-50. We will discuss this detector more in Chapter 5. Since, as we will discuss later, neutrons can produce signals that look identical to WIMPs in a dark matter detector and are therefore among the most concerning types of backgrounds, most of this document will focus on the design and performance of the LSV.

The innermost detector is the Time Project Chamber (TPC). This detector contains 50 kg of liquid argon (LAr) which is viewed by an array of PMTs. This is the detector with which we are looking for WIMP-nucleus interactions.

2.1 TPC Detector Design

Figure 2.2 shows a drawing of the DarkSide-50 TPC. This apparatus contains a total of ~ 150 kg of LAr, the inner ~ 50 kg of which is used to look for WIMP recoils. The innermost volume of liquid argon, referred to as the active volume, is in a cylindrical vessel surrounded by Teflon reflector around the sides, while fused silica windows

cover the top and bottom. An array of 38 photomultiplier tubes (PMTs) looks at the active volume through the fused silica windows. The inner surfaces of the fused silica windows and Teflon reflector are coated in TPB, a wavelength shifter that shifts the UV argon scintillation light into the visible spectrum so that the PMTs can detect it.

The entire volume is in an electric field, maintained by an array of copper field cage rings surrounding the Teflon, and by a conductive layer of ITO that coats both fused silica windows to hold them either at ground or high voltage, constituting the cathode and anode, respectively. There is also a thin layer of gaseous argon, supported by the “diving bell” structure in the anode fused silica, above the liquid argon volume. A mesh metal grid just below the gas-liquid interface holds a potential, allowing for a much stronger electric field between the grid and the anode than between the grid and the cathode. The field between the cathode and the grid is known as the drift field, and the field between the grid and the anode is the extraction field.

When a particle scatters in the liquid argon, it will produce fast scintillation light, and argon ions. Electrons from the argon ions are drifted to the gaseous layer, where the extraction field greatly accelerates them; the accelerated electrons then scintillate in the gaseous layer, producing a secondary burst of light. These two scintillation signals are called S1 and S2, respectively. The time profile of the S1 signal depends on whether the particle depositing energy in the argon scattered off of argon nuclei or off of argon electrons (known as nuclear and electronic recoils, respectively). This method is called pulse shape discrimination. Since we expect WIMPs to primarily scatter off of nuclei, this discrimination method is a very powerful tool for identifying electronic recoil backgrounds. The time difference between the S1 and S2 pulses as well as the distribution of light on the PMTs during the S2 pulse tell us the position of the recoil in the TPC, which allows us to make position cuts to remove nuclear recoils from surface backgrounds.

As a side note, liquid xenon TPCs, like LUX [83] and XENON100 [84], tend to use a very similar design (though without the TPB wavelength shifter). However, the S1 signal in liquid xenon does not have such powerful pulse shape discrimination, meaning that electron recoil backgrounds are more of a concern for these experiments. Instead, these experiments tend to rely on the ratio of S2/S1, which is another (albeit weaker) discriminant between electron and nuclear recoils. While liquid argon can typically suppress electron recoil backgrounds by up to a factor of 10^{10} [85] using pulse shape discrimination, liquid xenon TPCs typically expect to be able to suppress these backgrounds by closer to a factor of 200 using S1/S1 [86].

The DarkSide-50 TPC sits inside of a larger volume of liquid argon. This region is inside of two nested stainless steel cryostats separated by multilayer insulation.

2.2 TPC and Cryostat Composition

Since we will be focusing on the backgrounds in DarkSide-50, it is useful to understand where the origins of the backgrounds. To understand the source of the backgrounds, we must first discuss the composition of the detector itself. In this description, we will focus on the components in near proximity to the TPC, ignoring components of the




Figure 2.3: (*Left*) The DarkSide-50 outer cryostat as it was lowered into the LSV. (*Right*) The inner cryostat, removed from the outer cryostat, which contains the TPC. Photographs by Y. Suvarov.

LSV and WCV, because any backgrounds coming from them will likely be blocked by the large buffer of liquid scintillator and other materials they must pass through before they can reach the TPC.

The DarkSide-50 TPC is primarily contained in two nested stainless steel cryostats separated by a layer of Mylar multilayer insulation. The cryostats are made of 304L stainless steel, and have an approximately cylindrical body with a curved bottom dome and removable top dome that connects to the body around a flange via a series of nuts, bolts, and washers (with a total of 80 of each between the flanges of the inner and outer cryostats). The purpose of the cryostats is to thermally and hermetically isolate the TPC from the surrounding detector. A series of feedthroughs connect to organ pipes in order to fill and empty the TPC and supply the high voltage and readout system for the TPC PMTs and field cage rings. The cylindrical cryostats and top and bottom domes were fabricated out of sheets of 5 mm stainless steel, while the flanges were cut out of a different stainless steel sample with an original thickness of 25 mm. This geometry is illustrated in Figure 2.3.

In order to form a vacuum seal between the cryostats, a titanium o-ring was used inside the flanges of the inner cryostat, while a Viton o-ring was used inside the flanges of the outer cryostat. The Viton o-ring can be seen in Figure 2.4.

As mentioned above, Mylar multilayer insulation is used to thermally separate the inner and outer cryostats. Mylar is a highly reflective, thin material with a high ther-



Figure 2.4: (*Left*) The inner cryostat top dome covered in embossed multilayer insulation with the outer cryostat dome above it. (*Right*) The inside of the outer cryostat with netted multilayer insulation along the walls and a Viton o-ring in the flange. Photographs by Y. Suvarov.




Figure 2.5: (*Left*) The cryostat opened up to reveal the TPC. (*Right*) Top-down view of the TPC. Photographs by Y. Suvarov.

mal resistance, especially when several non-touching layers are used to separate two materials. A thin ~ 100 nm aluminum coating on the Mylar enhances the reflectivity to suppress radiative heat transfer. There are 24 and 25 layers of Mylar around the inner cryostat flange and cylindrical body, respectively, with a layer of polyethylene netting separating each layer. Additionally, 30 layers of embossed Mylar separate the top and bottom domes of the cryostats. Both kinds of Mylar sheets can be seen in Figure 2.4.

Figure 2.5 shows the structure of the TPC. Two different kinds of Teflon were used: the structural supports of the TPC and the PMT mounts were made of a Teflon manufactured by Industrial Fluoro-Plastics (IFP), while the Teflon cylinder used for the reflector was made by Applied Plastics Technology (APT). Not shown in this figure is a fused silica disk at the bottom of the TPC; this was coated in a conductive layer of Indium Tin Oxide (ITO) to form the TPC cathode, while the top was shaped like a cylindrical disk with a smaller, shallow cylinder carved out of the bottom to




Figure 2.6: (*Left*) The grid that maintains the extraction field in the TPC. (*Right*) TPC field cage rings and resistors. Photographs by Y. Suvarov.

form the diving bell. The bottom surface of the diving bell was also coated with ITO to make the anode of the TPC. The shape of the diving bell allows it to hold a gas pocket, forming the layer of argon gas at the top of the TPC. A structure was molded out of the fused silica in the diving bell to make the bubbler, which is used to maintain the gas layer at approximately constant size. The inner surfaces of the APT Teflon reflector and fused silica cathode and diving bell were also coated with the wavelength shifter TPB in order shift the UV argon scintillation light into the visible range.

The volume inside the reflector and fused silica windows is filled with liquid argon, and is the sensitive volume of the experiment.

When a particle scatters in the TPC it will produce a burst of light (called S1) and ionize some number of electrons. These electrons are drifted up to the top of the TPC via the drift field, which is maintained by a set of 16 copper field cage rings at different voltages. The voltage of each field cage ring is set by a series of resistors connecting each of them, for a total of 64 resistors. A stainless steel grid rests in a 304L stainless steel support ring near the top of the TPC, just below the layer of gaseous argon maintained by the diving bell. This grid is held at a constant voltage and is used to maintain a much higher electric field between itself and the ITO coating on the bottom of the diving bell. This stronger electric field is called the extraction field. When electrons reach the extraction field, they are accelerated to a very high




Figure 2.7: The top array of PMTs in their Teflon support mount, with circuitry and cold amplifiers visible. Photographs by P. Meyers.

speed and produce the S2 signal. The field cage rings and grid that maintain this electric field are shown in Figure 2.6.

Light from both scintillation pulses is collected by two arrays of PMTs that sit at the top and bottom of the TPC, just beyond the fused silica windows. These PMTs have a fused silica faceplate and a Kovar shell around them. A Cirlex circuit board containing 24 resistors and 5 capacitors operates each PMT. A specially designed amplifier that can operate at liquid argon temperatures was designed by A. Razeto, and consists of a Cirlex circuit board and additional circuitry. This cold amplifier facilitates readout of the analog signal from the PMTs via a set of male and female MCX connectors, which are primarily made of a Beryllium-Copper alloy. Each PMT is held into a Teflon support structure by three copper screws loaded with one brass spring each.

Tables 2.1 and 2.2 summarize all of the components in the DarkSide-50 TPC, along with how much there is (either the mass in grams or kilograms or the number of pieces) and the material composition. Many of the masses reported here were determined by P. Meyers or were measured by M. Laubenstein when he measured the radioactivity of each component.

While some materials are labeled as being made of 304L stainless steel, we do not know the kind of steel used for some of the others, though it is likely that these were made with either 304L or 316 stainless steel. The main difference between 304L and 316 stainless steels is that 304L has an additional 2–3% composition of molybdenum.

Table 2.1: Summary of DarkSide-50 TPC components.

Component	Amount	Material	Composition (% mass)						
Cryostat Body	142±1 kg	304L Stainless Steel (5 mm)	C 0.04 N 0.1	Cr 18 Fe 70.9	Mn 2	Ni 8	P 0.05	S 0.03	Si 1
Cryostat Flanges	30.0±0.3 kg	Stainless Steel (25 mm)	C 0.04 N 0.1	Cr 18 Fe 70.9	Mn 2	Ni 8	P 0.05	S 0.03	Si 1
Organ Pipe Tubes	1.92±0.02 kg	Stainless Steel	C 0.04 N 0.1	Cr 18 Fe 70.9	Mn 2	Ni 8	P 0.05	S 0.03	Si 1
Cryostat Bolts	2.92±0.03 kg	Stainless Steel	C 0.04 N 0.1	Cr 18 Fe 70.9	Mn 2	Ni 8	P 0.05	S 0.03	Si 1
Cryostat Nuts	826±8 g	Stainless Steel	C 0.04 N 0.1	Cr 18 Fe 70.9	Mn 2	Ni 8	P 0.05	S 0.03	Si 1
Cryostat Washers	468±5 g	Stainless Steel	C 0.04 N 0.1	Cr 18 Fe 70.9	Mn 2	Ni 8	P 0.05	S 0.03	Si 1
Outer Cryostat O-ring	1 pc	Viton	H 0.9	C 28.1	F 71.0				
Multilayer Insulation	740±120 g	Mylar	H 4.2	C 62.5	O 33.3				
MLI Netting	1.04±0.01 kg	Polyethylene	H 4.2	C 62.5	O 33.3				
MLI Al Layer	22.3±0.1 g	Aluminum	Al 100						
Liquid Argon	153±3 kg	Argon	Ar 100						
Teflon Support Structures	35.88 kg	IFP Teflon	C 24.0	F 76.0					

Table 2.2: Summary of DarkSide-50 TPC components.

Component	Amount .	Material	Composition (% mass)							
Teflon Reflector	22.34 kg	APT Teflon	C 24.0	F 76.0						
Cathode Window	1.977 kg	Fused Silica	Si 46.7	O 53.3						
Diving Bell	1.194 kg	Fused Silica	Si 46.7	O 53.3						
Field Cage Rings	16.5 kg	Copper	Cu 100							
Field Cage Resistors	64 pc	Alumina	Al 52.9	O 47.1						
Extraction Grid	1.9 g	Stainless Steel	C 0.04 N 0.1	Cr 18 Fe 70.9	Mn 2	Ni 8	P 0.05	S 0.03	Si 1	
Grid Support Ring	461 g	304L Stainless Steel (5 mm)	C 0.04 N 0.1	Cr 18 Fe 70.9	Mn 2	Ni 8	P 0.05	S 0.03	Si 1	
PMT Shell	5.31 kg	Kovar	Fe 52.5 Mg 0.1	Ni 29 Zr 0.1	Co 17 Ti 0.1	Mn 0.3 Cu 0.2	Si 0.2 Cr 0.2	C 0.02 Mo 0.2	Al 0.1	
PMT Support Screws	887 g	Copper	Cu 100							
PMT Stems	38 pc	Borosilicate Glass	Si 31.3	Al 3.2	B 5.6	Li 0.5	Na 4.6	Ba 1.8	O 53.2	
PMT Circuit Board	38 pc	Cirlex	H 2.6	C 69.1	N 7.3	O 20.9				
PMT Resistors	3.91 g	Alumina	Al 52.9	O 47.1						
PMT Capacitors	3.21 g	Alumina	Al 52.9	O 47.1						
PMT Barrel Connectors	108.5 g	Brass	Cu 63	Zn 37						
Cold Amp Circuit Board	18 pc	Cirlex	H 2.6	C 69.1	N 7.3	O 20.9				
Cold Amp Circuitry	8.17 g	Alumina	Al 52.9	O 47.1						
Male MCX Connectors	40.3 g	Be-Cu Alloy	Be 0.4	Ni 1.8	Cu 97.8					
Female MCX Connectors	29.6 g	Be-Cu Alloy	Be 0.4	Ni 1.8	Cu 97.8					
Solder	50±25 g	Solder	Sn 96.5	Ag 3.0	Cu 0.5					

Since molybdenum has a very low (α, n) yield, and the uncertainty in the amount of iron is much greater than 2–3%, we assume here that all stainless steel samples are made of 304L stainless steel.

In general, the specifications for 304L stainless steel composition show small amounts of variance across different companies’ reports. The numbers reported in Tables 2.1 and 2.2 mostly come from [87], using “typical” values consistent with numbers cited in other sources. The concentrations of carbon, manganese, phosphorous, sulfur, silicon, and nitrogen are all upper limits, while the concentrations of the others vary by up to $\sim 5\%$.

We also assumed that the PMT barrel connectors are made entirely of brass and that the MCX connectors are made entirely of high conductivity beryllium-copper alloy. We used similarly “typical” values for their chemical compositions. We also model the PMT resistors and capacitors and all of the cold amplifier circuit elements as being made entirely of alumina. While these assumptions introduce uncertainty into the (α, n) yield estimations, they err on the side of overestimating the abundance of the nuclei with the highest (α, n) yields. The (α, n) yields that we calculate for these materials are therefore conservative.

The uncertainty in the mass reported for stainless steel components, as calculated by P. Meyers, is around 1% for all components.

For the multilayer insulation, we assume an uncertainty of $\sim 1\%$ on the mass; [88] does not report uncertainty in the thickness of the mylar, but since these materials are made using precision techniques, we assume the uncertainty in the thickness is small compared to the uncertainty in the area. The sheets used were measured by A. Fan to ~ 1 cm, implying an uncertainty of $\sim 1\%$ on the area of each sheet. We estimate a similar uncertainty on the mass of the aluminum layer, which has the same surface area as the mylar and is also made using precision techniques.

The uncertainty in the mass of the liquid argon was reported by H. Beck in his report on the argon distillation.

The total solder mass was estimated by weighing two PMT bases and subtracting the mass of the circuit board, resistors, and connectors. We then assumed the mass of the solder was proportional to the number of electronic components and scaled to determine the mass of solder in the cold amplifiers. Overall, we will assume an uncertainty of 50% in the mass of the solder.

All components measured in pieces rather than weight (grams or kilograms) are assumed to have no uncertainty, since we can count how many of these components are in the detector. For all other components, we assume that the uncertainty is small compared to the precision with which the mass is reported. As we will see in Chapter 3.2, the uncertainty in the contamination of each of these components will be large compared to the uncertainty in their masses.

2.3 LAr Scintillation

The process of LAr scintillation is described in great detail in [89, 90]. When a particle scatters in liquid argon, it will either scatter off of an electron orbiting an argon atom


or off of the argon nucleus. In either case, a charged electron or nucleus will be sent recoiling through the liquid argon, or the atom will be left in an excited state.

As the charged particle moves through the argon, it will continually lose energy, either through ionization or excitation of argon electrons, or by scattering off of more argon nuclei, causing them to recoil as well. In this latter case, some of the energy given to argon nuclei will end up being transferred to electrons (directly by the recoiling nucleus or by a nucleus sent recoiling by this nucleus, and so on in a cascade of recoils) while the remainder will ultimately be dissipated as heat.

When an electron moves through a material, virtually all of its energy will be lost to other electrons; kinematic constraints severely limit the amount of energy that an electron can transfer to a nucleus in a single recoil. For a nucleus moving through a homogeneous material, the energy lost directly to electrons and the energy lost directly to nuclei are described by Lindhard theory [91].

Figure 2.8 illustrates the process of liquid argon scintillation. An incident particle (such as a WIMP or a neutron) scatters off of an argon nucleus. As the recoiling argon nucleus slows to a stop, it will lose some fraction $(1 - f)$ of its energy by recoiling off of other argon nuclei; some of these recoiling nuclei will have too little energy to scintillate and dissipate their kinetic energy as heat, while others will slow down and potentially scintillate through the same process as the original recoil nucleus. The remaining fraction f of the argon nucleus's energy will go into ionizing and exciting nearby argon atoms, forming N_{ion} ions and αN_{ion} excitons, where α is the ratio of excitons to ions.

The argon exciton has one valence electron promoted to the first excited orbital. This vacancy in the highest orbital allows the exciton to dimerize with a nearby ground state argon atom; the valence electron of the ground state atom and the excited electron of the exciton become bonding pairs in the Ar_2 dimer. These excimers then decay to the ground state and release a photon. This process, known as exciton self-trapping, takes place on the order of picoseconds and is described by



where $h\nu$ represents a 128 nm ultraviolet photon.

Alternatively, an argon ion can also form a dimer with a ground state argon atom with one of the atom's valence electrons being shared. A free electron, typically ionized from a nearby argon atom by the same charged projectile, can then be captured by the charged dimer. This process is known as recombination and splits the dimer into a doubly excited atom and a ground state atom. This doubly excited atom will then non-radiatively decay to a single excited state and then decay through exciton

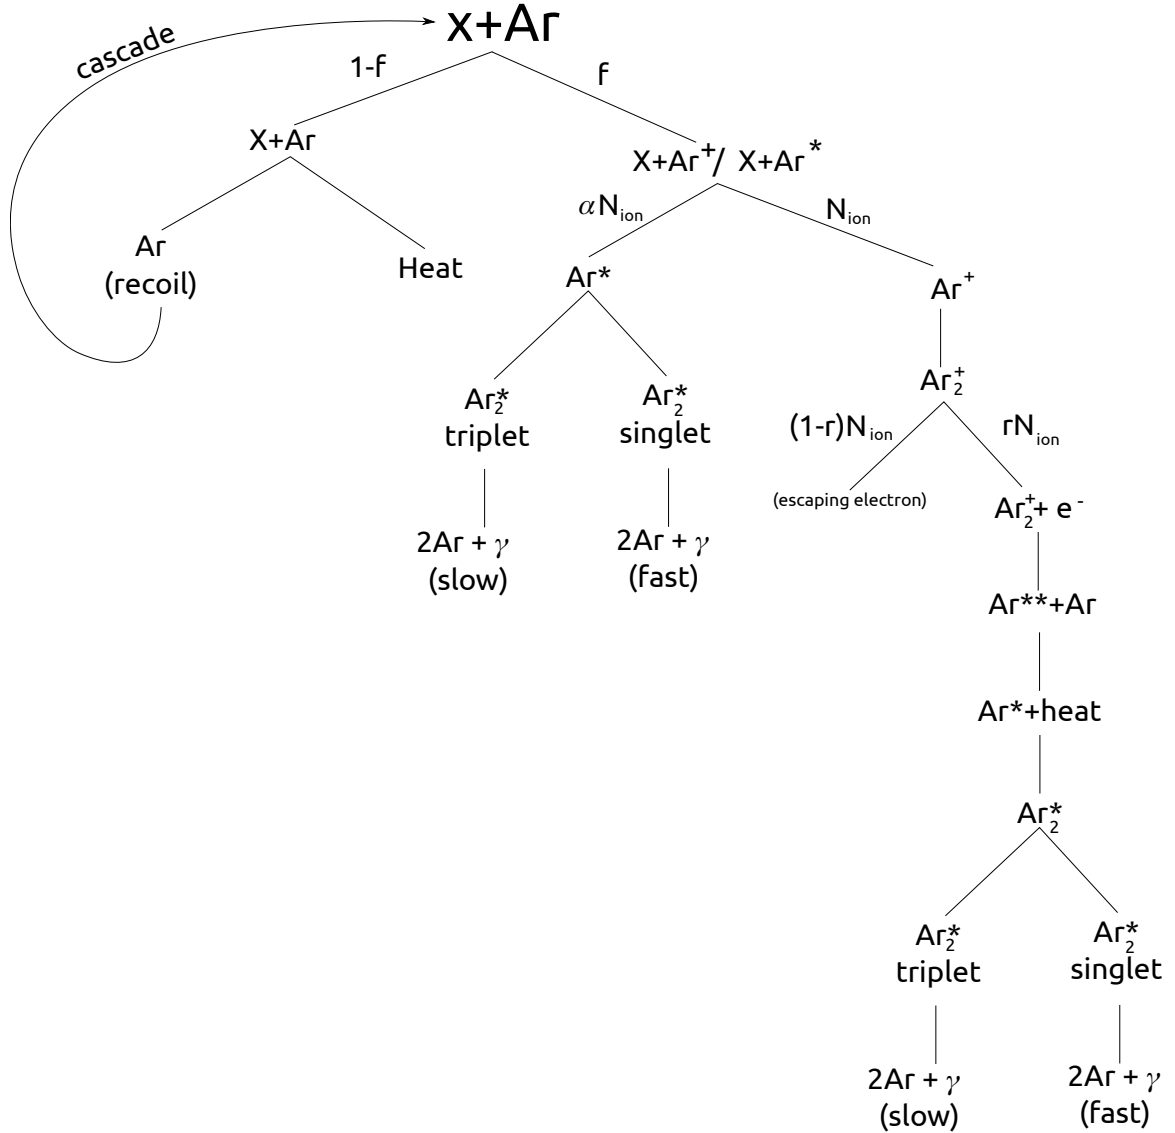



Figure 2.8: Diagram showing the process of liquid argon scintillation after a particle X has scatters on an argon nucleus. f is the fraction of energy that the recoiling argon nucleus loses to electrons; $(1-f)$ is the fraction lost to other argon nuclear recoils. N_{ex} and N_{ion} are the number of excitons and the number of ions produced, respectively, and $\alpha = N_{\text{ex}}/N_{\text{ion}}$. r is the recombination rate for argon ions capturing an ionized electron. Electrons that fail to recombine escape, and can either be collected as an S2 signal or be absorbed by another atom (such as an impurity), or they may be recaptured by another ion a long time later.

self-trapping:



The probability of an electron capturing on the Ar_2^+ dimer is referred to as the recombination probability, and varies with the density of ions in the LAr [89]. Generally, in the absence of an electric field, this probability will be a little less than 1. However, as was explored by [92], this probability goes down in the presence of an electric field and becomes smaller at stronger fields, as more electrons are drifted away before they can recombine.

The singly excited dimer produced by these processes is in a Rydberg state, where the Ar_2^+ core is orbited by a bound electron. The core and the bound electron each have a spin of magnitude 1/2; however, since the directions of the spins relative to each other can vary, we find the following four possible spin configurations

$$\begin{aligned}
&\left. \frac{1}{\sqrt{2}} (|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle) \right\} \quad \text{singlet} \\
&\left. \begin{array}{l} |\uparrow\uparrow\rangle \\ \frac{1}{\sqrt{2}} (|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle) \\ |\downarrow\downarrow\rangle \end{array} \right\} \quad \text{triplet}
\end{aligned}$$


Since the ground state, 2Ar , and the singlet state both have total spins of 0, the transition from the singlet state to the ground state is allowed by all selection rules and therefore happens on a very short timescale of 6 ns. The decay of the triplet state, however, is forbidden by conservation of angular momentum due to the state's total spin of 1, and therefore occurs on the much slower timescale of $\sim 1.5 \mu\text{s}$ [50]. When these dimers split apart, they release a 128 nm photon.

Since excitons and ions both form excimers through different mechanisms—in particular, the excited electron from the exciton channel necessarily has the same spin as the promoted electron, while the recombination electron will have a spin largely uncorrelated with that of the ionized electron—both channels form triplets and singlets with different probabilities. The exciton-to-ion ratio in LAr is generally assumed to be ~ 0.21 for electron recoils [93], and ~ 1 for nuclear recoils [92].

Since electron recoils and nuclear recoils lose energy to atomic excitation and ionization through different mechanisms, they result in different exciton-to-ion ratios. The non-unitary recombination probability also decreases the number of ions that eventually produce photons; since the recombination probability increases with the ionization density, the singlet-to-triplet ratio can also vary with the ionization density of the particle's track. This effect becomes more pronounced in the presence of an

electric field and tends to make the singlet-to-triplet ratios for electron and nuclear recoils more similar than they are without an electric field.

Scintillation light can also be suppressed by competing processes that allow argon excitons to decay non-radiatively. This can happen through processes such as biexcitonic collisions, photo-ionization, and the Penning process [89]. The rate of these processes is proportional to the square of the exciton density, and therefore suppresses more scintillation at higher densities. Biexcitonic quenching, photo-ionization, and the Penning process allow argon excitons or excimers to non-radiatively de-excite through the following three mechanisms, respectively:



Alternatively, a free electron can interact with the bound electron in the Ar_2^* excimer through a spin-exchange interaction. This process can convert a triplet dimer into a singlet dimer or vice versa, and is described in detail in [94].

These interactions all occur at rates proportional to the ionization density of the projectile electron or nucleus, and therefore have bigger effects for projectiles with high stopping powers (*e.g.* ^{40}Ar nuclei) than for those with lower stopping powers. The kinematics of these reactions means that for a given ionization density, there is an approximately constant probability per unit time of these reactions occurring. Since some of these reactions happen to argon dimers, this means they will disproportionately suppress the triplet state scintillation over the singlet state scintillation. These effects are described in detail in [94], which includes a list of rate constants for each of these processes.

The result of these effects, along with the different exciton-to-ion ratios for electron and nuclear recoils, is that a higher fraction of the scintillation light produced by nuclear recoils, with their higher stopping powers, will come from singlet state dimers compared to electron recoils. This difference means that the scintillation pulse produced by nuclear recoils will be much faster than that of electron recoils. At high stopping powers, this difference is predominantly driven by spin-exchange interactions [94].

By looking at the time profile of LAr scintillation pulses, we can therefore distinguish between electron and nuclear recoils. A common parameterization of this is f_{90} , the fraction of the scintillation light that occurred in the first 90 ns of the scintillation pulse. By 90 ns, virtually all of the singlet dimers will have decayed, while very few of the triplets will have done so, motivating the use of this parameter. This technique is generally called Pulse Shape Discrimination, or PSD.

Since the presence of an electric field can reduce the recombination rate of argon ions, the drift field in a TPC suppresses the total amount of light produced and decreases the rejection power of PSD. Nevertheless, measurements made in DarkSide-50 with 200 V/cm found that an f_{90} cut with a 90% nuclear recoil acceptance had a rejection power $> 1.5 \times 10^7$ for ^{39}Ar decays in the 8.6–65.6 keV energy range. This limit on

the PSD rejection power is limited by the statistics measured in [50]. Measurements reported in [85] estimate a rejection power $> 10^{10}$ for electron recoils above 20 keV without an electric field.

The distribution of energy lost directly to argon electrons for a charged particle in LAr compared to the energy lost to nuclei is generally described by Lindhard theory [91]. The fraction of the projectile's energy that goes to electrons is given by f in Figure 2.8.

In general, the stopping power of a projectile in a medium can be expressed as the sum of the electronic stopping power and the nuclear stopping power $S = S_e + S_n$, where the electronic stopping power is the energy lost to atomic electrons in the medium, and the nuclear stopping power is the energy lost to nuclei.

At low energies, the energy that an ion transfers to electrons is proportional to the momentum transferred in a collision. As a result, the electronic stopping power is generally proportional to the velocity of the projectile. To normalize for different sets of projectile and target nuclei, Lindhard and Scharff use dimensionless range and energy variables ρ and ϵ , defined by

$$\begin{aligned}\rho &= r4\pi\eta M_2 a^2 \frac{M_1}{(M_1 + M_2)^2} \\ \epsilon &= T \frac{aM_2}{Z_1 Z_2 k q_e^2 (M_1 + M_2)}\end{aligned}\tag{2.1}$$

where r is the ion's range, η is the number density of the target atoms, T is the kinetic energy of the projectile ion, k is Coulomb's constant, q_e the charge of an electron, M_1 and M_2 are the masses of the projectile and target nuclei, respectively, and $a = 0.8853a_0 \left(Z_1^{2/3} + Z_2^{2/3}\right)^{-1/2}$ is the screened potential radius with Bohr radius a_0 and Z_1 and Z_2 the atomic numbers of the projectile and target, respectively.

Lindhard et al. [91] then write the electronic stopping power for low energy ions as

$$S_e = \kappa \epsilon^{1/2}\tag{2.2}$$

where

$$\kappa = \xi_e \frac{0.0793 Z_1^{1/2} Z_2^{1/2} (A_1 + A_2)^{3/2}}{\left(Z_1^{2/3} + Z_2^{2/3}\right)^{3/4} A_1^{3/2} A_2^{1/2}}$$

where $\xi_e \approx Z_1^{1/6}$ and A_1 and A_2 are the mass numbers of the projectile and target nuclei, respectively.

At low energies, energy lost to nuclear recoils is generally described by hard-sphere scattering, and so a first order approximation can be made in which the nuclear stopping power is constant with respect to the ion energy

$$S_n \approx \left(\frac{\pi^2}{e}\right) \frac{q_e^2 a_0 Z_1 Z_2 M_1}{Z_1^{1/3} (M_1 + M_2)}\tag{2.3}$$

A more complete description of the nuclear stopping power is given in [91].

In general, Lindhard et al. report that this theory tends to agree with data to within $\sim 20\%$, though the accuracy breaks down at very low energies where $\epsilon \lesssim 0.01$. Notably, this model assumes that ions interact with a free gas of electrons, and does not account for atomic excitations and ionization binding energies that may affect these interactions. Sorensen [95] has shown that some of the disagreement between Lindhard’s model and data at low energies can be reconciled by accounting for the electrons’ binding energy.

Nevertheless, the model described by Lindhard et al. performs well for $\epsilon \gtrsim 0.01$, and so the equations presented in [91] remain the commonly used standards.

When describing noble liquid scintillation, we often care about the total amount of energy that is eventually dissipated through electronic ionization and excitation, as this is the energy that is free to scintillate. The parameter f in Figure 2.8 describes the fraction of energy that goes directly to these electronic processes; however, a full description of the amount of scintillation light produced by an ion must account for the cascade as well.

The total fraction of deposited energy that is ultimately dissipated by electron excitation and ionization after accounting for the full cascade is referred to as the Lindhard factor, f_L . When the projectile and target nuclei are the same (*e.g.* an ^{40}Ar nucleus recoiling in LAr), the effects of the cascade can be estimated by the semi-empirical equation [91]

$$f_L = \frac{\kappa g(\epsilon)}{1 + \kappa g(\epsilon)} \quad (2.4)$$

where $g(\epsilon) = 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon$.

However, when the projectile nucleus is different from the target material, a simulation must be performed to account for the full cascade. J. Ziegler’s tool “Stopping and Range of Ions in Matter” (SRIM) [96] uses this theory, with corrections made to account for the target material’s electron density distribution based on Hartree-Fock theory, to simulate ions slowing down in a material. This tool calculates the effects of the full cascade and estimates the effect of the projectile ion on the target material, including the total amount of energy dissipated as heat and the amount of energy ultimately lost through electron ionization. Since this model is based on Lindhard theory and it does not account for electron binding energies, the accuracy of these calculations also breaks down at low projectile energies.

Finally, a complete description of LAr scintillation must also include light lost through bi-excitonic quenching processes. These processes scale with the square of the projectile’s stopping power, and are therefore described well by Birks’s saturation law, presented in [97]. While this saturation law was developed to describe nonlinearities in organic liquid scintillator scintillation by describing the probability of a molecule being “damaged” and the “damaged” molecule non-radiatively dissipating excitation energy, it has been shown to describe the saturation of noble liquids well, too [92]. This saturation law is given by

$$f_B = \frac{1}{1 + kB \frac{dE}{dx}} \quad (2.5)$$




Figure 2.9: Drawing of Hamamatsu R11065 PMT detecting a photon.

where kB is Birks' constant, which parameterizes the saturation effects.

The SCENE collaboration has shown that a Lindhard-Birks factor given by $f_{LB} = f_L \times f_B$ with $kB = (5.0 \pm 0.2) \times 10^{-4} \text{ MeV}^{-1} \text{ g cm}^{-2}$ describes the scintillation of neutron- ^{40}Ar recoils well. The ratio of the amount of scintillation light produced by a projectile particle in the medium to the amount of light one would expect of $f_{LB} = 1$ is generally called the quenching factor and is given by $QF(T)$ where T is the kinetic energy of the projectile.

The quenching factor generically describes the nonlinearity of a scintillator's response. In the case of a noble liquid scintillator, $QF \approx f_{LB}$. However, for different scintillators, we often need to model the quenching differently.

Since the nonlinearity of a scintillator's response varies with the energy and identity of the incident particle, a common convention is to normalize the scintillator's response to that of an electron. When discussing scintillation signal strengths, it is therefore common to use units of "keV electron equivalent", or keV_{ee} which express the energy an electron would need to have to produce the same amount of scintillation light. Generally electrons are quenched very little, so the ratio between an energy measurement in units of keV and units of keV_{ee} is often close to the quenching factor.

2.4 Photomultiplier Tubes

The DarkSide-50 TPC uses Hamamatsu R11065 PMTs to detect photons in the visual spectrum. A drawing of such a PMT is shown in Figure 2.9. These phototubes have a flat 3 inch fused silica faceplate with a very thin 2.5 inch diameter bialkali photocathode evaporated on the back of the faceplate.

When an incident photon hits the photocathode, it can eject an electron from the photocathode through the photoelectric effect. A constant potential held between the photocathode and the anode plate causes the photoelectron (PE) to accelerate as it approaches the anode. When the electron strikes the anode, it can then ionize




Figure 2.10: (*Black*) The SPE charge distribution readout from a Hamamatsu R11065 PMT. (*Red*) A Gaussian distribution plus an exponential is fit to this curve above the pedestal, taken from [13].

multiple electrons. These electrons are then accelerated to the first dynode, where they may each ionize multiple electrons, and so on, with a steadily increasing potential accelerating electrons between each dynode in the chain.

This cascade causes the number of electrons to grow exponentially with each dynode in the so-called “avalanche”, until finally the current is collected as it flows through the readout cable. This signal is then digitized by an analog-to-digital-converter (ADC) and recorded to disk as some number of ADC counts per unit sample time. The total number of electrons finally collected from a single photoelectron is the gain of the PMT, and depends on the high voltage supplied to the PMT.

When discussing PMT signals, it is common to refer to the total charge collected. This is the integral of the current read out by the electronics over the duration of the pulse, and is a measure of the total charge of all of the electrons finally collected. Since the total gain of a PMT may vary with the voltage or with other properties of the PMT that may change over time, it is common to normalize the charge collected by the charge of a single photoelectron (SPE). Doing so allows one to measure the charge collected by a PMT in terms of the number of photoelectrons it detected, which is proportional to the number of photons that hit the photocathode.

An example of a photoelectron charge distribution is shown in Figure 2.10. The peak around 0 ADC counts is the pedestal and is a measure of the noise in the PMT and electronics—it is the amount of charge recorded when no photoelectrons were produced in the PMT. Above the pedestal, the SPE charge distribution can be described well by a Gaussian distribution with a small exponential component at lower charges. The Gaussian distribution is a result of the mean gain of each dynode, which may vary as a result of fluctuations in the incident electrons’ kinetic momentum and thermodynamic effects. The nature of the exponential tail is not as




Figure 2.11: The Hamamatsu R11065 quantum efficiency as a function of incident photon wavelength, as reported by Hamamatsu [14].

well understood, though one explanation is that it may simply be a mathematical artifact of the multiplication statistics.

When a photon reaches the photocathode, the photon may pass through the photocathode without producing a PE if the mean free path of the photon in the photocathode material is long compared to the thickness of the photocathode. Alternatively, if the mean free path is too short, the PE will be produced near the outer surface of the photocathode and have to traverse the entire thickness of the photocathode before reaching the anode; in this case the PE may be reabsorbed by the photocathode and therefore go undetected. The result of these two effects is that the probability of a photon producing a PE varies with the photon's wavelength; this probability is referred to as the quantum efficiency of the PMT. The quantum efficiency of the Hamamatsu R11065 PMTs as reported in [14] is shown in Figure 2.11.

Additionally, if the PE fails to hit the first dynode in the chain, the PMT can also fail to produce a signal. The probability of the PE hitting the first dynode in the chain is called the collection efficiency of the PMT and has been reported by Hamamatsu to be $\sim 85\%$.

If there are any impurities inside the PMT, either free gases or impurities adsorbed onto the component surfaces, they may be ionized by the ionized electrons in a signal. If this happens, the freed ions may drift towards the photocathode. The time it takes for the ion to reach the photocathode will vary with its mass and its location inside the PMT; each impurity may therefore reach the photocathode with a characteristic time delay after the initial signal, and they may produce signals ranging from 1 to several PE. These delayed signals are referred to as afterpulses and generally happen on timescales ranging from 100 ns to $\sim 20 \mu\text{s}$.

Chapter 3

Radioactive Contaminants

Now that we have established the importance of minimizing backgrounds in a dark matter experiment and discussed the particular context of the DarkSide-50 detector, it makes sense to next consider the actual backgrounds that we expect to be present.

Backgrounds can typically be categorized as either external or internal. External backgrounds are those coming from the environment. These can include radioactivity naturally present in the detector’s laboratory setting, such as radioactivity in the rocks, or cosmogenic backgrounds coming from the atmosphere. Typically the latter backgrounds are greatly reduced by going far underground and using the earth as shielding. Internal backgrounds are those coming from the detector components themselves. Most things, either through exposure to the air or inherent to how they were produced, contain trace amounts of radioactive contamination. While these contaminant levels are low on an absolute scale, the fact that they are right next to the detector, past all of its shielding, means that they can still pose a hazard to the experiment.

This chapter will begin by discussing the various radioisotopes that are typically present, along with where they come from and what sort of signals they may produce in the detector. We will then discuss their presence in DarkSide-50 as determined by radioisotope counting results that we have obtained, in order to see what these background rates will likely look like and understand their origin more fully.

3.1 Radioisotopes of Interest

There are many radioisotopes that can be found in materials commonly used in dark matter detectors. While a comprehensive list of every possible isotope would get impractically long and include many isotopes that end up being irrelevant for any given experiment, it is worth discussing the radioisotopes that appear in DarkSide-50 and produce measurable signals.

The radioisotopes we will discuss here include ^{39}Ar and ^{85}Kr , which are intrinsically present in the LAr volume. Contamination from ^{39}Ar can be reduced by using argon extracted from underground sources rather than the atmosphere, while ^{85}Kr backgrounds must be distilled out of the ^{39}Ar sample. The remaining signal from

these backgrounds can be very highly suppressed using PSD. These backgrounds both β -decay and produce signals only seen in the TPC.

We will also discuss ^{60}Co and ^{40}K , which are typically found in the detector structural materials and PMTs. In DarkSide-50, most of the ^{60}Co and ^{40}K can be found in the stainless steel of the cryostats or in the borosilicate glass of the Hamamatsu R11065 PMTs used in the detector. These isotopes β -decay; however, the β s from these materials are unlikely to escape the materials to produce a directly visible signal. Instead, these isotopes tend to produce high energy γ -rays. These isotopes therefore can produce electron recoil signals in both the TPC and the LSV. The contamination from these isotopes is primarily reduced by measuring the radioactivity of the materials used to make the various parts of DarkSide-50 and choosing components with little contamination. The electron recoils produced by these isotopes can then be heavily suppressed by PSD.

^{14}C and ^3H are intrinsically present in the organic liquid scintillator used in the LSV. Since they β -decay without producing any γ -rays, they do not produce any signal in the TPC. However, their signals can clearly be seen in the LSV as random coincidences, and at too high a rate they may reduce the LSV's capacity for detecting true coincidences. The contamination from these isotopes can be reduced by deriving the compounds used in the scintillator from petroleum rather than plant methanol.

^{232}Th , ^{238}U , and ^{235}U are each at the start of different decay chains, and are primarily found in the stainless steel of the DarkSide-50 cryostats and in the borosilicate glass of the PMTs. Much like ^{60}Co and ^{40}K most of the β - and α -emitters in these decay chains are unlikely to produce signals directly in the TPC or LSV. However, many of these decays produce γ -rays, which may be seen in both detectors. Additionally, some concentration of radon daughters in these chains may be deposited on the inner surfaces of the detectors (*i.e.* the Teflon reflector and the fused silica windows). These contaminants can constitute surface backgrounds; the most problematic of which are the α -emitters, which may decay and send a nucleus recoiling into the LAr. These nuclear recoils can be suppressed by making position cuts in the TPC to remove events near the surfaces, and these backgrounds can be reduced by avoiding radon exposure for the various components. This is accomplished in DarkSide-50 through the use of a radon-suppressed clean room, discussed in [50]. ^{232}Th , ^{238}U , and ^{235}U can also produce neutron backgrounds, either through spontaneous fission processes or through the (α, n) reaction when α -emitters in these decay chains decay and react with light elements in the various materials. These neutron backgrounds can generally be suppressed by making multiple scattering cuts in the TPC or by looking for a veto signal in the LSV. In general, backgrounds from these decay chains can also be reduced by carefully measuring their contamination in the materials used for various components of the detector and choosing materials with low levels of contamination.

3.1.1 Argon-39

The most prominent background for DarkSide-50, as well as for any other liquid argon-based experiment, comes from ^{39}Ar , which β -decays with a half-life of 269 years and an endpoint energy of 565 keV [15]. Since this background is intrinsic to the argon

itself, ^{39}Ar decays within the detector’s sensitive argon volume will deposit all of their energy into the argon. While this β spectrum covers most of the WIMP energy region of interest, it can very effectively be eliminated through PSD, as discussed in Chapter 2.3.


Another danger of ^{39}Ar backgrounds is that too high of an ^{39}Ar concentration in the detector may result in pileup of ^{39}Ar decays. This can drastically reduce the live time of the experiment by constantly triggering the detector on these kinds of decays, and it can interfere with the detector’s ability to reconstruct other kinds of events. This is a bigger problem for two phase detectors like DarkSide-50 that see two pulses for every event, often separated by many microseconds, and therefore need longer acquisition windows. This problem arises when the ^{39}Ar decay rate becomes comparable to the maximum drift time of ionized electrons in the TPC. DarkSide-50 has measured atmospheric argon to have $\sim 1 \text{ Bq}^{39}\text{Ar}/\text{kg}$ [50]. With electron drift speeds around $1 \text{ mm}/\mu\text{s}$, as observed by DarkSide-50, the ^{39}Ar pileup becomes comparable to the electron drift times at LAr masses ~ 10 tonnes.

High rates of ^{39}Ar decays can typically be mitigated by using argon derived from underground sources, where the ^{39}Ar concentration is highly suppressed [98, 42, 99].

Origin

Argon is the third most abundant element in the earth’s atmosphere, and most samples of LAr are simply condensed out of the air. Measurements indicate that the natural concentration of ^{39}Ar in atmospheric argons samples is about $1.01 \pm 0.08 \text{ Bq}/\text{kg}$ [100].

Most of this ^{39}Ar is cosmogenic, produced through interactions with cosmic rays in the upper atmosphere. ^{39}Ar is typically formed when a cosmogenic neutron hits an ^{40}Ar target and knocks out a neutron, through the reaction



Since ^{39}Ar is produced with a roughly constant rate and decays over time, the concentration of ^{39}Ar in atmospheric argon tends to be in steady equilibrium.

While this process produces the vast majority of the ^{39}Ar found in atmospheric argon, argon extracted from underground sources still has a small but nonzero concentration of this isotope. ^{40}Ar can be found underground or in the atmosphere as a result of ^{40}K decay. However, where ^{40}K can be found, so can ^{39}K . While ^{39}K is stable, it can undergo the $^{39}\text{K}(n,p)^{39}\text{Ar}$ reaction with neutrons produced from nuclear decays in the surrounding rocks. Since these reactions require bombardment of neutrons, which are typically produced by reactions from the uranium and thorium chains in the mantle, both of which have ppb levels in the mantle but ppm levels in the crust, ^{39}Ar levels in the mantle are about 1000 times lower than those in the crust. DarkSide-50 has measured the ^{39}Ar concentration in underground argon to be lower than that in atmospheric argon by a factor of $(1.4 \pm 0.2) \times 10^3$ [42].




Figure 3.1: Decay diagram of ^{39}Ar . Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q_{β^-} -value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^- decays, the given energy is the endpoint energy of the β^- [15].

Decay

Figure 3.1 shows the β decay of ^{39}Ar through the reaction



Since this reaction flips the parity of the nucleus and changes the angular momentum by 2, it is first forbidden, explaining the long half-life of the decay.

^{39}Ar produces no γ -rays when it decays, leaving the only resulting signal in the liquid argon that from the β itself.

3.1.2 Krypton-85

Since krypton only constitutes $\sim 10^{-4}\%$ of the earth's atmosphere at sea level—much smaller than argon, which constitutes $\sim 0.9\%$ [101]—the radioisotope ^{85}Kr is typically not present in large quantities in samples of atmospheric argon. Typically, the process of purifying liquid argon should further reduce the amount of ^{85}Kr contamination in it.

However, samples of underground argon, which have drastically reduced levels of ^{39}Ar , actually may have higher activities of ^{85}Kr than of ^{39}Ar [12]. While these levels are low enough to not be a dominant background for a dark matter search, especially since it is possible that additional purification processes may be able to further reduce this contaminant and the electron recoils it produces from its β -decay can be effectively identified using PSD, this isotope is worth discussing because it produces a clear signal in DarkSide-50.




Figure 3.2: Decay diagram of ^{85}Kr . Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q_{β^-} -value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^- decays, the given energy is the endpoint energy of the β^- [16].

Origin

While ^{85}Kr is generally produced in the atmosphere through cosmogenic neutron activation of ^{84}Kr through the reaction $^{84}\text{Kr}(n, \gamma)^{85}\text{Kr}$, the amount seen by DarkSide-50 in the underground argon sample is higher than can be explained by atmospheric ^{85}Kr contamination.

^{85}Kr can also be produced underground from fission—fissile nuclei in the crust such as ^{235}U have a 0.1–0.9% chance of producing ^{85}Kr when they undergo thermal fission [102]. While the origin of the ^{85}Kr seen by DarkSide-50 is not yet fully understood, it is likely that one of these sources might explain the observed contamination [12].

Decay

Figure 3.2 shows the β decay of ^{85}Kr to ^{85}Rb . The primary decay mode goes directly to the ground state of ^{85}Rb , and produces a β spectrum with an endpoint of 687 keV. Since the ^{39}Ar β spectrum has an endpoint energy of 565 keV, the ^{85}Kr spectrum can make it hard to see the ^{39}Ar β spectrum if the ^{85}Kr rate is comparable to or greater than the rate of ^{39}Ar decays in the detector.

^{85}Kr also has a 0.434% probability of β decaying to ^{85m}Rb with an endpoint energy of 173 keV. This excited state of ^{85}Rb will then decay by isomeric transition down to the ground state of ^{85}Rb with a half-life of 1 μs . This slightly delayed coincidence

allows ^{85}Kr decays to be identified separately from other backgrounds so that the level of ^{85}Kr contamination can be measured. By identifying ^{85}Kr decays in this fashion, DarkSide-50 has measured $(2.05 \pm 0.13) \text{ mBq/kg}$ of ^{85}Kr in the underground argon [12].

3.1.3 Cobalt-60

While ^{60}Co is not the most abundant contaminant in DarkSide-50, it is particularly noteworthy because it produces two relatively high energy γ -rays in coincidence with each other, making it the strongest signal in prompt coincidence between the LSV and the TPC.

Origin

^{60}Co is not a naturally occurring isotope of cobalt. While small amounts of it may be produced due to cosmogenic activation, its relatively short half-life of 5.27 years means that it does not naturally accumulate in any appreciable amount. Similarly, this short half-life means that it is not a remnant of nuclear weapons testing, which produced many radioisotopes currently in the atmosphere; any ^{60}Co produced then would have decayed away by now.

Most ^{60}Co present today is man-made, either as a byproduct in nuclear reactors or neutron bombardment of ^{59}Co in a dedicated beam.

Since ^{60}Co produces two high energy γ -rays, it is useful for many technological applications where radioactive tagging is useful. One such application that is directly relevant to building detectors is tracing cobalt in steel alloy production. Manufacturers will often activate trace amounts of ^{59}Co in stainless steel alloys as a way of tagging their cobalt concentration. However, the quantities used are often small and likely do not account for all of the ^{60}Co present in modern steel.

Another very common application of ^{60}Co is in medical applications, where it can be used for sterilizing equipment, radiation therapy, or imaging. It is also often used outside of medical applications for pest sterilization, food irradiation, and radiography.

These applications make the demand for ^{60}Co reasonably high, meaning that much of it is intentionally produced to be used in these ways. However, a side effect of this production is that uncontrolled disposal of ^{60}Co often results in the isotope getting mixed in with iron. When steel or other sources of iron are melted down to be reused in other applications, any ^{60}Co unintentionally disposed with it will get mixed in, contaminating the batch of steel [103]. In addition to cases where this uncontrolled disposal is a result of carelessness on someone's part, there are also many cases where this has been an unintended consequence of more elaborate chains of events. For example, in 2013 a truck carrying 111 TBq of ^{60}Co from a hospital in Tijuana to a radioactive waste disposal site was hijacked and abandoned in an empty lot [104]. In this case, the ^{60}Co was found by authorities and properly handled before the car containing it got impounded and the ^{60}Co got mixed with the iron in the car. However, a number of cases in the past have found hazardously high




Figure 3.3: Decay diagram of ^{60}Co . Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Blue arrows and text refer to internal de-excitation processes within a nucleus, with the relative intensity, γ -ray energy, and lifetime of each transition. Different shades of blue were used to help identify which text describes which arrow. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q_{β^-} -value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^- decays, the given energy is the endpoint energy of the β^- [17].

levels of ^{60}Co contamination in other stainless steel samples. For example, in 2012 Petco issued a recall on their stainless steel dog food bowls due to dangerously high ^{60}Co contamination resulting from a similar uncontrolled ^{60}Co disposal [105]. The prevalence of cases like these is good evidence that uncontrolled ^{60}Co disposal is a large contributor to the levels of ^{60}Co contamination found in recycled iron samples today.

Decay

^{60}Co typically β decays through the reaction




Figure 3.3 shows the full β decay of ^{60}Co .

Since the ground state of ^{60}Co has a total angular momentum $J_\pi = 5+$, it will preferentially decay to $J_\pi = 4+$ energy level of ^{60}Ni , rather than to one of the lower energy levels with lower spin.

Not shown in Figure 3.3 is an energy level at 2284.8 keV with $J_\pi = 0+$. Due to angular momentum conservation and the small energy gap, the decay of the $J_\pi = 4+$ level to this level is highly suppressed. Instead, the $J_\pi = 4+$ level will decay to one of the two $J_\pi = 2+$ levels, or to the ground state, which is at a much lower energy. Of the two $J_\pi = 2+$ levels, the $J_\pi = 4+$ level will typically relax to the lower energy of the two and produce a 1173.2 keV γ -ray. The nucleus will then decay a second time to reach the ground state and produce a 1332.5 keV γ -ray in the process.

Since the two γ -rays produced in the ^{60}Co β -decay are produced by two different de-excitations, their momenta are totally uncorrelated. However, their coincidence as well as their relatively high energy compared to other relevant signals makes them an important background.

3.1.4 Potassium-40

While ^{40}K only produces electron recoils, it is one of the most abundant γ -ray-emitters naturally present in common detector materials. Since liquid argon-based detectors can use PSD to reject the electron recoils produced by these decays, ^{40}K does not pose a serious danger. Nevertheless, the signal it produces is clearly visible as an electron recoil due to its high abundance.

As will be discussed below, about 10.8% of ^{40}K decays directly produce ^{40}Ar . While ^{40}K is not a dangerous background in DarkSide-50, it is worth mentioning that the high abundance of ^{40}K in the earth's crusts results in ^{40}Ar production, which is the ultimate source of DarkSide-50's underground argon.

Origin

Potassium is the eighth most abundant element in the Earth's crust, with a typical abundance of 2.09% by weight [101]. It is also the eight most abundant element in the sea, with a typical concentration around 399 mg/L [101].

Due to its long half-life of over 1 billion years, most of the ^{40}K in the earth and sea water has not yet decayed, giving ^{40}K a natural abundance of 0.012% [106]. Since potassium has a very high natural abundance, ^{40}K tends to be ubiquitous in the world around us: ^{40}K from seawater and earth gets into most samples of water, is absorbed by plants, and is typically present in most metal ores. While chemical and industrial processes tend to purify their materials to various degrees, trace amounts of ^{40}K will usually survive the purification processes, allowing it to be found in most materials.

Decay

As shown in Figure 3.4, ^{40}K can undergo three different decay mechanisms. Since all three decay mechanisms involve a change in nuclear spin by $+4$ and a change in nuclear parity, the decay of ^{40}K is third forbidden, giving it a long half-life. The first





Figure 3.4: Decay diagram of ^{40}K . Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission (“EC” refers to electron capture), its energy, relative intensity, and half-life. Red lines and text show Q-value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^\pm decays, the given energy is the endpoint energy of the β^\pm [3]. *The decay of ^{40}Ca to ^{40}Ar has not been observed, but is believed to be possible due to the positive Q-value of this reaction and the fact that it does not violate any conserved quantities.

and most likely mechanism is a straightforward β -decay to the ground state of ^{40}Ca :



Since this reaction does not produce any γ -rays, and the ^{40}K is present in the non-scintillating components of the detector (*e.g.* the stainless steel and the phototubes), the decays through this channel are mostly invisible to the experiment. This is similarly true for the rarest of the three decay mechanisms, where ^{40}K β^+ -decays to ^{40}Ar through the reaction



However, in this case the β^+ will annihilate with an electron shortly after the reaction occurs, and this reaction will produce two γ -rays that can be seen by the detector with energies of 511 keV. This reaction has a branching ratio of 0.001%.

Lastly, ^{40}K can undergo electron capture through the reactions




Table 3.1: List of x-rays produced by the electron capture decay of ^{40}K [3].

x-ray line	Energy (keV)	Intensity (%)
$k_{\alpha 2}$	2.955	0.30
$k_{\alpha 1}$	2.957	0.61
$k_{\beta 1}$	3.19	0.048
$k_{\beta 3}$	3.19	0.025

This latter reaction happens in 98.2% of all ^{40}K electron captures and produces a 1460.8 keV γ -ray that can be seen by the detector. Additionally, since electron capture reactions remove an atomic electron from one of the atom’s inner shells, these reactions may be followed by x-rays as the atom relaxes or by the emission of a 2.66 keV Auger electron. The x-ray energies and their intensities produced by this reaction are given in Table 3.1.


These x-rays are very low energy and are generally not a problem for experiments like DarkSide-50 with energy thresholds above 10 keV_{ee}. However, they may be more important for experiments looking for low mass WIMPs with lower energy thresholds.

3.1.5 Tritium


Tritium naturally constitutes a small fraction of all hydrogen and is therefore present in trace concentrations in organic compounds like the liquid scintillator in the DarkSide-50 neutron veto. Since the endpoint of the ^3H β spectrum is 18.6 keV, ^3H decays are generally low enough energy to not interfere with the veto’s sensitivity. Nonetheless, it produces a visible signal in the detector and is worth mentioning.

Origin

^3H is produced in the upper atmospheres from activation of ^{14}N from fast cosmogenic neutrons through the reaction



It can also be produced underground from neutron capture on ^6Li through the reaction



or by fission (spontaneous or thermal neutron-induced) of heavy isotopes in the earth’s crust, such as ^{238}U and ^{235}U , among others [107]. This mechanism is supported by observations of tritium in petroleum products as measured by [108].

Decay

Figure 3.5 shows the β decay of ^3H , described by the reaction






Figure 3.5: Decay diagram of ${}^3\text{H}$. Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q -value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β decays, the given energy is the endpoint energy of the β [18].




Figure 3.6: The ${}^3\text{H}$ β^- spectrum (*blue*) with the Fermi function correction for the nucleus's Coulomb potential and (*green*) without. The correction decreases the mean β^- energy by 0.28 keV.

This decay is allowed, since it changes neither the nuclear spin nor the parity. The small Q -value of 18.6 keV for the reaction means that the decay has a relatively small phase space factor and therefore a half-life of 12.3 years.

Since this decay is allowed, it is simple to calculate the energy spectrum of the outgoing β , following the prescription described in Chapter 4.1.

Since this decay is at a low energy, the Coulomb effects of the nucleus on the electron are clear to see. Since the nucleus exerts an attractive force on the β^- , it has the effect of shifting the β spectrum to lower energies. This is seen in Figure 3.6, which shows the ${}^3\text{H}$ β spectrum with and without the Fermi correction for these Coulomb effects.


3.1.6 Carbon-14

Since ^{14}C β decays with no accompanying γ -rays, it is unlikely to produce a signal in the active volume of a liquid argon-based detector. However, since it is an unavoidable component of carbon, and the LSV of DarkSide-50—discussed in Chapter 5—is based on an organic liquid scintillator rich in carbon, ^{14}C plays an important role in the signal seen by the veto.

A high enough rate of ^{14}C decays in the veto may easily produce too much background for the veto to efficiently see low energy coincidences between the TPC and veto that may result from a neutron event. A low ^{14}C background in the veto is therefore crucial for the veto to achieve a very high efficiency, making the ^{14}C background especially important to understand.

Origin

Most ^{14}C is generated in the upper atmosphere from cosmogenic activation. The most common production mode is from cosmogenic neutrons undergoing the $^{14}\text{N}(\text{n,p})^{14}\text{C}$ reaction. It can also be produced by various neutron capture reactions on ^{17}O and ^{13}C . These reactions are described by



These processes produce ^{14}C in the atmosphere at an approximately constant rate. Due to the isotope's long half-life of 5700 years [19], ^{14}C tends to accumulate to an approximately constant concentration in atmospheric CO_2 . An exception to this approximately constant rate occurred during nuclear weapons testing in the mid-1960's, which caused a large spike in the concentration of atmospheric ^{14}C that can be seen in carbon dating measurements.

Organic compounds are frequently derived starting with a methanol base. One common source of the methanol is plant matter. Since plants spend their lives exchanging carbon with the atmosphere, they tend to have approximately atmospheric concentrations of ^{14}C . For modern samples of carbon, the ^{14}C concentration is typically $^{14}\text{C}/^{12}\text{C} \sim 10^{-12}$ [109].

Another method of producing these compounds is to derive the methanol from petroleum rather than plant matter. In this case, the carbon in the petroleum has been deep underground, shielded from cosmic radiation, for far longer than the ^{14}C half-life. This means that nearly all of the ^{14}C in the petroleum will have decayed away, leaving very isotopically pure samples. However, some amount of ^{14}C may still be produced, albeit at very low rates, resulting from similar neutron capture reactions to those mentioned above resulting from neutrons produced via the (α, n) reaction on various nuclei present underground. Borexino has measured a ^{14}C concentration of $^{14}\text{C}/^{12}\text{C} \sim 10^{-18}$ in petroleum-derived pseudocumene [109].




Figure 3.7: Decay diagram of ^{14}C and ^{14}O to ^{14}N . Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Blue arrows and text refer to internal de-excitation processes within a nucleus, with the relative intensity, γ -ray energy, and lifetime of each transition. Different shades of blue were used to help identify which text describes which arrow. Red lines and text show Q-value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^\pm decays, the given energy is the endpoint energy of the β^\pm [19]. The three $T = 1$ multiplet states are labeled.

Decay

As shown in Figure 3.7, ^{14}C decays through the interaction

$$^{14}\text{C} \rightarrow ^{14}\text{N} + \beta^- + \bar{\nu}_e$$

with a half-life of 5700 years and an endpoint energy of 156.6 keV.

In order to understand the decay of ^{14}C , it is helpful to consider the isospin triplet (^{14}C , ^{14}N , ^{14}O). ^{14}C , ^{14}N , and ^{14}O each have total isospin $T = 1$ with isospins $t_3 = -1, 0$, and $+1$, respectively, as shown in Figure 3.7. ^{14}C and ^{14}O both have ground states with spin $0+$ and a ground state with spin $0+$; the ground state of ^{14}N has spin $1+$, but the first excited state in the $T = 1$ multiplet has spin of $0+$, as well.

As shown in Figure 3.7, the energy of the $T = 1$ isospin analog states increases with Z due to Coulomb effects in the nucleus. Therefore, since the spin, parity, and T does not change for the β^+ decay of ^{14}O to ^{14}N , and the $0+$ state is higher for the parent nucleus, this decay is super allowed, and happens on a short time scale with a half-life of 70.7 s. However, the ground state of ^{14}C is lower energy than the isospin analog state in ^{14}N , which means ^{14}C cannot decay to this state. Instead, it must decay directly to the ground state of ^{14}N .




Figure 3.8: The ^{14}C β spectrum (*blue*) with the shape factor correction predicted by [20], and (*green*) without this correction.

This decay conserves nuclear parity but changes the nuclear spin by $+1$, and is therefore described by a Gamow-Teller interaction. Gamow-Teller interactions include contributions from an axial-vector interaction and allows for contributions from a weak-magnetism vector term. While we would expect the decay of ^{14}C to be slower than that of ^{14}O due to its smaller phase space and it being merely allowed rather than super allowed, the long half-life observed is on scale with many second forbidden β -decays. It has been suggested that for the $A = 14$ case, the Gamow-Teller matrix element may be highly suppressed, making it comparable to the weak-magnetism vector current, which could result in a non-trivial shape factor $C(E)$ given by

$$C(E) = 1 + aE + \mu_1 \sqrt{1 + (\alpha Z)^2} b/E + cE^2 \quad (3.2)$$

where E is the β energy, μ_1 is a special Coulomb function described in [110], Z is the nuclear charge, α is the fine-structure constant, and a, b , and c are constants [111]. The values of these constants are discussed in more detail in [112]. It has been shown in [20] that it is a good approximation for ^{14}C to assume that $b \approx c \approx 0$, and that we can expect $a = -0.38 \pm 0.04 \text{ MeV}^{-1}$. In this case, the shape factor may suppress the β decay and make the ^{14}C decay behave like a forbidden interaction. The effects of the shape factor on the spectrum are shown in Figure 3.8, where we can see that accounting for a distorts the ^{14}C spectrum at higher energies.

3.1.7 Thorium-232

^{232}Th , as well as ^{238}U and ^{235}U , which we will discuss later, is the head of the ^{232}Th decay chain, which includes a wide array of important radioisotopes. In particular,

most internal neutrons will come from these three decay chains: each of these three nuclei can undergo spontaneous fission, which can produce multiple neutrons. This process is discussed in more detail in Chapter 4.4.2. In addition, there are many α -emitters in these decay chains; these α particles may react with other nuclei to undergo the (α, n) reaction, which can produce fast neutrons, as discussed in Chapter 4.4.3. While fission neutrons are often accompanied by high energy γ -rays and may come in higher multiplicities, making them easier to veto, (α, n) neutrons tend to be produced alone, often without any accompanying γ -rays, making them an even more dangerous background.

Origin

Thorium is naturally present in the crust of the earth at a concentration of about 9.6 ppm and in seawater at a concentration of 1×10^{-6} mg/L [101]. With a half-life of 1.4×10^{10} years [21], the ^{232}Th in the crust is of primordial origin. Due to its presence in rock, trace amounts of ^{232}Th tend to be present in most metal samples.

Decay Chain

Figure 3.9 shows the decay chain for ^{232}Th , along with decay modes and γ -rays associated with the decays. The decay chain starts with ^{232}Th and ends with ^{208}Pb , which is stable. In general, since ^{232}Th has a half-life far longer than the age of the earth, it is safe to assume that this entire decay chain is in secular equilibrium, and that measuring the activity of one part of the chain is sufficient for determining the activity of each isotope in the chain. While various chemical and mechanical processes may remove some isotopes and break secular equilibrium, most of the progeny of ^{232}Th have short enough half-lives that secular equilibrium will likely have been restored by now.

Indeed, as will be discussed in Section 3.2, we see no evidence of the ^{232}Th decay chain being broken anywhere—not only does this support the assumption that the chain is in secular equilibrium, but it also means that if secular equilibrium is broken later down the chain, we cannot account for it, anyway.

Neutrons

In addition to α decaying, as shown in Figure 3.9, ^{232}Th can also undergo spontaneous fission with a branching ratio of 1.4×10^{-11} . While this probability is very small, it is worth accounting for because these spontaneous fission reactions can produce neutron backgrounds.

The energy spectrum of these neutrons is shown in Figure 3.10, where we can see the mean neutron energy is 1.697 MeV, and the energy spectrum drops off approximately exponentially from the peak at 0.85 MeV. Each of these spontaneous fission events produces an average of 2.14 neutrons [4]. The total fission neutron yield for ^{232}Th is 3.0×10^{-14} n/s/mBq of ^{232}Th . As will be discussed in Section 3.2, there is a total of 524 mBq of ^{232}Th in the components of DarkSide-50 surrounding the TPC. This gives a total fission neutron production rate 1.6×10^{-11} n/s. Over the span of

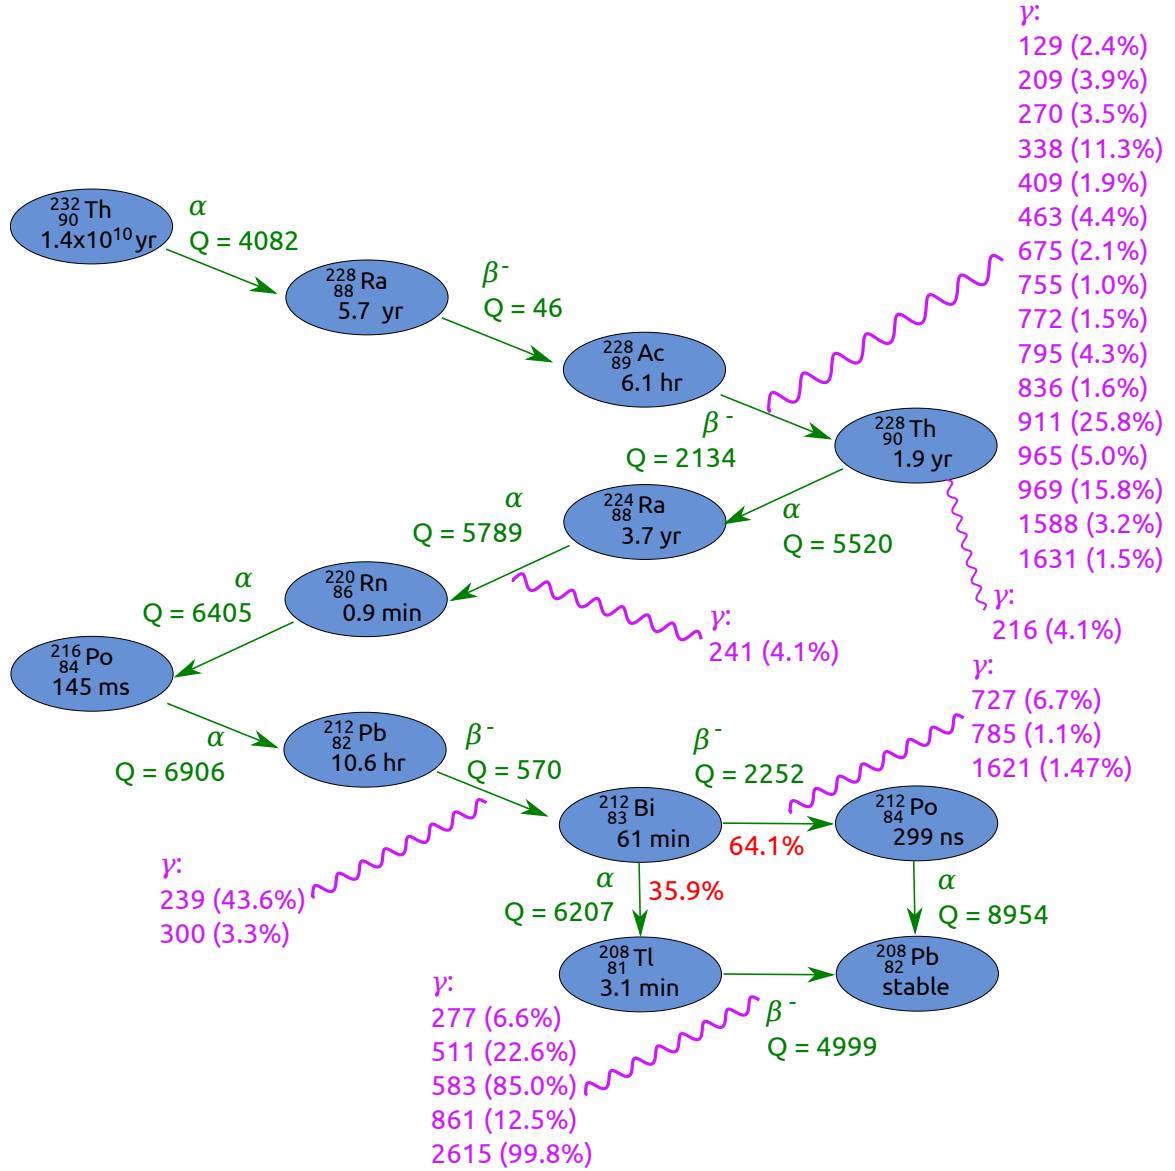


Figure 3.9: The ^{232}Th decay chain. Blue circles are each isotope in the chain above the isotope's half-life. Green arrows and text represent decay modes and their associated Q-values in keV. Red numbers show the branching ratio whenever the chain splits. Purple wavy lines and text represent γ -rays produced by the decay they are associated with, written as the γ -ray energy in keV followed by the intensity of that γ -ray in parentheses. The intensity of a γ -ray is the probability of that γ -ray appearing in a decay. Only γ -rays with intensities greater than 1% and energies greater than 100 keV are included. Only branches with a branching ratio above 1×10^{-6} are included [21].




Figure 3.10: Fission neutron production rate per mBq of ^{232}Th , as simulated by SOURCES4C [4].

a three year campaign, this means we would expect 1.5×10^{-3} spontaneous fission neutrons to be produced by ^{232}Th . Since many of these neutrons will not even leave a signal in the TPC, and those that do will be vetoed with a very high efficiency, this particular background is not a great concern.

Neutrons produced by the (α, n) reaction in the ^{232}Th chain will be discussed in more detail in Chapter 4.4.3.

Thallium-208

While going through every radioisotope in Figure 3.9 is beyond the scope of this dissertation, it is worth talking about some of the more prominent decays. In particular, ^{208}Tl , near the bottom of the ^{232}Th decay chain, produces a 2.6 MeV γ -ray in 99.8% of its decays. This γ -ray is the highest energy γ -ray produced in any of the naturally present radioisotopes. This high energy, in conjunction with the 500-800 keV γ -rays that it is frequently in coincidence with, makes a very distinctive high energy electronic recoil signal that can easily be identified in the detectors, and it can therefore be used as a useful calibration point, similar to ^{60}Co .

3.1.8 Uranium-238

Much like ^{232}Th , the ^{238}U decay chain contains many radioisotopes that can constitute various γ -ray and surface backgrounds. It can also produce fission neutrons and (α, n) neutrons, making this chain important to understand.

Origin

Uranium is naturally present in the crust of the earth in typical concentrations of 2.7 ppm and in seawater with concentrations around 0.0032 mg/L [101]. ^{238}U has a natural abundance of 99.27% [101]. As a result, most metals mined from the earth will have some trace level of ^{238}U contamination.

Additionally, the concentration in seawater means that most things washed in water will end up with some amount of ^{238}U and its daughters deposited on their surfaces, though the concentration of ^{238}U in the water may vary depending on the quality and purity of the water. Additional ^{238}U daughters may be found on inner surfaces of detector materials resulting from radon naturally present in the air. ^{238}U underground will eventually produce ^{222}Rn , which can diffuse through the crust and into these air. When these ^{222}Rn nuclei decay, they may deposit their daughters on the surfaces exposed to air. Most of the daughters below ^{222}Rn have half-lives on the scale of minutes or shorter, meaning that they will be depleted before long. However, ^{210}Pb has a half-life of 22.2 years, meaning that it and its daughters can persist throughout the duration of an experiment.

Decay Chain

Figure 3.11 shows the ^{238}U decay chain. As will be discussed in Section 3.2, there is evidence that this decay chain breaks above ^{226}Ra . Radium is a highly reactive element, and its precursor is ^{230}Th , which has a half-life of 7.5×10^4 years. Since ^{226}Ra has a half-life of 1600 years, it is possible that ^{226}Ra contamination in various materials can be removed, either naturally or when the material is being processed. The long half-life of ^{230}Th means that ^{226}Ra that is chemically removed from a sample will not be restored to secular equilibrium within the timescale of an experiment. We therefore expect secular equilibrium in the ^{238}U decay chain to be broken between ^{230}Th and ^{226}Ra ; we refer to the precursors of ^{226}Ra as the upper ^{238}U decay chain, and ^{226}Ra and its progeny as the lower ^{238}U decay chain.

Additionally, since ^{210}Pb is a natural component of lead found in the crust, many materials that have lead in them may have additional ^{210}Pb radioactivity. Since ^{210}Pb has a half-life of 22 years, this additional radioactivity will not deplete within the timescale of an experiment, and so ^{210}Pb and the daughters below it may be out of secular equilibrium with the rest of the lower part of the ^{238}U decay chain. While we do see evidence of this in some materials, it is worth noting that most materials for which we have directly measured the ^{210}Pb contamination in DarkSide-50, we find the levels consistent with the rest of the lower chain.

Neutrons

Similar to ^{232}Th , ^{238}U can also undergo spontaneous fission, but with a much higher branching ratio of 5.45×10^{-7} [4].

The energy spectrum of these neutrons is shown in Figure 3.12. These neutrons have a mean energy of 1.788 MeV, with a peak around 0.85 MeV. The spectrum then drops off approximately exponentially after the peak. Each spontaneous fission

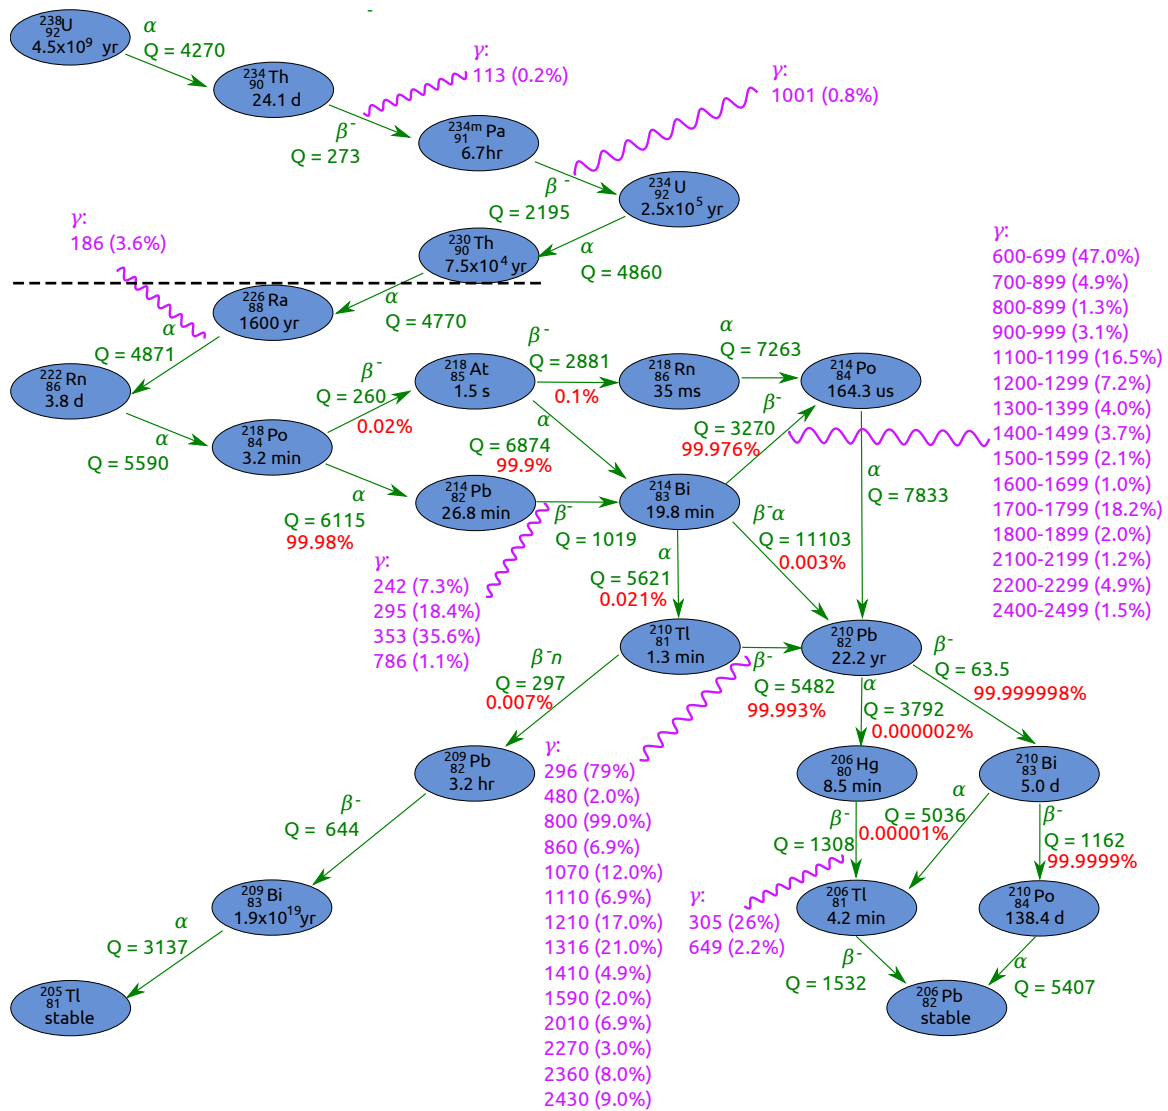


Figure 3.11: The ^{238}U decay chain. Blue circles are each isotope in the chain above the isotope's half-life. Green arrows and text represent decay modes and their associated Q-values in keV. Red numbers show the branching ratio whenever the chain splits. Purple wavy lines and text represent γ -rays produced by the decay they are associated with, written as the γ -ray energy in keV followed by intensity of that γ -ray in parentheses. Only γ -rays with intensities greater than 1% and energies greater than 100 keV are included. Only branches with a branching ratio above 1×10^{-6} are included [22]. The black dashed line above ^{226}Ra shows where secular equilibrium is likely broken.




Figure 3.12: Fission neutron production rate per mBq of ^{238}U , as simulated by SOURCES4C [4].

event produces an average of 2.01 neutrons [4]. The total fission neutron yield for ^{238}U is 1.10×10^{-9} n/s/mBq of ^{238}U . As will be discussed in Section 3.2, there is a total of 7.2 ± 0.9 Bq of ^{238}U in the components of DarkSide-50 surrounding the TPC. This gives a total fission neutron production rate $7.9 \pm 1.0 \times 10^{-6}$ n/s. Over the span of a year, this means we would expect 250 ± 30 spontaneous fission neutrons to be produced by ^{238}U . During a three year campaign, this comes to 750 ± 90 fission neutrons. However, it should be noted that most of these neutrons may not leave any signal in the detector due to geometrical constraints.

Neutrons produced by the (α, n) reaction in the ^{238}U chain will be discussed in more detail in Chapter 4.4.3.

Delayed Neutrons

As can be seen in Figure 3.11, in addition to β -decaying, ^{210}Tl can also β -decay with a delayed neutron emission to ^{209}Pb . Overall, ^{210}Tl will be produced in 0.021% of all decays of the lower ^{238}U decay chain, and it will decay via delayed neutron emission with a branching ratio of 0.007%. In total, we then expect this delayed neutron to be produced in $1.5 \times 10^{-6}\%$ of all decays of this chain. For comparison, as will be discussed in Chapter 4.4.3, the lower ^{238}U chain in general tends to have neutron yields ranging from 10^{-3} – $10^{-8}\%$, depending on the material; the delayed neutron yield for ^{210}Tl falls on the lower end of this range.

Figure 3.13 shows this decay process, glossing over the β -decays to lower energy levels of ^{210}Pb . This diagram shows that this decay process can produce neutrons with kinetic energies of either 260 keV or 201 keV.




Figure 3.13: Decay diagram of ^{210}Tl via the β^-n channel to ^{209}Pb . The much more likely branches that ^{210}Tl can decay to via normal β^- decay to ^{210}Pb are not shown. Black horizontal lines show nuclear energy levels of the parent and daughter nuclei. To the left of the lines is the total angular momentum and parity of the energy level; to the right is the energy level (in keV) above the ground state. Green arrows and text show nuclear decay emission, its energy, relative intensity, and half-life. Red lines and text show Q-value for the nuclear decay in keV, separating the parent and daughter ground state energy levels. For β^\pm decays, the given energy is the endpoint kinetic energy of the β^\pm . Spin and parity of the excited states shown for ^{210}Pb are not shown because they were not listed in the ENSDF dataset. Similarly, the relative intensities of both of these branches were not specified, but their total intensity is 0.007% [23, 24].

As will be discussed in Section 3.2, we expect a total of 1.5 ± 0.4 Bq of the lower ^{238}U chain in DarkSide-50. Accounting for the intensity of this particular decay relative to the rest of the lower chain, we expect these delayed neutrons to be produced at a rate of $(2.3 \pm 0.6) \times 10^{-8}$ Hz. This comes to a total of 0.7 ± 0.2 n/year, or 2.1 ± 0.6 neutrons total in a three year WIMP search.

Surface Backgrounds

As mentioned above, ^{222}Rn daughters may be deposited on the inner surfaces of detector materials as a result of exposure to air and the ^{222}Rn naturally present in the air. Of these daughters, those that α -decay are generally more of a concern to liquid argon-based dark matter detectors than those that β -decay, since PSD can effectively remove the β s. As can be seen in Figure 3.11, most of the ^{222}Rn daughters above ^{210}Pb have a very short half-life and are therefore likely not present in high activities. However, the long half-life of ^{210}Pb means that the daughters below it may persist throughout the duration of the experiment. While ^{210}Pb itself can α -decay to ^{206}Hg , it is very highly suppressed by its low branching ratio. This argument similarly holds for ^{210}Bi decaying to ^{206}Tl . This leaves the most concerning daughter as ^{210}Po .

^{210}Po may α -decay to ^{206}Pb , producing a 5.304 MeV α and a 103 keV recoiling ^{206}Pb nucleus. While either of these products may deposit energy in the active LAr volume of the detector, the high energy of the α means that it will generally be above the WIMP energy region of interest—however, if decays earlier in the chain have embedded the ^{210}Po nucleus deeper in the surface, the degraded α may pose a greater risk. On the other hand, the much lower energy of the ^{206}Pb nucleus means that it may pose a concern for detectors with low energy thresholds, where this surface background may create a signal similar to a WIMP event. Since both of these products are nuclei, the recoils that they produce will be nuclear recoils like we expect from WIMPs.

3.1.9 Uranium-235

Much like ^{238}U , ^{235}U is a naturally occurring isotope of uranium and is nearly inevitably present in trace amounts of metal samples. Although it is always present in smaller concentrations than ^{238}U , the ^{235}U decay chain contains higher energy α decays, making the backgrounds in this chain important to understand.

Origin

^{235}U is present in materials for the same reason that ^{238}U is. The fact that it is very hard to separate these isotopes means that wherever one is present, the other will be as well, and they will generally be present in a fixed ratio.

Since ^{235}U has a natural abundance of 0.720% [101] and a half-life of 7.04×10^8 years, compared to ^{238}U with a natural abundance of 99.275% and a half-life of 4.47×10^9 years, we can calculate the ratio of their activities as

$$\frac{A_{238\text{U}}}{A_{235\text{U}}} = \frac{99.275 \times t_{235\text{U}}^{1/2}}{0.720 \times t_{238\text{U}}^{1/2}} \approx 21.7 \quad (3.3)$$

Decay Chain

Figure 3.14 shows the decay chain for ^{235}U . Much like ^{232}Th , while it is possible that the ^{235}U chain may break between ^{223}Ra and ^{219}Rn due to the ^{219}Rn diffusing out of the material, ~ 11 day half-life of ^{223}Ra means that the decay chain is likely to be brought back into secular equilibrium within a matter of months.

Additionally, since all of the isotopes in the decay chain below ^{219}Rn have very short half-lives (the highest being 31.2 minutes), any ^{219}Rn daughters that may have deposited on the inner surfaces of the detector would decay away in under a day, meaning that there are no significant surface backgrounds from the ^{235}U decay chain.

Neutrons

^{235}U can also undergo spontaneous fission with a branching ratio of 2.011×10^{-9} [4].

The energy spectrum of these neutrons is shown in Figure 3.15. These neutrons have a mean energy of 1.99 MeV, with a peak around 0.85 MeV. The spectrum then

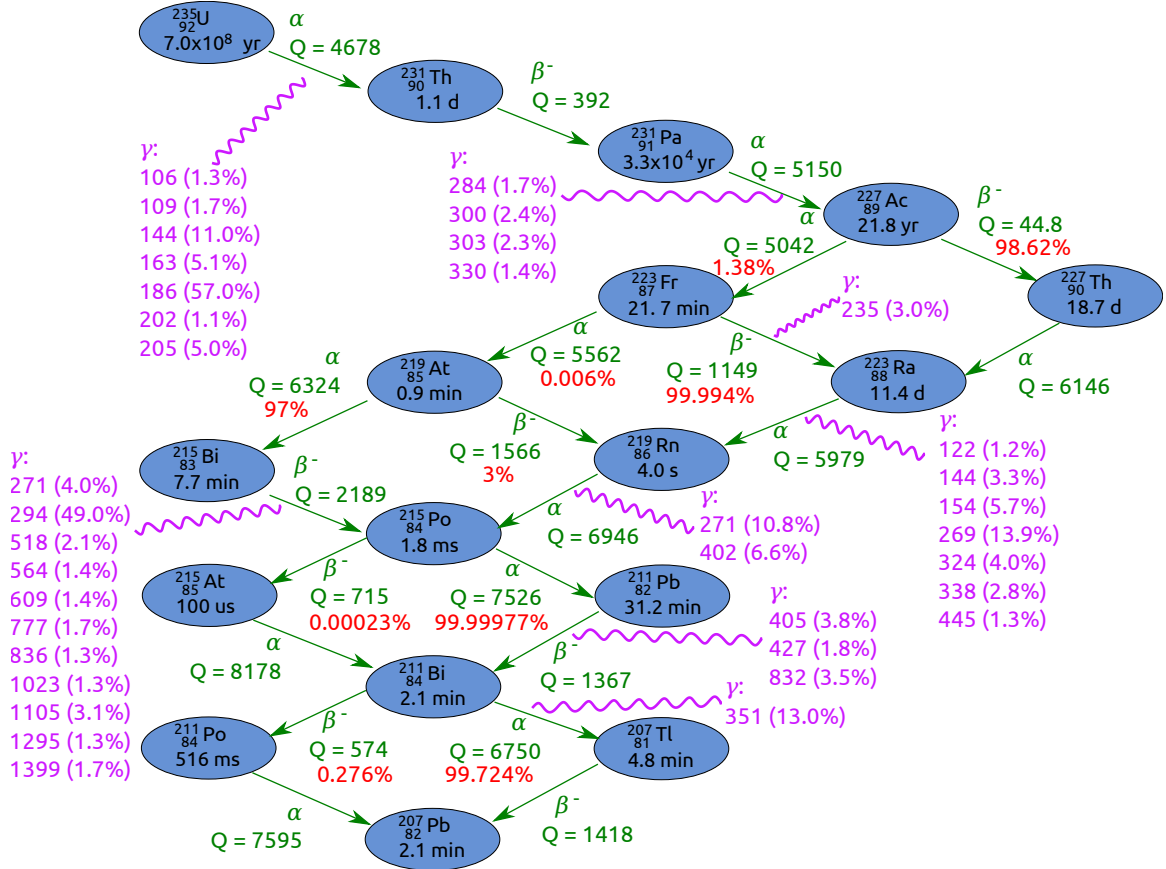


Figure 3.14: The ^{235}U decay chain. Blue circles are each isotope in the chain above the isotope's half-life. Green arrows and text represent decay modes and their associated Q-values in keV. Red numbers show the branching ratio whenever the chain splits. Purple wavy lines and text represent γ -rays produced by the decay they are associated with, written as the γ -ray energy in keV followed by intensity of that γ -ray in parentheses. Only γ -rays with intensities greater than 1% and energies greater than 100 keV are included. Only branches with a branching ratio above 1×10^{-6} are included [25].




Figure 3.15: Fission neutron production rate per mBq of ^{235}U , as simulated by SOURCES4C [4].

drops off approximately exponentially after the peak. Each spontaneous fission event produces an average of 1.86 neutrons [4]. The total fission neutron yield for ^{235}U is 3.75×10^{-12} n/s/mBq of ^{235}U . As will be discussed in Section 3.2, there is a total of 332 mBq of ^{235}U in the components of DarkSide-50 surrounding the TPC. This gives a total fission neutron production rate 1.25×10^{-9} n/s. Over the span of a three year campaign, this means we would expect 0.04 spontaneous fission neutrons to be produced by ^{235}U , most of which may not leave any signal in the detector due to geometrical constraints. Therefore, we expect very little fission neutron background resulting from ^{235}U .

Neutrons produced by the (α, n) reaction in the ^{235}U chain will be discussed in more detail in Chapter 4.4.3.

3.2 DarkSide-50 Counting Results

Having established which radioactive contaminants may be present in the detector, the next step to building a background model is to establish the total contamination of each radioisotope. These contamination levels can be established *in situ* by analyzing the spectrum in the TPC and in the neutron veto, or it can be done *ex situ* by isotopically analyzing the materials used for as many individual components as possible.

Ideally, these two techniques will be used in conjunction with each other to establish and validate a full background model. Much of the work being done to determine the radioactive contamination *in situ* is being done by H. Qian by analyzing the

signals that the neutron veto sees in coincidence with the TPC, while P. Agnes is working on studying the signals in the TPC. We will discuss this technique in more detail later. Here, we will focus on the *ex situ* measurements made of the various DarkSide-50 component materials.

The one exception is that we will discuss the *in situ* measurements made by C. Stanford of the uranium and thorium decay chain contamination in liquid argon. These measurements were made by counting α -decays in the bulk of the liquid argon, which are easily identified by their high energy above other backgrounds, and by looking for bismuth-polonium coincidence signals.

3.2.1 Counting Techniques

Ex situ measurements of the radioactive contamination of various materials were performed via Germanium Counting, Glow Discharge Mass Spectrometry (GDMS), Neutron Activation Analysis (NAA), and Inductively Coupled Plasma Mass Spectrometry (ICP-MS).

Germanium Counting measurements were performed at LNGS by M. Laubenstein and are discussed in [103]. This technique uses an array of germanium crystal detectors that can detect γ -rays emitted from decays. By analyzing the γ -ray spectrum seen by the crystals, one can therefore either deduce the concentration of various γ -ray-emitting isotopes or place an upper limit on the concentration. As a result, this technique is only sensitive to the contamination levels of isotopes that emit γ -rays when they decay. However, as discussed in Section 3.1, we can often assume that decay chains are mostly in secular equilibrium, with a few points where equilibrium may be broken, and infer the contamination levels of the rest of the chain. These results were typically sensitive on the scale of ~ 1 ppb. Measurements of the ^{232}Th decay chain are typically inferred from measurements of the 911 keV line in ^{228}Ac , the 239 keV line in ^{212}Pb , and the 583 and 2615 keV lines in ^{208}Tl . Measurements of the ^{235}U chain are inferred from measurements of the 186 keV γ -ray line from ^{235}U . Measurements of the upper ^{238}U decay chain are inferred from measurements of the 1001 keV γ -rays emitted by ^{234m}Pa and the 113 keV line emitted by ^{234}Th , while measurements of the lower ^{238}U decay chain are inferred from the 186 keV line emitted by ^{226}Ra , the 295 and 353 keV lines emitted by ^{214}Pb , the 609, 1120, and 1764 keV lines emitted by ^{214}Bi .

GDMS measurements were made by the Evans Analytical Group. This technique analyzes a solid sample by making it the cathode of a chamber of positive argon ions. These ions are accelerated towards the cathode with a kinetic energy ~ 100 keV, and atomize the surface of the cathode when they hit it. The escaping atoms are then ionized in the argon gas and then drifted into a mass spectrometer, where trace amounts of different elements can be measured. This technique is sensitive to the identity of the contaminants, and therefore directly measures the amount of uranium and thorium present. Since it is only used to measure the tops of the decay chains, we must assume the rest of the chain is in secular equilibrium with the top. Nevertheless, this technique is very sensitive and can measure concentrations below 1 ppb.

NAA measurements are performed by the North Carolina State University Nuclear Reactor Program Nuclear Services. This analysis technique involves placing a sample in a neutron beam, where they are bombarded by high energy neutrons. Many of the isotopes that capture these neutrons will then become radioactive. After neutron activation, the samples are transferred to a high purity germanium target, where γ -rays from the activated radioisotopes are then counted, and a spectral analysis reveals the contamination levels of various radioisotopes in the sample. This technique is sensitive to ~ 1 ppt concentrations, though reports from the Nuclear Services group only provide the activity of the tops of the decay chains.

The ICP-MS technique is approximately a generalization of the GDMS technique using different plasmas to atomize the sample. For DarkSide-50, ICP-MS measurements were performed by Pacific Northwest National Laboratory and LNGS. These measurements typically had sensitivity on the scale of ~ 1 ppt and only reported measurements of the tops of the decay chains.

3.2.2 Summary of Measurements

Samples of several different materials were sent to be counted using the above techniques. We summarize the results of these measurements in Tables 3.2 and 3.3. For many materials, we had measurements taken using multiple techniques. In these cases, if one technique reported an upper limit (typically to 1σ confidence) while the other reported a measurement, we used the measured value. Otherwise, we combined the results by averaging the reported values together, weighing the measurements by a factor of $1/\sigma$. Germanium counting measurements that have multiple measurements of the same part of the decay chain were treated the same way. If both measurements gave upper limits, the lower of the two limits was used.

Typically, GDMS measurements did not report any uncertainty. For these measurements, we assumed an uncertainty of ± 1 in the least significant digit.

Since GDMS, NAA, and ICP-MS techniques report the total concentration of uranium present, Equation 3.3 is used to determine how much of the measured uranium is ^{238}U and how much is ^{235}U . These measurements also reported contamination levels in units of ppb or ppt, while Germanium Counting measurements typically used units of mBq/kg. To convert all measurements to the latter units, the following conversion factors were used for the uranium and thorium decay chains:

- ^{235}U : 0.51 mBq/kg/ppb
- ^{238}U : 12.35 mBq/kg/ppb
- ^{232}Th : 4.07 mBq/kg/ppb

These conversion factors result from multiplying the specific activity of the relevant isotope by a factor of 1×10^{-9} .

In addition to the measurements listed in Tables 3.2 and 3.3, some isotopes were found to have excess activity below ^{210}Pb in the ^{238}U decay chain. The excess ^{210}Pb about the rest of the lower ^{238}U chain is reported in Table 3.4.

Table 3.2: Counting results measured per unit mass. Measurements' sources are marked with superscripts: ^a Germanium Counting, ^b GDMS, ^c NAA, ^dICP-MS, ^e *in situ* measurement.

Component	²³⁸ U _{upper}	²³⁸ U _{lower}	²³⁵ U	²³² Th
	[mBq/kg]			
Cryostat Body	2.5±1.2 ^b	0.4±0.2 ^a	0.10±0.05 ^b	0.8±0.3 ^a
Cryostat Flanges	12±12 ^b	12±12 ^b	0.51±0.51 ^b	0.8±0.4 ^a
Organ Pipe Tubes	<100 ^a	1.7±1.0 ^a	<5.3 ^a	<3 ^a
Cryostat Bolts	50±12 ^b	2±1 ^a	2.04±0.51 ^b	16.5±1.4 ^b
Cryostat Nuts	37±12 ^b	37±12 ^b	1.5±0.5 ^b	32.6±4.1 ^b
Cryostat Washers	12±12 ^b	4±2 ^a	0.51±0.51 ^b	17.4±1.7 ^a
MLI Insulator	93±28 ^d	93±28 ^d	3.8±1.2 ^d	3.6±1.1 ^d
MLI Netting	0.80±0.24 ^d	0.80±0.24 ^d	0.033±0.010 ^d	0.11±0.03 ^d
MLI Al Coating	62±19 ^d	62±19 ^d	2.6±0.8 ^d	<24 ^d
Argon (in μ Bq/kg)	0 ^e	1.31±0.01 ^e	0 ^e	0.17±0.04 ^e
Teflon Supports	<0.073 ^c	<0.073 ^c	<0.003 ^c	<0.004 ^c
Teflon Reflector	<0.028 ^c	<0.028 ^c	<0.001 ^c	<0.005 ^c
Cathode Window	0.37±0.04 ^d	0.37±0.04 ^d	0.017±0.002 ^d	0.041±0.004 ^d
Diving Bell	0.37±0.04 ^d	0.37±0.04 ^d	0.017±0.002 ^d	0.041±0.004 ^d
Field Cage Rings	1.15±0.12 ^b	1.15±0.12 ^b	0.05±0.005 ^b	0.40±0.04 ^b
Extraction Grid	<1.2 ^b	<1.2 ^b	<0.05 ^b	1.6±0.4 ^b
Grid Support Ring	2.5±1.2 ^b	0.4±0.2 ^a	0.10±0.05 ^b	0.8±0.3 ^a
PMT Shell	<0.12 ^b	<0.12 ^b	<5.7×10 ^{-3b}	<0.04 ^b
PMT Supp. Screws	<0.0012 ^b	<0.0012 ^b	<5.1×10 ^{-5b}	<0.0004 ^b
PMT Resistors	5300±600 ^a	270±40 ^a	190±30 ^a	<120 ^a
PMT Capacitors	(7±4)×10 ^{3a}	(4±2)×10 ^{3a}	80±40 ^a	800±100 ^a
PMT Barrel Conn.	<0.62 ^b	<0.62 ^b	<0.026 ^b	<0.20 ^b
Cold Amp Circuitry	180±90 ^a	300±22 ^a	33±7 ^a	113±9 ^a
Male MCX Conn.	180±90 ^a	46±8 ^a	<12 ^a	18±7 ^a
Female MCX Conn.	391±109 ^a	11±5 ^a	18±5 ^a	<15 ^a
Solder	<12.35 ^b	<12.35 ^b	<0.51 ^b	<4.04 ^b

Table 3.3: Counting results measured per piece. Measurements' sources are marked with superscripts: ^a Germanium Counting, ^b GDMS, ^c NAA, ^dICP-MS, ^e *in situ* measurement.

Component	²³⁸ U _{upper}	²³⁸ U _{lower}	²³⁵ U	²³² Th
	[mBq/pc]			
Outer Cryostat O-ring	206±12 ^a	206±12 ^a	9.5±0. ^a	44±3 ^a
Field Cage Resist. ×10 ⁻³	22±7 ^a	16.4±1.4 ^a	1.3±0.5 ^a	8.0±0.9 ^a
PMT Stem	150.5±21.5 ^a	18.2±1.2 ^a	7±1 ^a	6±1 ^a
PMT Circuit Board	<0.16 ^a	0.0344±0.0004 ^a	<1×10 ^{-5a}	<0.011 ^a
Cold Amp Circuit Board	<0.16 ^a	0.0344±0.0004 ^a	<1×10 ^{-5a}	<0.011 ^a

Table 3.4: Counting measurements of ^{210}Pb about the rest of the lower ^{238}U chain, as measured by the Germanium Counting technique.

	$^{210}\text{Pb}_{\text{excess}}$	Units
Argon	15.9 ± 0.5	$\mu\text{Bq/kg}$
Field Cage Resistors	9 ± 1	Bq/kg
PMT Resistors	14 ± 1	Bq/kg
Cold Amp Circuitry	13.8 ± 1.1	Bq/kg
PMT Barrel Connectors	280 ± 90	mBq/kg
Male MCX Connectors	11 ± 1	Bq/kg
Female MCX Connectors	8 ± 1	Bq/kg

Table 3.5: Germanium Counting results measured per unit mass or per piece for ^{60}Co and ^{40}K contamination.

Component	^{60}Co	^{40}K
	[mBq/kg]	
Cryostat Body	13.3 ± 1.0	< 2.5
Cryostat Flanges	11.9 ± 1.0	< 1.8
Organ Pipe Tubes	24 ± 2	< 15
Cryostat Bolts	4.4 ± 0.7	< 15
Cryostat Nuts	< 3.4	< 120
Cryostat Washers	1.3 ± 0.5	< 35
Grid Support Ring	13.3 ± 1.0	< 2.5
PMT Resistors	< 44	4.4 ± 0.8
PMT Capacitors	< 73	< 1.1
PMT Barrel Conn.	< 2.1	130 ± 40
Cold Amp Circuitry	< 6.2	400 ± 100
Male MCX Conn.	38 ± 8	< 210
Female MCX Conn.	< 5.4	< 150
	[mBq/pc]	
Outer Cryostat O-ring	< 0.76	100 ± 20
Field Cage Resistor	< 0.0011	0.019 ± 0.008
PMT Stem	8.8 ± 0.8	79 ± 11
PMT Circuit Board	< 0.001	< 0.051
Cold Amp Circuit Board	< 0.001	< 0.051

Contamination from the uranium and thorium decay chains explains a large fraction of the electromagnetic backgrounds and all of the radiogenic neutron backgrounds. Nevertheless, ^{60}Co and ^{40}K may significantly contribute to the γ -ray backgrounds present in the detector. Germanium Counting measurements for these isotopes are shown in Table 3.5, for the materials for which these measurements are available.

The contamination reported for the PMT stem in these tables come from measurements of the full PMT. The assumption that all of the radioactivity in the PMTs comes from the stem is conservative, since the borosilicate glass has the highest (α, n) neutron yield, and is justified by measurements reported in [113]. In this study, the XENON1T collaboration dissected a Hamamatsu R11410-21 PMT and measured the contamination in each part. The main difference between this model of PMT and the Hamamatsu R11065 PMTs used by DarkSide-50 is the borosilicate glass in the stem. The measurements in [113] show very low contamination in all of the components that the R11065 and R11410-21 PMTs have in common, supporting the hypothesis that the radioactivity in the DarkSide-50 PMTs is concentrated in the stem.

As can be concluded from these measurements, the biggest contributors to the background are the stainless steel in the cryostats, due to their high total mass, the PMT components—especially the borosilicate glass— and Viton o-ring in the outer cryostat flange due to their high levels of contamination.

Chapter 4

Background Sources

Having established what contaminants might be in the detector, the next step is to understand the physics that underlies them as well as what sort of signals they may produce, and how we can eliminate them as best as possible.

This chapter will discuss electromagnetic backgrounds coming from β -decays and γ -rays, as well as surface backgrounds from α -decays and neutron backgrounds from cosmogenic and radiogenic sources.

4.1 Beta-decay Backgrounds

In general, if a nucleus would have a lower ground state energy by changing one neutron into a proton, it will eventually undergo a β -decay to do so. This process of converting a proton to a neutron is described by

$$n \rightarrow p + e^- + \bar{\nu}_e \quad (4.1)$$

Since this interaction involves a flavor change it must be mediated by the weak force through the exchange of a W^\pm boson.

The theory of β -decay is discussed in detail in a series of publications by D. H. Wilkinson [114, 115, 116, 117, 118, 119, 120]. Much of the discussion of β -decay here will follow the formulation presented in this series. Fermi's Golden Rule tells us that the decay rate λ can be given by

$$\lambda = \frac{2\pi}{\hbar} |H_{if}|^2 \frac{dn}{dE} \quad (4.2)$$

where $\frac{dn}{dE}$ is the density of states of final states and $H_{if} = g \int d^3\vec{r} [U_f^* \phi_e(\vec{r}) \phi_n u(\vec{r})] \mathcal{O}_x U_i$ is the matrix element for the transition, g is a coupling constant, U_f and U_i are the final and initial nuclear wave functions, $\phi_e(\vec{r})$ and $\phi_\nu(\vec{r})$ are the wave functions for the outgoing electron and neutrino, and \mathcal{O}_x is an operator describing the transition.

There are five possible interactions that be involved in mediating a β -decay, depending on the spin and parity of the parent and daughter nuclear states:

1. \mathcal{O}_s : Scalar interaction

2. \mathcal{O}_v : Polar vector interaction
3. \mathcal{O}_t : Tensor interaction
4. \mathcal{O}_a : Axial vector interaction
5. \mathcal{O}_p : Pseudoscalar interaction

Since the nuclear wave functions for nuclei with radius R vanish when $r > R$, the low energy nature of the β -decay (being a weak interaction) means that we can consider the electron and neutrino fields around $r = 0$. In the case of the neutrino field, this allows us to take the leading (constant) term of a plane-wave expansion, while in the electron case we have a term accounting for the Coulomb force felt by the electron. The leading order expansion of the Coulomb force in ϕ_e for an electron with energy E is generally described by the Fermi function $F(E, Z)$.

Generally speaking, if the β -decay does not change the nuclear parity or change the spin by more than 1 (at most, the neutrino and electron can each carry a spin of $\pm\frac{1}{2}$), the transition is said to be *allowed*. If this is not the case, the integral of the leading order expansion of these fields will go to zero, and higher order terms will become relevant. In this case, the β -spectrum will be more sensitive to the operators mediating the decay.

If one additional term is needed to make the integral nonzero, the decay is said to be *first forbidden*; if two additional terms are needed, the decay is said to be *second forbidden*, and so on. The matrix element for these transitions is then generally written as the product of a constant term describing the coupling strength ξ and the shape factor $C(E)$, which is typically approximately constant for an allowed decay in which we evaluate the electron and nucleon wave functions for a point-like nucleus.

For a full treatment of the β -decay process, it is necessary to consider the spatial variation of the nucleon and electron wave functions over the nucleus by convolving them with the nuclear charge distribution. Doing so for a uniform spherical distribution, we find that we can approximate the shape factor for vector transitions, $C^V(E, Z)$ by [115],

$$\begin{aligned}
C^V(E, Z) &\approx 1 + C_0^V + C_1^V E + C_{-1}^V/E + C_2^V E^2 \\
C_0^V &= -233(\alpha Z)^2/630 - (E_0 r_N)^2/5 - 6E_0 r_N \alpha Z/35 \\
C_1^V &= -13r_N \alpha Z/35 + 4E_0 r_N^2/15 \\
C_{-1}^V &= 2\gamma E_0 r_N^2/15 + \gamma r_N \alpha Z/70 \\
C_2^V &= -4r_N^2/15
\end{aligned}$$

and we can approximate the shape factor for axial transitions, $C^A(E, Z)$ by [115],

$$\begin{aligned}
C^A(E, Z) &\approx 1 + C_0^A + C_1^A E + C_2^A E^2 \\
C_0^A &= -233(\alpha Z)^2/630 - (E_0 r_N)^2/5 + 2E_0 r_N \alpha Z/35 \\
C_1^A &= -21r_N \alpha Z/35 + 4E_0 r_N^2/9 \\
C_2^A &= -4r_N^2/9
\end{aligned}$$

where $r_N \approx (1.25 \text{ fm})A^{1/3}$ is the nuclear radius.

These factors allow us to write the decay rate λ for a β with kinetic energy T and momentum p as [111]

$$\frac{d\lambda(T)}{dT} = \xi C(E) F(E - V_0, Z) p E (Q_{\beta^\pm} - T)^2 R(E, E_0, M) \quad (4.3)$$

where Q_{β^\pm} is the Q-value for the β^\pm -decay, $E = T + m_e c^2$ and E_0 is the maximum value for the β^\pm -decay, m_e is the mass of an electron, and M is the mass of the recoiling nucleus. $V_0 \approx 1.13\alpha^2 Z^{4/3}$ is a correction to the energy in the Fermi function that accounts for atomic electron screening effects. The term $pE(Q - T)^2$ accounts for the phase space size for the particular decay. $R(E, E_0, M)$ accounts for the phase space factor contribution from the recoiling nucleus, and is typically given by [115],

$$\begin{aligned} R(E, E_0, M) &= 1 + r_0^{V,A} + r_1^{V,A}/E + r_2^{V,A}E + r_3^{V,A}E^2 \\ r_0^V &= \frac{E_0^2}{2M^2} - \frac{11}{6M^2} & r_0^A &= -\frac{2E_0}{3M} - \frac{E_0^2}{6M^2} - \frac{77}{18M^2} \\ r_1^V &= \frac{E_0}{3M^2} & r_1^A &= -\frac{2}{3M} + \frac{7E_0}{9M^2} \\ r_2^V &= \frac{2}{M} - \frac{4E_0}{3M^2} & r_2^A &= \frac{10}{3M} - \frac{28E_0}{9M^2} \\ r_3^V &= \frac{16}{3M^2} & r_3^A &= \frac{88}{9M^2} \end{aligned}$$

where r_i^V and r_i^A are coefficients used for vector and axial decays, respectively. This correction typically changes the phase space factor by $\lesssim 0.1\%$ for β -decays with endpoint energies below 1 MeV.

The Fermi function $F(E, Z)$ is generally given by [20]

$$F(E, Z) = 2(1 + \gamma)(2pr_N)^{2(\gamma-1)} e^{\pi\eta} \left(\frac{|\Gamma(\gamma + i\eta)|}{|\Gamma(1 + 2\gamma)|} \right)^2 \left(\frac{2L_0(E, Z)}{1 + \gamma} \right). \quad (4.4)$$

where $\gamma = \sqrt{1 - (\alpha Z)^2}$ with α being the fine structure constant, $\eta = \pm \frac{\alpha Z E}{pc}$ (positive for a β^- and negative for a β^+), p is the momentum of the β^\pm , and Γ is the complex gamma-function. Without the last term in this equation, the Fermi function written above would represent the Coulomb effects of a point-like nuclear charge evaluated at a distance r_N from the center. The last term in this equation account for the finite charge distribution of the nucleus, where $L_0(E, Z)$ represents the numerically integrated Dirac equation over the nucleus's charge distribution. Values of L_0 for a uniformly spherical charge distribution for various stable nuclei are tabulated in [115].

In the limit where the β is non-relativistic, this can generally be simplified to [121]

$$F(E, Z) = \frac{2\pi\eta}{1 - e^{-2\pi\eta}} \quad (4.5)$$

Figure 4.1 shows the Fermi function evaluated as a function of β^\pm energy for the $Z = 1$ and $A = 3$ and for the $Z = 6$ and $A = 14$ case. In general, the Fermi function




Figure 4.1: Fermi functions showing the Coulomb effects on β spectra as a function of β kinetic energy for β^- emission in a (blue) ${}^3\text{H}$ -like nucleus and a (green) ${}^{14}\text{C}$ -like nucleus, and for β^+ emission in a (purple) ${}^3\text{H}$ -like nucleus and a (orange) ${}^{14}\text{C}$ -like nucleus. Dashed lines show the full Fermi function given in Equation 4.4, and solid lines show the approximation given in Equation 4.5.

enhances the low energy part of the β^- spectrum as the Coulomb force of the nucleus retards the motion of the outgoing β^- , while it suppresses the low energy part of the β^+ spectrum as the nucleus's positive charge repels the outgoing β^+ . The effects of the Fermi function become stronger as Z increases. However, the function tends to level off to a constant value at higher energies. This leveling means that for a β -decay with a high endpoint energy, the Fermi correction will only have a strong effect on the spectral shape at the lower β energies; for a low endpoint energy β -decay, on the other hand, the Fermi function will vary strongly across more of the spectrum. Since ${}^3\text{H}$ has a low endpoint energy of 18.6 keV and the ${}^3\text{H}$ Fermi function flattens out around 20 keV, we therefore expect to see the deformation of the Fermi function more strongly in these decays than we would from ${}^{14}\text{C}$ decays, which have an endpoint energy of 156.6 keV and whose Fermi function flattens out around 100 keV. Indeed, the effects of the Fermi function on the ${}^3\text{H}$ spectrum can be seen in Figure 3.6.

As can be seen in Figure 4.1, the approximation given in Equation 4.5 agrees with the full expression given in Equation 4.4 quite well—even at higher energies where the functions become constant, the two equations generally agree to within 1%. However, the discrepancies are more extreme for higher values of Z , and so for larger nuclei we often need to use Equation 4.4.

In addition to the factors influencing the β -spectrum discussed here, higher order corrections account for inner and outer radiative effects as well as other higher order effects that may further alter the β -spectrum. These effects are discussed in the aforementioned series by Wilkinson, and are typically subdominant to the effects discussed here.

Table 4.1: Selection rules for various β -decay transitions with nuclear spin changes of ΔS and parity changes $\Delta\pi$.

Transition	ΔS	$\Delta\pi$	Operators
Fermi	0	0	$\mathcal{O}_s, \mathcal{O}_v$
Gamow-Teller	$0, \pm 1$ (no $0 \rightarrow 0$)	0	$\mathcal{O}_a, \mathcal{O}_t$
First Forbidden	$0, 1, 2$	1	\mathcal{O}_p , higher order terms
Second Forbidden	$2, 3$	0	higher order terms
Third Forbidden	$3, 4$	1	higher order terms

4.1.1 Forbidden Transitions

As alluded to above, some β decays are classified as being “forbidden” and typically have longer half-lives with a distorted spectrum. In general, allowed transitions are classified as either Fermi transition or Gamow-Teller transition depending on the nature of the decay.

Table 4.1 summarizes the selection rules for the various kinds of β -decay transitions. Fermi transitions typically have the shortest half-lives and are often referred to as “super allowed” transitions.

For the most part, allowed transitions are ones in which the lowest order term in the β^\pm and neutrino wave functions dominate, and each level of forbiddenness requires the next leading order term. However, it should be noted that for some nuclei, higher order terms may still contribute significantly to the decay.

As discussed in Chapter 3.1.6, the decay of ^{14}C is a $0+ \rightarrow 1+$ transition, and therefore satisfies the selection rules for being a Gamow-Teller transition. However, minor distortions are seen in the spectral shape and the half-life is 5700 years, whereas Gamow-Teller transitions typically have half-lives on scales ranging from minutes to a couple years. It is generally believed that ^{14}C behaves like a second forbidden transition due to the nuclear structure of the ^{14}C and ^{14}N nuclei suppressing the Gamow-Teller matrix element. Gamow-Teller transitions are generally carried out by axial-vector interactions, however it has been noted that this matrix element is anomalously small for this decay. Effects from weak-magnetism, a higher vector interaction term in the weak force, may therefore be able to explain the long half-life and distortion seen in the ^{14}C β -spectrum [111]. Where there is some discrepancy among the various measurements of the spectrum, there is some evidence to support the hypothesis that the weak-magnetic current plays an important role in the ^{14}C decay.

4.1.2 Cherenkov Backgrounds

Many of the β -emitters that are present in DarkSide-50 are located in volumes where the β s themselves are unlikely to directly produce a signal. One exception could arise if the decay occurs in a transparent medium with a high enough index of refraction for the β to produce Cherenkov light, and if the β is accompanied by a γ -ray that scatters in the TPC. Normally such events would not provide a concerning background to

an experiment like DarkSide-50, which can use PSD to eliminate backgrounds from these γ -rays. However, in this case, the Cherenkov light produced by the β is very fast, and this light may add a fast component to the pulse shape of the γ -ray, making it look like a nuclear recoil.

In the DarkSide-50 analysis, we introduce an additional cut to remove events where the S1 signal has an unusually high concentration in a single PMT, as one might expect from a Cherenkov event happening in the fused silica. It is likely that such events in the Teflon and fused silica can be further reduced by making cuts on the scintillation pulse shape, as the Cherenkov light will have a much faster time profile than the fast component of the liquid argon scintillation. If further studies find that these two cuts are insufficient for removing these events with the needed efficiency, we can likely increase the analysis energy threshold of the dark matter search to further suppress them.

In order to fully understand these backgrounds more rigorously, detailed optical simulations need to be performed. However, this analysis goes beyond the scope of this document. G. Koh and X. Xiang are currently working on a much more detailed analysis of this background and the efficiency with which we can remove it.

4.2 Gamma-ray Backgrounds

In Chapter 2.3, we discussed the physics underlying LAr scintillation and saw that only charged particles directly produce scintillation light. However, uncharged particles may still interact with atoms to indirectly produce light. This is the case for γ -rays, which may scatter off of electrons. These electrons will then scintillate as they slow down in the LAr.

Because these electron recoils are easily removed by pulse shape discrimination, we do not expect a very dominant background from them. However, we do expect many γ -ray backgrounds to have a high enough rate that they will produce a clearly visible signal in both the TPC and LSV. By studying these γ -ray signals, then, we can calibrate the time offset between the inner and outer detectors of DarkSide-50, and we can use these signals to validate our background model. Understanding these signals is therefore a useful step in running a background-free experiment.

4.2.1 Types of X-rays and Gamma-rays

In DarkSide-50, γ -rays come from nuclear transitions after some sort of nuclear decay or reaction. X-rays, on the other hand, typically come from lower energy atomic transitions as atomic electrons relax to lower energy levels. γ -rays coming from nuclear relaxation tend to have higher energies than x-rays from electronic transitions; γ -rays are typically on the scale of 100–1000 keV, while x-rays are usually on the 1–100 keV scale.

Nuclear de-excitations that produce γ -rays typically follow β -decays that go to an excited state of the daughter nucleus. These γ -rays can also be produced by α -decays,

Table 4.2: Selection rules for γ -ray decays for which the nuclear angular momentum changes by ΔL and the parity changes by $\Delta\pi$.

Transition		ΔL	$\Delta\pi$
Type	Multipole		
Electric	Dipole	1	1
Magnetic	Dipole	1	0
Electric	Quadrupole	2	0
Magnetic	Quadrupole	2	1
Electric	Octupole	3	1
Magnetic	Octupole	3	0

as is the case for the ^{235}U decay, though since α -decays are more likely to go directly to the daughter nucleus's ground state, these transitions are less common.

When a nucleus is in an excited state, it can decay by γ -ray emission to a lower energy level depending on the magnetic and electric charge structure of the parent and daughter nuclei. Nuclear transitions by γ -ray emission are generally characterized as either being electric or magnetic and by some multipole moment mediating the transition. The selection rules for these various kinds of decays are summarized in Table 4.2. In general, the ΔL term refers to the angular momentum carried out by the emitted γ -ray. However, since the angular momentum is a vector, a decay with a given ΔL may correspond to a change in nuclear states where the angular momentum changes by less than ΔL . For example, γ -ray decays may leave the angular momentum of the nucleus unchanged. The exception is that $0 \rightarrow 0$ transitions are not allowed unless an $e^+ - e^-$ pair is also created. This would correspond to a monopole transition, but since monopoles do not produce electromagnetic radiation, such a transition cannot occur without such a pair-creation process accompanying it.

Typically, the higher the multipole moment of a γ -ray decay, the less likely the transition is to happen. The effect of these selection rules is that some nuclear de-excitations will preferentially go straight to the ground state, releasing one high energy nucleus, while other decays may involve multiple stages before the nucleus reaches the ground state.

The timescale for these transitions is typically on the order of picoseconds, so nuclear decays that result in excited nuclear states will usually see all of their γ -rays emitted in a prompt coincidence with the decay. This effect is pronounced in ^{60}Co , as seen in Figure 3.3, where the β -decay to the 4+ excited state of ^{60}Ni preferentially decays to the 2+ excited state, which then decays to the ground state, emitting two γ -rays in prompt coincidence with each other.

Another source of γ -rays, which typically have lower energies, can come from internal conversion. Internal conversion occurs when an excited nucleus de-excites by directly giving the excitation energy to an inner shell electron. This process can happen when the electron orbitals overlap with the nuclear wave function. When nuclei de-excite via internal conversion, the energy can be carried away by multiple electrons, as long as the electrons' ionization and kinetic energy all add up to the

energy of the transition. Since this process does not involve the creation of a γ -ray at this stage, the $0 \rightarrow 0$ transition can take place through this mechanism.

While the internal conversion process itself does not produce a γ -ray, it leaves the nucleus in an excited state, with some of the lower atomic orbitals left unfilled. The atom will then de-excite as electrons from higher levels fall down to fill the holes produced by the internal conversion. Each of these atomic transitions will produce an x-ray or an Auger electron. As a result, internal conversions tend to be accompanied with a cascade of x-rays, which may be seen as low energy electron recoils in a detector.

4.2.2 Gamma-ray Sources

In order to understand the signals seen in the TPC and neutron veto, it is helpful to be able to identify the sources of various γ -ray signals. Tables 4.3 and 4.4 list the various γ -rays and their sources. Since the ^{232}Th , ^{238}U , and ^{235}U decay chains tend to have a very large number of x-rays that will largely be blocked by the cryostat materials, we restrict this table to γ -ray lines with energy over 100 keV and with intensities over 1%.

Tables 4.3 and 4.4 also list correlated γ -rays, along with the correlation probabilities. The correlation probability tells you the probability of seeing a particular correlated γ -ray given that you have seen the given γ -ray. Correlations generally arise if a γ -ray line corresponds to the nucleus de-exciting to another (lower) excited state, which can then decay, or if other (higher) excited states can decay to the given state. Many of the nuclei here, especially those with higher A , have a wide array of isomers that can all decay to other ones, and often this creates large cascades of coincident γ -rays. Since listing all of these correlations in a table would be impractical, we mark isotopes with such cascades with a \mathcal{C} under the “Correlated lines” column.

These γ -rays will primarily produce signals by Compton scattering off of atomic electrons, making an electron recoil. Since electron recoil backgrounds are very efficiently removed by pulse shape discrimination, γ -rays do not provide very concerning backgrounds. However, these lines provide a way of measuring the concentration of the various contaminants in the materials used in DarkSide-50, making them very important for quantifying the concentrations of the various radioisotopes and building a background model. These lines were used during the material selection phase of the experiment, where these γ -rays were counted in order to determine the contamination of the various radioisotopes present, and they can also be observed in the LSV or neutron veto energy spectra to verify the background model and reduce the systematic errors.

4.2.3 Cherenkov Signals

Much like we saw for β -decays, γ -rays can produce signals with coincident electron recoils in the liquid argon and Cherenkov light in the Teflon or fused silica, which may add together to have a pulse shape resembling a nuclear recoil.

A full analysis of these backgrounds would require a full Monte Carlo simulation to determine what fraction of all such γ -rays scatter in either the Teflon or fused

Table 4.3: γ -ray lines in DarkSide-50. Energies are in keV. Intensities relative to the top of the decay chain and correlation probabilities (in parentheses) relative to γ -ray line intensity are in percent. Only γ -ray lines with intensities greater than 1% and energies over 100 keV are included. Isotopes marked with a * have many γ -ray, and only the three highest intensity lines are included. Correlated γ -rays are only listed if they have a probability greater than 1%. γ -ray lines that are part of a cascade with many lines are labeled with \mathcal{C} , and only correlations with listed lines are explicitly included.

Chain	Isotope	Energy	Intensity	Correlated lines
^{60}Co	^{60}Co	347.1	0.0075	826.1 (86.4), 1332.5 (86.4), 2158.6 (13.4)
	^{60}Co	826.1	0.0076	347.1 (100), 1332.5 (100)
	^{60}Co	1173.2	99.85	1332.5 (100)
	^{60}Co	1332.5	99.9826	1173.2 (99.87)
	^{60}Co	2158.6	0.0012	347.1 (100)
	^{60}Co	2505.7	2.0×10^{-6}	—
^{40}K	^{40}K	511.0	0.0020	511 (100)
	^{40}K	1460.8	10.7	—
^{232}Th	$^{228}\text{Ac}^*$	338.3	11.3	\mathcal{C}
	$^{228}\text{Ac}^*$	911.2	25.8	\mathcal{C}
	$^{228}\text{Ac}^*$	969.0	15.8	\mathcal{C}
	^{224}Ra	241.0	4.1	—
	^{224}Ra	292.7	0.0062	241.0 (100)
	^{224}Ra	404.2	0.0022	241.0 (100)
	^{224}Ra	422.0	0.0030	241.0 (100)
	^{224}Ra	645.5	0.0054	—
	^{212}Pb	115.2	0.60	300.1 (100)
	^{212}Pb	176.7	0.052	238.6 (100)
	^{212}Pb	238.6	43.6	—
	^{212}Pb	300.1	3.3	115.2 (100)
	^{212}Pb	415.2	0.013	—
	$^{212}\text{Bi}^*$	727.3	4.3	\mathcal{C} : 785.4 (16.9)
	$^{212}\text{Bi}^*$	785.4	0.71	727.3 (100)
	$^{212}\text{Bi}^*$	1620.5	0.96	—
	$^{208}\text{Tl}^*$	510.8	8.1	\mathcal{C} : 583.2 (100), 2614.5 (100)
	$^{208}\text{Tl}^*$	583.2	30.5	\mathcal{C} : 510.8 (29.0), 2614.5 (100)
	$^{208}\text{Tl}^*$	2614.5	35.8	\mathcal{C} : 510.8 (13.5), 583.2 (85.0)

Table 4.4: γ -ray lines in DarkSide-50. Energies are in keV. Intensities relative to the top of the decay chain, and correlation probabilities (in parentheses) relative to γ -ray line intensity are in percent. Only γ -ray lines with intensities greater than 1% and energies over 100 keV are included. Isotopes marked with a * have many γ -ray, and only the three highest intensity lines are included. Correlated γ -rays are only listed if they have a probability greater than 1%. γ -ray lines that are part of a cascade with many lines are labeled with \mathcal{C} , and only correlations with listed lines are explicitly included.

Chain	Isotope	Energy	Intensity	Correlated lines
^{238}U	$^{234}\text{Pa}^*$	131.3	18.9	\mathcal{C}
	$^{234}\text{Pa}^*$	883.2	10.0	\mathcal{C}
	$^{234}\text{Pa}^*$	946.0	14.0	\mathcal{C}
	^{226}Ra	186.2	3.6	—
	^{226}Ra	262.3	0.005	186.2 (100)
	^{226}Ra	414.6	3.0×10^{-4}	186.2 (100)
	^{226}Ra	449.4	1.9×10^{-4}	186.2 (100)
	^{226}Ra	600.7	4.9×10^{-4}	—
	$^{214}\text{Pb}^*$	242.0	7.3	\mathcal{C}
	$^{214}\text{Pb}^*$	295.2	18.4	\mathcal{C}
	$^{214}\text{Pb}^*$	351.9	35.6	\mathcal{C}
	$^{214}\text{Bi}^*$	609.3	45.5	\mathcal{C}
	$^{214}\text{Bi}^*$	1120.3	14.9	\mathcal{C}
	$^{214}\text{Bi}^*$	1764.5	15.3	\mathcal{C}
^{235}U	$^{235}\text{U}^*$	143.8	11.0	\mathcal{C}
	$^{235}\text{U}^*$	163.4	5.1	\mathcal{C}
	$^{235}\text{U}^*$	185.7	57.0	\mathcal{C}
	$^{231}\text{Pa}^*$	283.7	1.7	\mathcal{C}
	$^{231}\text{Pa}^*$	300.1	2.4	\mathcal{C}
	$^{231}\text{Pa}^*$	302.7	2.3	\mathcal{C}
	$^{223}\text{Ra}^*$	154.2	5.7	\mathcal{C}
	$^{223}\text{Ra}^*$	269.5	13.9	\mathcal{C}
	$^{223}\text{Ra}^*$	323.9	4.0	\mathcal{C}
	^{219}Rn	271.2	10.8	\mathcal{C}
	^{219}Rn	401.8	6.6	\mathcal{C}
	^{211}Pb	404.9	3.8	\mathcal{C} : 427.1 (55.3)
	^{211}Pb	427.1	1.8	\mathcal{C} : 404.9 (100)
	^{211}Pb	832.0	3.5	—
	^{211}Bi	351.1	13.0	—

silica with enough energy to produce Cherenkov light and then scatter in the TPC exactly once with little enough energy to not overpower the Cherenkov signal. Much like we discussed in Section 4.1, we expect that a combination of fiducial, multiple scattering, pulse shape, energy, and S1 distribution cuts will effectively eliminate these backgrounds. However, this analysis is still a work in progress. While G. Koh and X. Xiang are investigating these backgrounds in much greater detail, we will not discuss them further in this document.

4.3 Alpha-decay and Surface Backgrounds

Contamination from heavy nuclei can lead to signals from α -decays in the detector. Some of these heavy nuclei may be found in higher concentrations on the inner surfaces of the detector; these nuclei can give rise to surface backgrounds. The backgrounds that a detector sees from these decays can result directly from signals produced by the emitted α particles, or by the recoiling daughter nucleus after the decay. Understanding the α -decays in the experiment and the signals they may produce in the detector is therefore an important part in understanding the experiment’s backgrounds. In particular, these signals often appear as nuclear recoils, which may produce WIMP-like signals inside of a detector.

4.3.1 The Alpha-decay Process

Many nuclei with $Z > 82$ can become more stable by emitting an α particle. This process is ultimately a competition between the Coulomb repulsion of the α and the daughter nucleus and the potential well created by the nuclear binding force. If the electrostatic repulsion of the daughter nucleus overcomes the binding potential at some distance from the nucleus, the α can tunnel through the potential barrier and escape the nucleus, typically with a kinetic energy below the daughter’s Coulomb barrier. Typical α energies are therefore less than 10 MeV—usually ~ 5 MeV.

This process is more common in heavier nuclei, and decreases the atomic number by two and the mass number by four. Generally speaking, if the parent nucleus has an even number of protons and neutrons, it will decay directly to the ground state of the daughter nucleus. In the shell model of the nucleus, this just means that the emitted α will carry out nucleons from the highest shell, which leaves the daughter nucleus in the ground state. However, since α -decay removes paired protons and neutrons, nuclei with odd numbers of protons or neutrons will be left with the unpaired nucleon in an excited state of the daughter nucleus. These excited daughter nuclei will then decay down to their ground state, releasing a γ -ray on a timescale typically on the order of picoseconds and therefore in prompt coincidence with the α .

While α particles do not have any intrinsic spin or angular momentum, they can carry orbital angular momentum away from the nucleus. If the decay process preserves nuclear parity, the decay will change the angular momentum by an even amount; if it changes the parity, it will carry away an odd orbital momentum. The potential felt by the α increases in strength with its orbital angular momentum, and so the

more angular momentum carried away by the α , the more the decay process will be suppressed. Additionally, α -decays of nuclei with odd proton or neutron numbers are further suppressed, and such nuclei tend to have half-lives two to three orders of magnitude longer than those of even-even nuclei.

When an α -decay occurs with Q-value Q , some of the energy is carried away as kinetic energy of the α , and the rest by the recoiling daughter nucleus. The kinetic energy of the α , T_α and of the daughter nucleus T_{A-4} are given by

$$\begin{aligned} T_\alpha &= \frac{A-4}{A}Q \\ T_{A-4} &= \frac{4}{A}Q \end{aligned} \tag{4.6}$$

where A is the mass number of the parent nucleus.

4.3.2 Alpha Backgrounds

Due to their high mass and consequent high stopping power in most materials, α particles tend to travel very short distances. In most materials, a 5 MeV α will travel 1–10 μm [122]. This short distance means that most α -decays that take place in the various detector materials will not leave a signal in the TPC, with the exception of those that produce γ -rays, as shown in Figures 3.9, 3.11, and 3.14.

However, there are three ways in which an α can directly produce a signal in the TPC. Trace amounts of radon, especially ^{222}Rn and its daughters, are likely to get mixed in with the liquid argon itself. This can happen if trace amounts of air gets mixed in with the argon and the purification process leaves residual radon daughters behind. Alternatively, the pipes that make up all of the liquid argon fluid handling systems used in DarkSide-50 have been exposed to air, which means their inner surfaces likely have some contamination from radon daughters. Liquid argon flowing through the pipes can pick up contamination from these daughters and carry it into the bulk during recirculation. The third way α -emitters may contaminate the liquid argon is if uranium and thorium daughters near the inner surfaces of the various materials the argon comes into contact with decay and eject a nucleus into the argon.

Studying DarkSide-50 data, C. Stanford has identified various α -emitters from the lower ^{238}U decay chain and the ^{232}Th decay chain in the liquid argon. Measurements of the lower ^{238}U chain were made at ^{222}Rn , ^{218}Po , ^{214}Po , and ^{210}Po , and measurements of the ^{232}Th chain were made at ^{220}Rn and ^{216}Po . The results from these measurements, averaging together isotopes in the same chain, are shown in Table 4.5. In addition to seeing the ^{238}U chain split between ^{226}Ra and ^{222}Rn , we see an excess of ^{210}Po . This is likely due to ^{210}Pb naturally present on the surfaces exposed to the liquid argon mixing with the argon. We therefore refer to the $^{210}\text{Pb}_{\text{excess}}$ as the activity of ^{210}Po and its daughters, as measured by ^{210}Po , above the lower ^{238}U chain.

Since the α nucleus is very light compared to ^{40}Ar , and the α s involved in these decays have high energies around 5 MeV, we expect α particles to lose most of their

Table 4.5: Uranium and thorium contamination measurements in the DarkSide-50 underground liquid argon as measured by C. Stanford.

Chain	Concentration [$\mu\text{Bq/kg}$]	\pm
$^{238}\text{U}_{\text{upper}}$	0	0
$^{238}\text{U}_{\text{lower}}$	1.31	0.01
$^{210}\text{Pb}_{\text{excess}}$	15.9	0.5
^{235}U	0	0
^{232}Th	0.17	0.04

energy by exciting and ionizing atomic argon electrons. Indeed, this is verified by simulations in SRIM [96] and the NIST ASTAR database [122], where we see the electronic stopping power is over an order of magnitude greater than the nuclear stopping power down to ~ 20 keV, and by 3 orders of magnitude at 5 MeV. This means we expect the scintillation light from these α s to be quenched very little.

Given the ~ 5 MeV energy scale of these α -decays, we therefore expect the visible energy from these decays to be far above the DarkSide-50 WIMP search energy range, and so these backgrounds are unlikely to be a problem.

4.3.3 Surface Backgrounds

The other place where we may see α -decays is on the inner surfaces of the TPC. Radium naturally present in uranium and thorium in rocks underground decay to radon. Since radon is a gas, it can diffuse into the air. When this radon decays, it can deposit its daughters on surfaces exposed to the air, and we therefore expect a higher concentration of the radioisotopes following radon on the surfaces of materials.

For ^{232}Th and ^{235}U , all of the daughters below radon in the decay chain have very short half-lives, and are therefore likely to decay away too quickly to be a concern for the experiment. However, in the ^{238}U decay chain, ^{210}Pb has a half-life of 22.2 years. We therefore expect ^{210}Pb to accumulate on these surfaces and to eventually be in secular equilibrium with the rest of the chain below it. Since we know that pulse shape discrimination can easily eliminate the β backgrounds on the surfaces, we are primarily concerned with those that α -decay. This includes ^{210}Pb , ^{210}Bi , and ^{210}Po . However, ^{210}Pb only α -decays with a branching ratio of $2 \times 10^{-6}\%$, and ^{210}Bi with a branching ratio of $1 \times 10^{-5}\%$, so we expect the backgrounds from ^{210}Po decays to dominate.

^{210}Po α -decays with a Q-value of 5.407 MeV. According to Equation 4.6, we therefore expect the α to carry away 5.304 MeV of kinetic energy, with the remaining 103 keV being carried by the daughter ^{206}Pb nucleus. This prediction is consistent with the measurements reported in [123], which also reports a branch in which the α carries away 4.516 MeV with a branching ratio of $1 \times 10^{-3}\%$.

Much like in the case for bulk α -decays, we expect the high energy α produced in this reaction to produce so much visible energy that it is outside the WIMP search energy range. However, since these decays are isotropic and the α and recoil nucleus




Figure 4.2: A drawing of the surface background of interest to DarkSide-50, in which a ^{210}Po nucleus decays and sends a recoiling ^{206}Pb nucleus into the liquid argon and an α into the wavelength shifter (TPB for DarkSide-50), which may scintillate as well.

have opposite momenta. This means that in half of all α -decays, the daughter will be sent recoiling into the TPC, while the α goes into the surface it is deposited on. An illustration of this process is shown in Figure 4.2.

In order to assess the risk of these backgrounds, we need to know the visible energy of the ^{206}Pb recoil signal. Since ^{206}Pb is a nucleus, the energy it deposits will look like a nuclear recoil to the TPC, and so we cannot rely on the standard pulse shape discrimination methods to reject these events. As discussed in Chapter 2, we can normally predict the quenching factors for nuclear recoils by calculating the Lindhard-Birks factor. However, the standard treatment by Lindhard and Schraff to calculate Lindhard factors [91] is only valid when the project nucleus is the same as the target nucleus. Additionally, at the low energies relevant to these decays, $\epsilon \lesssim 0.01$ in Equation 2.1, meaning that the underlying theory fails to describe these interactions well.

The SRIM software [96] can generally be used in cases where the projectile and target nuclei are different, as it simulates the entire cascade. However, even in these cases, the basic theory on which it is based is known to break down for ions with low velocities, as reported by Paul [124]. Additional errors at low ion velocities likely arise from the fact that SRIM models do not account for atomic excitations and shell correction, as described by Sigmund and Schinner [125], or for the effects of atomic binding energies, as described by Sorensen [95]. As shown in [95], it is likely that these effects cause the SRIM model to under-predict the energy lost to heat in these interaction.

SRIM simulations indicate that a third of the energy deposited by a 103 keV ^{206}Pb nucleus in liquid argon will be dissipated as heat, and that nearly all of the energy that goes into ionization comes from the cascading argon ions rather than directly from the ^{206}Pb nucleus itself. This effect means that ^{206}Pb should qualitatively look like the sum of many lower energy argon nuclear recoils, and so we expect f_{90} to be nuclear recoil-like.




Figure 4.3: (*Left*) A schematic of the detector used to measure surface background scintillation (courtesy of C. Stanford). (*Right*) A photograph of the detector.

However, given the aforementioned inaccuracies in SRIM at these low velocities, these calculations cannot be taken at face value. Furthermore, there are no pre-existing measurements of the visible energy produced by a ^{210}Po surface background on which to base an assessment of these backgrounds. Instead, a full understanding of the surface background requires an independent measurement.

Unfortunately, the low kinetic energy of the ^{206}Pb nucleus makes it hard to identify energy depositions from this isotope above the noise. However, the earlier decays in the ^{238}U decay chain offer an alternative option: ^{214}Bi β -decays to ^{214}Po , which then α -decays to ^{210}Pb with a half-life of $164\ \mu\text{s}$. This relatively short half-life means that coincidences between the ^{214}Bi and ^{214}Po can be identified in order to make the signal from the ^{210}Pb scintillation visible above the background.

^{214}Po α -decays with a Q-value of 7.833 MeV. According to Equation 4.6, this means that the α will have a kinetic energy of 7.687 MeV, while the recoiling ^{206}Pb nucleus will have 146 keV. This energy is higher than the energy of the recoiling ^{210}Pb nucleus, but it is close enough to give us a meaningful estimate of the quenching factor.

The Detector and Lead Nuclear Recoil Quenching Measurement

In order to measure the signal from this ^{210}Pb nuclear recoil, we¹ designed a small single phase liquid argon detector. Figure 4.3 shows a schematic drawing of this detector. The detector consists of a cup of the Spectralon reflector. The Spectralon

¹This work was done in close collaboration with J. Xu and C. Stanford.

cup is viewed from the top by a quartz window, which separates the contents of the cup from a Hamamatsu R11065 PMT. The quartz window maintains a vacuum seal, separating the contents of the cup from the steel vessel containing the PMT. Teflon gaskets maintain this seal, as a fastenable steel ring holds the windows and gaskets tight against the Conflat flange. The PMT sits in the middle of this ring. The Spectralon cup sits snugly inside of a steel cup, and a Gas Feed Tube enters this cup from the bottom.

This steel cup is welded to the inner surface of a bore hole in a Conflat flange. This flange meets with a similar one that connects to the steel vessel containing the PMT. The interface between the flanges is maintained by a copper gasket. This vessel is typically filled with nitrogen gas during operation in order to maintain favorable operating conditions for the PMT.

A Teflon ring (not shown) sits around the base of the neck of the PMT, and a series of springs connect the Teflon ring to the steel ring above the quartz window. This connection helps hold the PMT in place so that it does not move during transportation.

A long, narrow steel tube sticks out of the steel vessel that contains the PMT. At the top of this neck are signal and high voltage feedthroughs. The purpose of the long tube is to separate the feedthroughs from the vessel, so that they can stay relatively warm while the Spectralon cup is held at cryogenic temperatures. A pair of cables connect the PMT bases to the high voltage feedthroughs.

Samples that we wish to measure are placed inside of the Spectralon cup when this setup is disassembled.

In order to measure these ^{214}Po decays, we deposited ^{222}Rn daughters onto a thin rectangular VUV mirror. The VUV mirror was made of a highly reflective aluminum layer on a quartz substrate. To protect the aluminum from oxidizing and losing its reflectance, the mirror was also coated in a 25 ± 10 nm layer of MgF_2 . Although we have observed that MgF_2 scintillates under α excitations with a decay time over $100 \mu\text{s}$, SRIM simulations indicate that a 7.8 MeV α will deposit far under 1 keV in such a short distance. Additionally, since we expect the slow component of the argon scintillation to have a decay time of $1.55 \mu\text{s}$, the timescales we are interested in are too small to see much of the MgF_2 scintillation light. We therefore conclude that the background from MgF_2 scintillation is likely to be negligible.

To deposit radon daughters onto the VUV mirror, we placed the mirror inside a cylindrical vacuum-sealed chamber, as shown in Figure 4.4. This chamber was connected to a ^{226}Ra source, which was connected to an argon gas tank. By flowing argon gas through the ^{226}Ra and into the chamber, we could transport ^{222}Rn from decayed ^{226}Ra into the chamber and near the mirror. After about three hours of exposure like this, we found that the mirror had accumulated ~ 2 MBq of ^{222}Rn daughters deposited onto it. The nitrogen flow was then turned off and the ^{226}Ra source was disconnected from the deposition chamber. Remaining ^{222}Rn was removed from the deposition chamber before opening it by opening the valve the connected the chamber to the charcoal trap. The charcoal trap was submerged into a liquid nitrogen bath, causing all of the air to get pulled into the trap.




Figure 4.4: The setup used to deposit radon daughters on samples.

The mirror was then inserted into the vessel shown in Figure 4.3, and the vessel was then evacuated and back-filled with argon gas for several cycles to remove impurities and cooled down to ~ 87 K inside a liquid argon bath before being filled with purified liquid argon.

The longest-lived precursor to ^{214}Bi in the ^{222}Rn decay chain is ^{214}Pb , with a half-life of 26.8 minutes. We therefore had to complete these measurements in a few hours before the ^{214}Bi levels depleted to levels too low for us to measure.

After filling the vessel with liquid argon, ^{214}Po decays were identified by looking for a coincidence with a ^{214}Bi scintillation signal within a timescale of a couple hundred microseconds.

Analysis and Results

The results of this plot are shown in Figure 4.5. In order to obtain this energy spectrum, three corrections had to be made to account for systematics of the experiment.

First, we had accounted for impurities in the liquid argon that might suppress the scintillation light. Since the long lifetime of the triplet state makes it more likely to be quenched by impurities, we were able to fit an exponential to the tail of the liquid argon scintillation pulse to measure the triplet state and use this measurement as a metric for the impurity quenching effects. Using an expected triplet lifetime of $1.55\ \mu\text{s}$ as a benchmark, we were able to correct the spectrum for this effect.

Second, the VUV mirror is not perfectly reflective, and so some amount of scintillation light was likely lost after being absorbed by the mirror. Since the ^{210}Pb scintillation events happened very close to the mirror, this could be a significant effect. At the 128 nm wavelength of liquid argon scintillation, the mirror has a re-




Figure 4.5: Measured visible energy spectrum of ^{214}Po α -decays in liquid argon. The high energy peak is from ^{214}Po α scintillation, while the lower energy peak is from ^{210}Pb scintillation. The inset plot shows visible energy vs. f_{90} for these events. At lower energies, we see a tail leading down to lower f_{90} , which we conclude is from decays of ^{214}Po embedded in the substrate. We therefore require that ^{210}Pb events be in the red box shown in this inset plot to appear in the mean energy spectrum.

flectance over 85%. Since about half of the scintillation light had to reflect off of the mirror, a correction factor was applied to account for light lost to the mirror.

Third, ^{218}Po decays earlier in the ^{222}Rn chain than ^{214}Bi with a Q-value of 6.115 MeV. If these decays send the α into the argon, the recoiling ^{214}Pb nucleus can become embedded into the substrate. ^{214}Bi nuclei from these embedded ^{214}Pb decays will therefore be below the surface of the substrate. When their daughter ^{214}Po nuclei decay and send ^{210}Pb nuclei into the argon, these ^{210}Pb nuclei will lose some of their energy to the substrate before they reach the liquid argon. This causes the ^{210}Pb to have a tail going down to low energies that can obfuscate the peak corresponding to these nuclei depositing all of their energy into the argon. However, since ^{210}Pb nuclei scintillate primarily by dividing their kinetic energy among multiple lower energy ^{40}Ar nuclei, and f_{90} for ^{40}Ar recoils decreases at lower energy, we find that the tail from these embedded nuclei also has a much lower f_{90} . The f_{90} vs. visible energy distribution is shown in the inset in Figure 4.5, and a cut was made to select events with f_{90} between 0.74–0.9, to select events within 2σ above the peak of the f_{90} distribution.

The result of these corrections is that we saw the 146 keV ^{210}Pb recoil peak at 45.7 ± 2.0 PE in our small detector. Since different detectors have different light collection efficiencies, it is useful to normalize this result to the light yield of electron recoils in the detector. Using an ^{241}Am source, which produces a 59.5 keV γ -ray, we




Figure 4.6: Molecular diagrams of (*left*) tetraphenyl butadiene (TPB) and (*right*) p-terphenyl (pTP).

measured a light yield of 6.15 ± 0.1 PE/keV. For comparison, DarkSide-50 has measured a light yield at null field of 8.1 ± 0.2 PE/keV [42].

Dividing the number of photoelectrons we saw for ^{210}Pb recoils by the light yield of the detector normalizes the visible energy of these recoils to 7.4 ± 0.3 keV_{ee}, corresponding to a quenching factor of $(19.7 \pm 0.9)^{-1}$. For comparison, the model presented by SCENE in [92] based on the Lindhard-Birks quenching model predicts a quenching factor of $\sim 3.2^{-1}$ for 146 keV ^{40}Ar nuclei recoiling in liquid argon, a factor of ~ 6 less heavily quenched than we see with ^{210}Pb scintillation. Simulations using SRIM indicate that ^{210}Pb has a stopping power about 5 times higher than that of ^{40}Ar at this energy, and so this increased quenching effect at least makes sense at an intuitive level, due to the dependence of the quenching factor on the stopping power. However, attempting to extrapolate this Lindhard-Birks model to ^{210}Pb under-predicts the effects of quenching by a factor of ~ 3 , showing that the usual models do not hold for these low energies.

If we assume that ^{206}Pb has the same quenching factor as ^{210}Pb , we would expect to see the signal from the 103 keV ^{210}Pb recoil at ~ 5 keV_{ee}. This is likely a decent assumption, because the stopping power of ^{206}Pb at 103 keV is close to that of 146 keV ^{210}Pb , and due to the lower energy of the ^{206}Pb nucleus, any deviations in the quenching factor between these two nuclei would likely suppress the signal from ^{206}Pb even more. This low energy 5 keV_{ee} signal is far below the energy region of interest for DarkSide-50.

Wavelength Shifter Scintillation Measurements

While the ^{206}Pb scintillation in LAr is likely too low energy to provide a hazardous background itself, Figure 4.2 shows that whenever the ^{210}Pb nucleus enters the LAr, the α will deposit energy in the wavelength shifter. In order to understand this signal, we must therefore also understand the signal produced by α scintillation in the wavelength shifter.




Figure 4.7: An illustration of the scintillation process of organic scintillators.

In these studies, we consider two wavelength shifters: para-terphenyl, or p-terphenyl (pTP), and tetraphenyl butadiene (TPB). The molecular structures of these molecules are shown in Figure 4.6. As can be seen in these diagrams, both molecules are made of a series of benzenoid rings. This structure is common among organic scintillators; the π bonds that connect the carbon atoms create a series of rings in which electrons can flow freely. When these rings are excited by some incident particle, excitations form in these rings corresponding to the normal modes of the wave functions of the rings.

An illustration of this scintillation process is shown in Figure 4.7. Standing waves in the π bond rings of these molecules carry some angular momentum. At a given energy level, there are degenerate modes that have angular momenta in opposite directions; the wave function corresponding to a given energy level is a linear combination of these standing waves. Depending on the linear combination of waves, the state may have a total angular momentum of 0 or 1. The former case constitutes a singlet state while the latter is a triplet state. Due to the different angular momenta of the singlet and triplet states, they tend to form different energy levels.

Within each energy level of the wavelength shifter molecule is a continuum of closely spaced vibrational energy levels. The spacing between these vibrational states varies between energy levels.

If an incident photon has energy close to one of the excited states of the wavelength shifter molecule, the molecule may absorb the photon and transition to one of the vibrational levels of its excited state, typically the first excited state. The range of photon energies that the molecule can absorb this way define the absorption spectrum of the wavelength shifter. Since the molecule's ground state has zero angular momentum, these transitions will take the molecule to the first excited singlet state.

After the molecule absorbs a photon, it will undergo vibrational relaxation to the lowest vibrational state of this state. This process is non-radiative and happens on a sub-nanosecond timescale. The molecule will then relax back down to the ground state. This can occur radiatively, through the process of fluorescence, or non-radiatively through some sort of quenched relaxation process, dissipating the excitation energy as heat. Fluorescence of the singlet state tends to take place on the timescale of ~ 11 ns [126].

When the wavelength shifter molecule fluoresces, it will decay down to one of the vibrational levels of the ground state. The photon emitted by this process will have an energy corresponding to the energy difference between the lowest vibrational level of the first excited state and one of the vibrational levels of the ground state. The set of photon energies accessible to this transition define the emission spectrum of the wavelength shifter. Because the vibrational levels of the ground state and excited states have different spacings, the absorption and emission spectra will be different. This difference is known as the Stokes shift; the result of this effect is that photons absorbed by the wavelength shifter are re-emitted at a longer wavelength.

Figure 4.8 shows the emission spectra of TPB and pTP. As can be seen in this figure, TPB emits light at a longer wavelength than pTP does.

When ionizing radiation deposits energy in the wavelength shifter, many molecules may be excited. At high excitation densities, two molecules in the first singlet excited state can interact via a spin exchange interaction and undergo intersystem crossing to enter the triplet state. Since the transition from the triplet state back down to the ground singlet state violates conservation of angular momentum, these states tend to have very long lifetimes. Radiative decay of the triplet state is often referred to as phosphorescence and tends to take place on much longer timescales than the singlet state decay; however, the triplets can undergo intersystem crossing a second time and decay to the ground state back from the singlet state. These decays happen on a timescale of ~ 275 ns for TPB [126].

In order to study the scintillation of TPB and pTP under α excitation, we used an evaporator to deposit ~ 0.3 mg/cm² of TPB and pTP onto quartz slides, as shown in Figure 4.9. We then used the same apparatus shown in Figure 4.4 to deposit radon daughters onto the slides.

Following the same procedure as described above, we transferred the slides into the detector shown in Figure 4.2 and filled the detector with LAr to study the scintillation signal.





Figure 4.8: (*Left*) Emission spectrum of TPB measured at different temperatures, taken from [26]. Vertical offsets between curves have been inserted for clarity. (*Right*) Emission spectrum of pTP, from [27].




Figure 4.9: Acrylic slide with pTP wavelength shifter evaporated onto it, glowing in a UV lamp.




Figure 4.10: α -decay scintillation spectra for ^{222}Rn daughters on LAr and either TPB or pTP. The inset plot shows the f_{90} distribution for the events where the α scintillated in the wavelength shifter and the recoiling ^{210}Pb nucleus scintillated in the LAr

Analysis and Results

Figure 4.10 shows the scintillation of spectra of ^{222}Rn daughters in LAr and either TPB or pTP. Each spectrum shows two peaks in this figure. At higher energies, above 10^4 PE, we see the peak caused by the α scintillating in the LAr, while the recoiling ^{210}Pb nucleus scintillated in the wavelength shifter. At lower energies, ~ 100 PE for TPB and ~ 250 PE for pTP, we see the peaks corresponding to the case where the α scintillated in the wavelength shifter while the recoiling ^{210}Pb nucleus scintillated in the LAr.

We find that the signal produced by the α scintillating in pTP is stronger than that produced by the α scintillating in TPB. This is consistent with similar studies we performed using a ^{210}Po needle source, where we saw that TPB scintillates more efficiently and with a lower f_{90} than pTP at room temperatures, while pTP scintillates more efficiency and with a lower f_{90} at cryogenic temperatures.

Both of these signals produce events that would fall in the WIMP search region of DarkSide-50, and therefore could provide a background. However, as can be seen in the inset plot in Figure 4.10, pTP and TPB both scintillate with lower f_{90} s than we typically expect for nuclear recoils in LAr at these energies. This means that it should be possible to use f_{90} to reject these backgrounds, in addition to fiducial cuts.

While we see the f_{90} of pTP scintillation differs from that of nuclear recoils in LAr by a greater margin than we see for TPB scintillation, the shorter wavelength of pTP scintillation makes it more complicated to use for many current LAr experiments, since, for example, the acrylic used by DEAP3600 and the ITO coatings used by DarkSide-50 are not transparent to pTP scintillation light.

While it may be possible to design an experiment that can use pTP as a wavelength shifter, C. Stanford is currently studying on the TPB scintillation signal observed here can be used to reject surface backgrounds in DarkSide-50. We will therefore not discuss these backgrounds further in this document.

4.4 Neutron Backgrounds

The final class of backgrounds we will discuss here comes from neutrons. Since neutrons are not charged, they produce signals in the TPC by scattering off of target argon nuclei, which then scintillate. This is the exact same process through which we expect WIMPs to scintillate, meaning that the signals produced by a neutron can exactly mimic the signals expected from a WIMP. However, neutrons have much higher coherent nuclear scattering cross sections, and so the neutron background of an experiment, while small compared to other backgrounds, has the ability to hide a potential WIMP signal or cause an experiment to make a false claim. Understanding these neutron backgrounds is therefore crucial to the success of a dark matter experiment searching for WIMPs.

Since neutrons have a much higher elastic scattering cross section than WIMPs are expected to have, they are also more likely to scatter multiple times within the TPC. Radiogenic neutrons typically have energies in the range of 1–10 MeV. In this

energy range, the elastic scattering cross section for neutrons in argon varies from about 1–10 barns. This corresponds to a mean free path ranging from 5–50 cm [7]. In a $\sim 30\text{ cm} \times 30\text{ cm}$ TPC, like in DarkSide-50, there is a significant chance that neutrons will scatter multiple times in the TPC. As detectors are made larger, the probability of seeing multiple scatters increases. Since WIMPs are expected to have much lower cross sections, one would only ever expect them to scatter once in the TPC, which means that these multiple scattering events can be rejected.

However, since the mean free path is on the same scale as the size of DarkSide-50, one would only expect to be able to reject around half of all neutron events by multiple scattering. Relying on such cuts alone is therefore not nearly enough to eliminate these backgrounds. The solution introduced by DarkSide-50 is to nest the TPC inside a series of nested detectors, which make up the neutron veto system. We will discuss this system in much more detail in the coming chapters. However, in order to be sure that one of these backgrounds will not limit the sensitivity of the experiment, it is important to understand the source of these neutron backgrounds in great detail.

The three primary sources of neutron backgrounds are cosmogenic, fission processes, and (α, n) reactions. As mentioned in Chapter 3.1, there are also some radioisotopes that may decay and produce delayed neutrons. However, these processes are exceedingly rare and tend to produce very low energy neutrons, and so these neutrons are mostly subdominant to the other three sources.

4.4.1 Cosmogenic Neutrons

One of the main reasons for building a detector underground is to escape cosmic ray and muon backgrounds. LNGS, where DarkSide-50 is running, is underneath the Apennines mountains in Gran Sasso National Park. The overburden of rock above the laboratory creates shielding equivalent to 3800 meters of water [127], a unit commonly referred to as Meters of Water Equivalent, or m.w.e.

While this overburden does greatly decrease the muon flux, some fraction of muons still penetrate into the laboratory, where they can interact with surrounding rock and other materials to produce a shower of particles, including neutrons. Understanding these cosmogenic backgrounds is therefore crucial for understanding the neutron background.

Muon Flux

High energy cosmic rays incident on atmospheric nuclei can create hadronic showers. These showers can lead to the production of charged pions, which will then decay to muons and anti-muon neutrinos. Despite this large overburden of rock blocking most of these products, many muons still penetrate down into the underground laboratory. On the surface of the earth, the cosmogenic muon flux is $\sim 6.5 \times 10^5 \text{ m}^{-2} \text{ h}^{-1}$ with a mean energy of $\sim 1 \text{ GeV}$. In Hall C of LNGS, these muons are attenuated down to a flux of $1.228 \pm 0.004 \text{ m}^{-2} \text{ h}^{-1}$ with a mean energy of 280 GeV [128]. The mean energy at LNGS is higher than that at the surface because the higher energy muons




Figure 4.11: Cosmogenic muon energy distribution calculated from Equation 4.7 at (*blue*) LNGS, (*green*) SNOLAB, and (*purple*) SURF, normalized to the laboratory’s muon flux, summarized in Table 4.6.

are much more likely to penetrate down into the laboratory than the lower energy muons are.

In order to determine the neutron kinetic energy spectrum at some depth, it is useful to define the “slant depth” h as the integral over the muons’ path of the density of the material they are traversing in. At a depth of 3800 m.w.e., an experiment at LNGS is at a slant depth of 3.8×10^5 g/cm². The spectrum of the muons’ kinetic energy T_μ is then given in [129] as

$$\frac{dN}{dT_\mu} = \text{const} \cdot (T_\mu + \epsilon(1 - e^{-\beta h}))^{-\alpha} \quad (4.7)$$

where $\alpha \approx 3.7$ is the surface muon spectral index, $\beta \approx 3.83 \times 10^{-6}$ cm²/g describes the rate of energy loss by radiative processes, and $\epsilon \approx 620$ GeV is the critical energy above which radiative processes become dominant [129]. The constant factor at the start of Equation 4.7 normalizes the rate of muons at the given depth.

β and ϵ should both be functions of the muon energy. However, since they vary slowly with energy [129], we treat them as constant here.

Figure 4.11 shows Equation 4.7 evaluated at LNGS where DarkSide and XENON run in comparison to SNOLAB (6010 m.w.e. [130]) where DEAP3600 will run and SURF (4300 m.w.e. [83]) where LUX/LZ runs. Averaging over Equation 4.7, we find that the average energy can be given by

$$\langle T_\mu \rangle = \frac{\epsilon(1 - e^{-\beta h})}{\alpha - 2} \quad (4.8)$$

Table 4.6: Muons

Laboratory	Depth (m.w.e.)	Φ_μ ($\mu/\text{m}^2/\text{h}$)	$\langle T_\mu \rangle$ (GeV)
LNGS	3800	1.2 [128]	280
SURF	4300	0.16 [131]	290
SNOLAB	6010	0.012 [132]	330




Figure 4.12: (*Left*) The muon azimuthal angular distribution and (*right*) muon zenith angular distribution, as measured by (*blue*) Borexino and (*green*) MACRO, and (*red*) simulated by FLUKA with a limited zenith angle to match MACRO, and (*black*) with a full zenith angular distribution to match Borexino. Plot from [28].

Equation 4.8 gives us an average muon kinetic energy of 280 GeV at LNGS, 330 GeV at SNOLAB, and 290 GeV at SURF. The muon flux at these three sites is summarized in Table 4.6. Since LNGS is relatively shallow compared to these other sites, the muon flux is higher there. It is therefore especially important to veto cosmogenic backgrounds at LNGS, in order to compensate for the shallow depth.

While the muon energy distribution is well-approximated by Equation 4.7, the muon flux truly depends on the details of the rock overburden topology and is therefore more complicated. The angular distributions for these muons have been measured at LNGS by the Borexino Collaboration (in [128] and by the MACRO Collaboration (in [129]). These angular distributions were compared to distributions generated by the FLUKA simulation code in [28], and were found to agree very well.

These angular distributions are shown in Figure 4.12. As can be seen by the azimuthal angular distribution, the topology of the mountain makes the angular distribution of the muons complicated. However, we can also see from the zenith angular distribution that most of the muons have momenta pointing mostly downwards.

Muons are also sometimes correlated with each other, with a mean distance between them ~ 4 m. It has been reported by Borexino that approximately 1.5% of all muon events involve more than one muon interacting in the detector, while about 12% of all single muon interactions in the Borexino inner detector belong to multi-muon events [28]. Additionally, the OPERA Collaboration has measured the ratio of positive to negative muons to be $R_\mu = N_{\mu^+}/N_{\mu^-} = 1.377 \pm 0.014$ (stat.) $^{+0.017}_{-0.015}$ (syst.) [133].

Muon Interactions

These muons are highly relativistic, and therefore produce a large amount of Cherenkov and scintillation light in the detector. They therefore are not directly a large source of background, since they produce too large of a signal to look like a WIMP. However, through their interactions with the rocks and the surrounding laboratory materials, they can produce many other particles, including high energy neutrons which may provide a source of background.

High energy muons can interact with nuclei directly, through photo-nuclear reactions in electromagnetic showers, or through nuclear cascades in hadronic showers. These reactions are generally of the form (γ, n) , $(\gamma, 2n)$, or (γ, np) . Since these muon interactions are mostly mediated by photons, we will discuss the various reactions in terms of the energy carried by the mediating photons.

Below the first excited energy level of the nucleus (typically around 100 keV), the photon has too little energy to inelastically scatter, and we will therefore only see Thomas elastic scattering. Above this energy but below the nuclear binding energy (typically 7–8 MeV), the photon has enough energy to excite the nucleus, and so we will see inelastic scattering with resulting γ -rays as the nuclei de-excite. However, in this energy region below ~ 10 MeV, Compton scattering on atomic electrons will far dominate photo-nuclear scattering, so most of the resulting particles from these interactions will be ionized electrons.

Photons up to about 30 MeV can produce neutrons via giant dipole resonances in which the collective dipole of the nucleus is excited. While the cross section and excitation energies vary with the energy of the photon and the mass of the nucleus, cross sections typically range from 10–500 mb and produce neutrons typically in the range of 1–10 MeV. The emitted neutrons tend to form a 90° angle with the incident photon, though at higher energies, the distribution tends to shift forward [134].

Above ~ 30 MeV, photons typically have enough energy to interact with individual nucleons rather than the nucleus as a whole. Up to photon energies ~ 140 MeV, these interactions occur primarily through the quasi-deuteron effect, in which the photon can be absorbed by a proton-neutron pair (effectively a deuteron within the nucleus) and eject a neutron. This process primarily produces neutrons in the forward direction at cross sections between 50–2000 μb , and tends to be stronger for larger nuclei. Neutrons produced through this mechanism tend to have energies in the range of 20–350 MeV, with higher energy neutrons produced at smaller angles relative to the incident photon [134]. The cross section for neutron production via the quasi-deuteron effect is described by

$$\begin{aligned}\sigma(T_\gamma) &= L \frac{(A-Z)Z}{A} \sigma_d(T_\gamma) \\ L &= \frac{13.82A}{r_N^3}\end{aligned}\tag{4.9}$$

where σ_d is the cross section for a free deuteron, L is the Levinger factor and varies slowly with A , since $r_N \sim A^{1/3}$ is the nuclear radius, and A and Z are the mass and

atomic numbers of the parent nucleus, respectively. The term $(A - Z)(A)$ accounts for the number of proton-neutron pairs inside the nucleus.

At energies above 140 MeV, the energies involved in these photo-nuclear interactions is greater than the energy of the pion, meaning that these processes begin producing pions as well as neutrons and other particles. Above 300 MeV the energy carried by the photon is enough to excite various baryon resonances, starting with the Δ -resonance. These resonances can then decay into pions, η mesons, and kaons, and may also result in an intranuclear cascade, producing multiple hadrons and spallation fragments (including neutrons), and evaporation reactions. These hadrons may also proceed to produce hadronic cascades outside of the nucleus. At sufficiently high energies, the photons themselves begin to mix with the ρ , ω , and ϕ vector mesons and behave like hadrons themselves, as described by the Vector Meson Dominance model. This effect hadronic behavior allows the photon to set off hadronic showers, though causes the inner nucleons to be shielded against the photons. This effective shallowing of the nucleus causes the neutron production cross section to level off to $\sim 100 \mu\text{b}$ above $\sim 10 \text{ GeV}$, while the cross section ranges from roughly 100–600 μb at lower energies [134]. The cross section scales with the nucleus much like in Equation 4.9; however, the aforementioned screening means that A must be replaced with a screened A_{eff} given by

$$A_{eff} = \frac{A\sigma_{\gamma A}}{Z\sigma_{\gamma p} + (A - Z)\sigma_{\gamma n}} \quad (4.10)$$

where $\sigma_{\gamma p} \approx 99.8 + 57.0/\sqrt{T_\gamma}$ and $\sigma_{\gamma n} \approx 98.5 + 45.2/\sqrt{T_\gamma}$ where the energy of the exchanged photon is given in MeV [134]. Below $\sim 1 \text{ GeV}$, the effects of this screening are relatively small and can be neglected.

In addition to these hadronic products that can result from muon-nuclear interactions, high energy γ -rays can be produced by Bremsstrahlung as the muon travels through materials, and high energy electrons can be produced through ionization as well as the photoelectric effect and pair production of these high energy γ -rays.

This cascade of high energy electrons and γ -rays constitutes the electromagnetic shower that accompanies muons, while the various mesons, neutrons, and protons make up the hadronic shower.

Figure 4.13 shows the secondary particle production as predicted by a FLUKA simulation [28]. These plots show that cosmogenic neutrons are very likely to be accompanied by an array of high energy electrons and positrons, charged pions, and charged kaons in addition to an even bigger flux of high energy γ -rays, all produced by the electromagnetic and hadronic showers discussed above. Additionally, since these particles are all produced along the muon track, most will be correlated in momentum.

Prompt Neutron Production

Most cosmogenic neutrons are produced from these aforementioned photo-nuclear reactions, either directly through spallation or as a product of a resulting hadronic

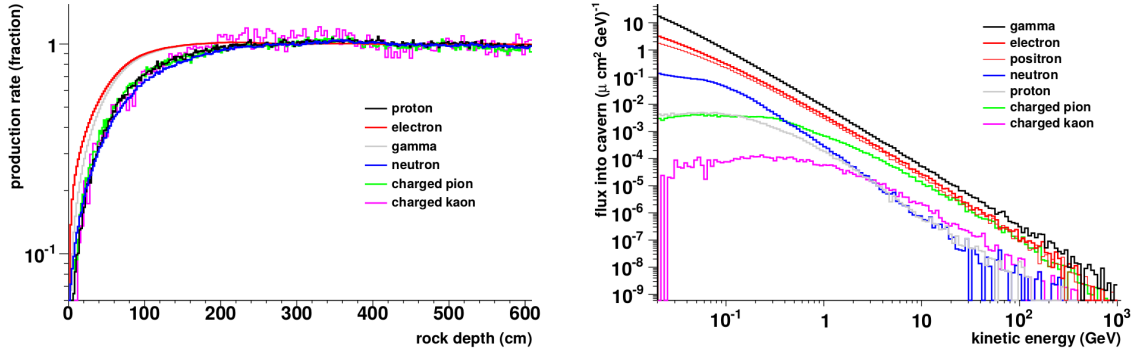


Figure 4.13: (*Left*) Particle production rates of 280 GeV muons as a function of how far the muon has traveled through rock, normalized to a maximum value of 1 for each particle species. (*Right*) The flux of particles produced by muons in Hall C of LNGS per muon as a function of the particle’s kinetic energy. Both plots were generated by a FLUKA simulation and taken directly from [28].

shower. These neutrons tend to be in prompt coincidence with the muon and its electromagnetic and hadronic showers. There is also a second class of cosmogenic neutrons that tend to be delayed from the prompt signals. Specifically, most of the above photonuclear reactions leave unstable nuclei in their wake. While many of these nuclei tend to decay with very short time constants and are therefore in prompt coincidence with the rest of the muon products, some nuclei will decay by neutron emission with timescales on the order of 10 ms or longer.

Borexino has measured the neutron yield per unit slant track length in pseudocumene given by $Y_n = (3.10 \pm 0.07 \text{ (stat.)} \pm 0.08 \text{ (syst.)}) \times 10^{-4} n/\mu(\text{g/cm}^2)$ [128]. The DarkSide-50 neutron veto, as will be discussed in Chapter 5, consists primarily of pseudocumene and trimethyl borate. The similarity between the DarkSide-50 neutron veto and the Borexino liquid scintillator means that Borexino’s measurements are likely a decent approximation of what DarkSide-50 may expect to see produced in the neutron veto. However, since DarkSide-50 has two outer detectors that would detect a muon passing through them with very high efficiency (as will be discussed in Chapter 5), neutrons produced by muons in either of these detectors are unlikely to provide a dangerous background to the experiment. Instead, we are primarily concerned with neutrons produced in the rock around the laboratory.

Figure 4.14 shows the energy spectrum for neutrons produced by muons in pseudocumene, as predicted by FLUKA simulations in [29]. Fitting a power law to this curve shows that the energy range $\sim 10\text{--}100$ MeV can be described as $\sim T_\mu^{-1/2}$, and the range from $\sim 100\text{--}1000$ MeV can be described as $\sim T_\mu^{-2}$. In the 10–100 MeV energy range, neutron production is dominated by the quasi-deuteron effect, while in the 100–1000 MeV range, the Δ -resonance effect dominates. In these energy regions, the cross section scales roughly like $\frac{(A-Z)Z}{A}$.

In the work reported by [29], the neutrons were all produced in pseudocumene; neutron production in pseudocumene will predominant occur on ^{12}C , which consti-




Figure 4.14: The cosmogenic neutron energy spectrum per meter of muon track length produced by a 285 GeV muon. From [29].

tutes $\sim 90\%$ of the mass. Since we are primarily concerned with neutrons produced in the rock surrounding the laboratory rather than those produced in the DarkSide-50 neutron veto, we can estimate the material dependence of Figure 4.14 with the aforementioned scaling laws.

The rock closest to the underground laboratories at LNGS have an average density of $2.71 \pm 0.05 \text{ g/cm}^3$ and is mostly composed of CaCO_3 and MgCO_3 [135]. This composition and density allow us to roughly scale the pseudocumene neutron production rate in Figure 4.14 by a factor of ~ 3.4 to estimate the neutron production rate in the rocks of LNGS.

FLUKA simulations of a detector like DarkSide-50 but with a fiducial mass of 3.3 tonnes of liquid argon show that 581 neutrons deposit between 0–1 MeV in the argon per year [28]. Since these high energy neutrons tend to have very long mean free paths compared to the size of these detectors, we can roughly assume that the event rate scales with the size of the fiducial volume. Doing so, we expect to see about 8.8 cosmogenic neutron events in DarkSide-50 per year. Due to the highly efficient veto system, which will be discussed in Chapter 5, we will very likely be able to veto these events with high efficiency.

Table 4.7: Cosmogenic β -delayed neutron emitters expected in DarkSide-50.

Nucleus	$Y_n (\mu\text{g}/\text{cm}^2)^{-1}$ $\times 10^{-7}$	Products	$t_{1/2}$ (ms)	I_n (%)	T_β (MeV)	T_n (MeV)
^8He	0.30 ± 0.15	$^7\text{Li}+\beta+n$	119	<0.6	0.981	6.682
		$^7\text{Li}+\beta+n$	119	<0.3	0.981	6.264
		$^7\text{Li}+\beta+n$	119	>5.4	5.501	2.946
		$^7\text{Li}+\beta+n$	119	>2.6	5.501	2.528
		$^7\text{Li}+\beta+n$	119	>5.4	7.571	1.030
		$^7\text{Li}+\beta+n$	119	>2.6	7.571	0.612
		$^3\text{H}+\alpha+\beta+n$	119	<0.9	0.981	0.08–1.99
^9Li	2.9 ± 0.3	$^8\text{Be}+\beta+n$	178	<2	11.922	0.017
		$^8\text{Be}+\beta+n$	178	26.3	11.177	0.679
		$^8\text{Be}+\beta+n$	178	13.7	10.826	0.991
		$^8\text{Be}+\beta+n$	178	<2	5.666	2.884
		$^8\text{Be}+\beta+n$	178	1.1	2.324	5.855
		$^8\text{Be}+\beta+n$	178	2.7	1.796	6.324
		$\beta+\alpha+\alpha+n$	178	3.8	11.177	~ 0.250
		$\beta+\alpha+\alpha+n$	178	2.3	10.826	0.46–1.02
^{13}B	~ 0.7	$^{12}\text{C}+\beta+n$	17.33	0.021	3.540	4.951
		$^{12}\text{C}+\beta+n$	17.33	0.001	3.540	0.512
		$^{12}\text{C}+\beta+n$	17.33	0.160	4.577	3.613
		$^{12}\text{C}+\beta+n$	17.33	0.094	5.890	2.401
^{17}N	~ 0.3	$^{16}\text{O}+\beta+n$	4173	38.5	4.126	0.387
		$^{16}\text{O}+\beta+n$	4173	0.6	3.595	0.887
		$^{16}\text{O}+\beta+n$	4173	49.5	3.303	1.161
		$^{16}\text{O}+\beta+n$	4173	6.8	2.745	1.687

Delayed Neutron Production

While, as mentioned above, most muons passing through the neutron veto or water tank will leave a sufficiently large signal that they are unlikely to provide a problematic background themselves, one danger does arrive from delayed neutrons that may be produced by unstable nuclei created by a passing muon. Some of the unstable isotopes have half-lives ranging from milliseconds to seconds, which means they can produce neutrons that are not in a fast coincidence with the passing muon. The muon yields of these isotopes in Borexino are reported in [128]. The yield for each of these isotopes as well as their half-lives and their neutron production branching ratios and neutron energies are reported in Table 4.7.

In general, β -delayed neutron energies and branching ratios were calculated by finding the energy levels and branching ratios that the parent nucleus can decay to in the intermediate nucleus and determining the energy levels that can be reached by neutron emission in the final daughter nucleus. The Q-value for each decay by neutron emission is just the mass difference of the intermediate nucleus and the neutron and final daughter nucleus plus the excitation energy of the intermediate state. We then

generalize Equation 4.6 as

$$\begin{aligned} T_E &= \frac{A_D}{A_I} Q \\ T_D &= \frac{A_E}{A_I} Q \end{aligned} \tag{4.11}$$

where T_E and T_D are the kinetic energies of the emitted particle (typically the neutron) and final daughter nucleus respectively, A_E , A_I and A_D are the mass numbers for the emitted particle, the intermediate nucleus, and daughter nucleus, respectively, and Q is the Q-value for the decay. Since these decays conserve nucleon number, we have the addition constraint that $A_E + A_D = A_I$.

Using Equation 4.11, we calculate the neutron kinetic energies shown in Table 4.7. Most of the data for this dataset came from the ENDF dataset and the data collected in [136], though wherever possible, we try to support these calculations with measurements.

^8He can decay through β -delayed neutron emission to ^7Li with a $\sim 16\%$ branching ratio by going through the $1+$ states of ^8Li at 3210 keV and 5400 keV. 9670 keV. This decay has been studied in [137], where it was found that $\sim 32\%$ of these decays go to the $1/2-$ excited state of ^7Li at 478 keV, while the rest go to the $3/2-$ ground state of ^7Li . Additionally, β -decays that go to the 9670 keV state of ^8Li may decay to ^7Li via neutron emission or to ^5He via triton emission. The total branching ratio for these two decays is 0.9%, though the fraction that decays through each mechanism is unknown. The ^5He produced by these triton decays will then decay to ^4He by neutron emission. This decay is nearly instantaneous, and so the ^5He will still be recoiling when it emits the neutron. The neutron's kinetic energy is therefore boosted when we shift from the ^5He rest frame to the laboratory frame, and we get a range of possible kinetic energies depending on the orientation of the emitted neutron.

The decay of ^9Li is much more difficult to describe; the energy levels of the intermediate ^9Be nucleus tend to be very broad, and measurements of the various energy levels and their branching ratios vary from source to source. Additionally, many of the higher energy levels accessible to β -decay in ^9Be can decay by particle emission to either the ground state or first excited state of ^8Be , with angular momentum $2+$ and energy 3.03 MeV, though measurements of the branching ratios between these two modes seem sparse. The decays with enough energy to reach this state all have angular momentum of $5/2-$ or $7/2-$, and so we will assume that they will dominantly decay to the higher spin excited state instead of the $0+$ ground state. In general, as shown in [138], the β -delayed neutrons from this decay have a broad, distribution extending up to ~ 3 MeV, with a peak slightly over 1 MeV. There is also a much sharper peak around 680 keV corresponding to neutron emission from the 2.429 MeV state, and a broader lower energy peak around 250 keV from the decay of the 2.429 MeV state into two α s and a neutron. Instead of calculating this peak directly, we estimate its value from the reported value in [138]. In general, the values given in Table 4.7 are an attempt at combining the data available as best as possible to create an effective picture of the β -delayed neutrons from ^9Li . For the most part, the values reported

Table 4.8: Rough estimates of the rate of delayed neutron precursors generated by muons within 30 cm of the TPC cryostat

	^8He	^9Li	^{13}B	^{17}N	TOTAL
precursors/year	0.03	0.33	0.08	0.03	0.47
precursors/3 years	0.10	0.99	0.24	0.10	1.43

here are in general agreement with those reported in [139], with corrections made to the neutron energy to account for the kinetic energy of the recoiling final daughter nucleus, as described in Equation 4.11.

Borexino did not report seeing ^{13}B . However, the intensity of the β -delayed neutrons from this isotope is small, so it is possible that Borexino did not see it. Nevertheless, FLUKA simulations reported in [139] estimate that the $^{13}\text{B}/^8\text{He}$ ratio should be ~ 2.4 . We therefore scale the muon yield reported in [128] for ^8He to estimate the yield for ^{13}B . The intermediate ^{13}C nucleus has an excited state of 9.897 MeV accessible to β -decay. This state can decay by neutron emission to either the ground state or first excited state of ^{12}C . The branching ratio used to calculate Table 4.7 was measured in [140].

^{17}N was also not observed by Borexino [128]. However, the FLUKA simulations in [139] report that the muon yield for ^{17}N is comparable to that of ^8He , and so the ^8He yield was used in Table 4.7. ^{17}N decays by emission of a β -delayed neutron, going through an intermediate ^{17}O nucleus. Seven β decay branches from ^{17}N to ^{17}O are reported in [141], while an eighth branch with a small branching ratio is mentioned in [139]. Of these eight branches, only four have enough energy to decay emit a neutron, and none have enough energy to decay to the first excited state of ^{16}O .

Overall, these delayed neutrons should be relatively rare compared to the other muon processes we expect to see in DarkSide-50. However, because these isotopes can be produced right outside the TPC and then decay a comparatively long time after the muon has passed, these isotopes could pose a potential source of background.

The mean free path for neutrons in this energy range in the liquid scintillator is ~ 15 cm, so delayed neutrons produced in the neutron veto far away from the TPC cryostat are unlikely to leave a signal in the TPC. Instead we are more concerned with neutrons within a roughly 30 cm distance from the cryostat. The TPC cryostat is approximately cylindrical with a height of 117.3 cm and a diameter of 64.2 cm. We therefore expect ~ 35 muons to come near the cryostat per day. Given the yields in Table 4.7, we can calculate the number of delayed neutron precursors we expect to provide potential neutron backgrounds out of a prompt coincidence with a muon signal. The results from this estimation are given in Table 4.8.

As shown in Table 4.8, the rate of these delayed neutrons is small and only a couple are expected to be produced over the course of a three year campaign of DarkSide-50. Since the energies of these neutrons are on the same scale as the expected energies of (α, n) neutrons, as will be discussed in Section 4.4.3, the same set of cuts that removes (α, n) neutrons can likely be used to efficiently remove cosmogenic delayed neutrons as well.

We will discuss all of the cosmogenic neutron backgrounds in more detail and how we plan on vetoing them in Chapter 5.1, and we will discuss the overall vetoing efficiency and expected residual background in more detail in Chapter 10.1. In summary, we expect to be able to efficiently veto cosmogenic neutron backgrounds by detecting the Cherenkov signal from muons and the products of their electromagnetic showers.

4.4.2 Fission Neutrons

The liquid drop model of the nucleus describes nuclear structure in terms of competing forces from the attractive surface potential of the residual strong force binding the nucleus together and the repulsive Coulomb force pushing the nucleons apart. When a nucleus gets sufficiently large, this long range repulsive Coulomb force can overpower the binding force and cause the nucleus to break up into two smaller fission fragments. This typically happens when $(Z^2/A) \gtrsim 45$ for a nucleus with atomic number Z and mass number A [142]. Since these nuclei can decrease their energy by undergoing fission, they will generally do so on a short timescale, and so such nuclei are rarely found in nature.

When this condition is not met but $A \gtrsim 85$, the fission process is still exothermic, meaning that fission would still be energetically favorable for the nucleus [142]. However, these reactions typically have some activation energy associated with them, coming from the binding surface potential in excess of the repulsive Coulomb potential. This activation energy provides a barrier preventing the nucleus from undergoing fission. However, quantum mechanics tells us that it is still possible for the nucleus to tunnel through this barrier and decay into two fission fragments anyway. This fission process is very suppressed and tends to happen rarely for these nuclei, which generally undergo an α -decay instead. Due to this barrier, the rate of fission for a such a nucleus is usually six or more orders of magnitude lower than its α -decay rate.

When a nucleus undergoes fission, it splits into two smaller nuclei. Conservation of baryon number in this process tells us that A must be conserved (*e.g.* the total number of nucleons among both fission fragments must equal the number of nucleons in the parent nucleus), but it does not specify how the nucleons are divided between the fragments. In fact, the two fragments that a given nucleus will produce are governed by a statistical distribution. Typically, both fragments will end up with different but relatively similar masses, with nucleon distributions that come close to filling each populated shell.

Due to the nature of this process, at least one of the fission fragments will be left unstable. This instability is necessary because the number of neutrons a nucleus needs to be stable increases super-linearly with Z . This growth pattern ensures that at least one fragment will be left with too many neutrons to be stable. The typically high energy of these fission fragments means that they will frequently emit neutrons via evaporation processes as cool down. These fragments often continue to be unstable after doing so, and will generally decay through a chain until they reach stable daughter nuclei.

The total energy released by a fission event is equal to the mass difference between the parent nucleus and the final state fission fragments. This generally results in

hundreds of MeV being carried away as kinetic energy in the released neutrons, α s, β s, neutrinos, and γ -rays produced in the decay of the resultant chains.

Neutron Energy Spectrum

The energy spectrum of neutrons produced in a fission event is generally described by the semi-empirical Watt distribution, given by [4]

$$P(T) = (\text{BR}_{SF} \langle \nu \rangle) e^{-T/a} \sinh \sqrt{bT} \quad (4.12)$$

where the term $(\text{BR}_{SF} \langle \nu \rangle)$ normalizes the rate of neutron production, with BR_{SF} is the parent nucleus’s branching ratio for undergoing spontaneous fission per α -decay and $\langle \nu \rangle$ is the average number of neutrons emitted per fission; the term $e^{-T/a}$ accounts for the Maxwellian distribution of the evaporated neutron kinetic energy in the fission fragment’s reference frame, with T being the neutron kinetic energy and a being a parameter related to the effective temperature of the neutron-emitting daughter nucleus; and the term $\sinh \sqrt{bT}$ accounts for the velocity of the parent nucleus relative to the laboratory frame, with b being a parameter related to the kinetic energy and effective temperature of the neutron-emitting daughter nucleus.

Most of the neutron-emitting daughters in the fission fragment decay chain have half-lives well below a nanosecond, and so most neutrons produced in a fission event will happen on a very short timescale. These neutrons are generally referred to as “prompt neutrons”. However, it is possible that some of the daughters may β -decay, which can sometimes happen on timescales ranging from milliseconds to seconds. If the decay product of such a daughter may emit a neutron, we can get a “delayed neutron”, separated from the primary fission event by such a timescale. These delayed neutrons are very similar to those discussed in Section 4.4.1, and need to be accounted for separately by the relative branching ratios of these delayed neutron precursors compared to the spontaneous fission branching ratio.

The parameters a and b in Equation 4.12 are typically fit to experimentally measured spectra, while BR_{SF} and $\langle \nu \rangle$ can be measured directly for a given nucleus. The SOURCES4A code [4] uses Equation 4.12 and measurements of all of these parameters, including delayed neutron branching ratios, to predict fission neutron yields for both prompt and delayed neutrons.

In DarkSide-50 the three nuclei that may undergo spontaneous fission in appreciable rates are ^{232}Th , ^{238}U and ^{235}U , with the vast majority of fissions coming from ^{238}U . Their neutron spectra, as calculated by SOURCES4A are shown in Figures 3.10, 3.12, and 3.15, respectively. SOURCES4A does not predict any delayed neutrons for any of these three nuclei.

Neutron Multiplicity

The mean number of neutrons produced by each fission, the mean neutron kinetic energy, the spontaneous fission branching ratio relative to α -decays, neutron yield, and the total number of neutrons produced in 1 and 3 years of running DarkSide-50 are shown in Table 4.9.

Table 4.9: Summary of the fission neutrons produced in the ^{232}Th , ^{238}U , and ^{235}U decay chains, as simulated by SOURCES4A [4]. $\langle\nu\rangle$ is the mean number of neutrons produced per fission and $\langle T\rangle$ is the mean neutron kinetic energy.

Chain	Isotope	$\langle\nu\rangle$	$\langle T\rangle$ (MeV)	BR_{SF}	Yield (n/s/mBq)	1 year	3 years
^{232}Th	^{232}Th	2.14	1.697	1.41×10^{-11}	3.02×10^{-14}	5.0×10^{-4}	1.5×10^{-3}
^{238}U	^{238}U	2.01	1.788	5.45×10^{-7}	1.10×10^{-9}	249	747
	^{234}U	1.81	1.989	1.20×10^{-11}	2.17×10^{-14}	4.9×10^{-3}	1.5×10^{-2}
	^{230}Th	2.14	1.808	5.33×10^{-13}	1.14×10^{-15}	2.6×10^{-4}	7.8×10^{-4}
^{235}U	^{235}U	1.86	1.990	2.01×10^{-9}	3.75×10^{-12}	0.04	0.12
	^{231}Pa	1.93	2.025	2.98×10^{-12}	5.75×10^{-15}	6.0×10^{-5}	1.8×10^{-4}




Figure 4.15: The distribution of the number of neutrons produced in ^{238}U spontaneous fission as reported by [30]. The data is well fit by a Gaussian with mean 2.05 ± 0.04 and standard deviation 1.04 ± 0.03 .

Most of the branching ratios in Table 4.9 agree with those in the ENDF databases [7]. The one exception is ^{235}U , which ENDF says has a branching ratio for spontaneous fission of 7.0×10^{-11} . Various measurements seem to predict values varying between the value in SOURCES4A and the one in the ENDF database, with the measurements in [143] agreeing with the prediction made by SOURCES4A. Since the fission neutron background is dominated by ^{238}U by multiple orders of magnitude, the disagreement in the yield of ^{235}U has little bearing on the total fission neutron background, and the use of the value predicted by SOURCES4A can be considered conservative.

Figure 4.15 shows the distribution of the number of neutrons produced per fission of ^{238}U , as reported by [30]. Notably, only $\sim 25\%$ of fissions have exactly one neutron, and $\sim 5\%$ produce no neutrons at all. It has been reported in [144] that the fission neutron multiplicity for a given nucleus is well fit by a Gaussian distribution. Fitting a Gaussian distribution to the data in [30] gives a mean multiplicity of 2.05 ± 0.04 , with a standard deviation of 1.04 ± 0.03 , consistent with the mean reported in Table 4.9.

Since we expect the neutron veto to have a very high efficiency for vetoing neutron backgrounds, we expect decays with higher multiplicity to be heavily suppressed by a quadratic factor in the inefficiency (*i.e.* if the efficiency of the neutron veto were 99%, we might expect fission events that produce two neutrons to be suppressed 1000 times more efficiently than those that produce 1 neutron). This means the primary source of fission neutrons will be those coming from ^{238}U decays in which exactly one neutron is produced. In three years of running, ~ 63 of these should be produced.

Coincident γ -rays

Neutrons produced by spontaneous fission tend to be accompanied by many high energy γ -rays, carrying away excess energy from the fission fragments' decay chains. In general, a positive correlation has been found between the number of neutrons produced by a fission event and the total amount of energy carried away by γ -rays. The total amount of energy carried away by γ -rays in these processes can be described by the empirical equation [145]

$$\begin{aligned} T_\gamma(\nu, Z, A) &= \phi(Z, A)\nu + 4.0 \\ \phi(Z, A) &= \alpha_0 - \beta_0 Z^2 A^{1/2} \end{aligned} \tag{4.13}$$

where T_γ is the total energy carried away by γ -rays, ν is the number of prompt neutrons produced, Z is the atomic number of the parent nucleus, A is the mass number of the parent nucleus, and $\alpha_0 = 2.51 \pm 0.01$ and $\beta_0 = 1.13 \times 10^{-5} \pm 7.2 \times 10^{-8}$ are fit parameters extracted from data.

Figure 4.16 shows the total amount of energy carried away by γ -rays, as predicted by Equation 4.13, for the spontaneous fission of ^{238}U . Considering this distribution in conjunction with Figure 4.15, we expect an average of 6.06 ± 0.03 MeV to be carried away by γ -rays.

The average γ -ray energy has also been modeled by [145] to obey

$$\langle T_\gamma \rangle = \alpha_1 + \beta_1 Z^{1/3} / A \tag{4.14}$$

with $\alpha_1 = -1.080 \pm 0.004$ and $\beta_1 = 106.9 \pm 0.2$ parameters fit to data. Dividing the total energy carried away by γ -rays by the average γ -ray energy gives the mean number of γ -rays produced. Figure 4.16 shows the mean number of γ -rays produced during a fission decay. For the typical ^{238}U fission event with ~ 2 neutrons, this corresponds to 6.36 ± 0.05 γ -rays, each with an average energy of 0.95 ± 0.01 .

This multiplicity of near-MeV γ -rays in prompt coincidence with the neutrons will make DarkSide-50's ability to veto fission neutrons even more efficient, as the

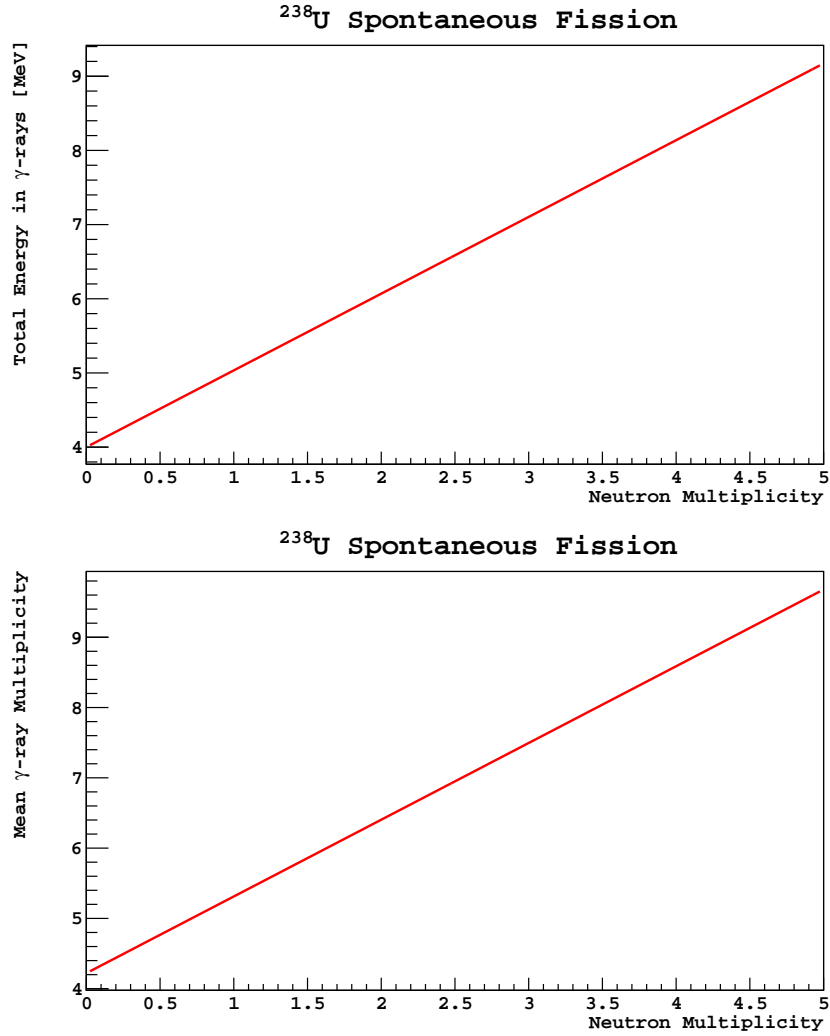


Figure 4.16: (*Top*) The total energy carried away by γ -rays as a function of neutron multiplicity for the spontaneous fission of ^{238}U , as described by Equation 4.13. (*Bottom*) The mean number of γ -rays produced in the fission fragment decay chains of ^{238}U as a function of the neutron multiplicity, as described by Equations 4.14 and 4.13.

probability of at least one of these γ -rays leaving a signal in either the TPC or in the neutron veto in prompt coincidence with the neutron is very high.

4.4.3 (α ,n) Neutrons

The third mechanism that can produce neutron backgrounds in a dark matter experiment is the (α , n) reaction. The (α , n)-yield varies between materials based on their composition in addition to their contamination. Neutron yields for this reaction are typically on the order of 10^{-6} per α -decay, and every α -emitter in the ^{232}Th , ^{238}U , and ^{235}U decay chains can produce a neutron through this reaction. While some neutrons may be accompanied by a γ -ray, many are not, and so since these neutrons only ever have multiplicity of 1, there may be no accompanying signal to make these neutrons easier to veto.

In general, calculating the (α , n) yield of a given material depends on the nuclear physics of the target nucleus and the incident α energy. The standard tool for calculating (α , n) yields is SOURCES4A [4], written by Los Alamos. However, the SOURCES4A dataset only extends to 6.5 MeV α energies, too low to fully study the relevant decay chains—especially higher energy α s are more likely to yield a neutron. Various extensions based on experimentally available data have been unofficially added to the SOURCES4A datasets, such as [146]. However, while the commonly used extensions extend the α energies for many isotopes up to 10 MeV, many important isotopes, such as ^6Li , are still entirely missing, and some isotopes retain their cutoff at 6.5 MeV. A full detailed analysis of an experiment’s (α , n) background can therefore not be carried out without accounting for the contributions of these isotopes.

Another popular tool for evaluating (α , n) yields is the nuclear reaction code TALYS [147], which integrates with the TENDL database to combine nuclear models with data from the Japanese Evaluated Nuclear Data Library (JENDL) [148] and the Evaluated Nuclear Data File B-VII.1 (ENDF/B-VII.1) [7]. The TALYS code simulates single projectile particles at a specified energy incident on a single target nucleus, and provides details of the nuclear reaction. Due to the low-level nature of the simulations performed by TALYS, extracting (α , n) yields is less straightforward than it is for SOURCES4A. However, the capabilities of TALYS are more complete, making it more apt to a full description of the (α , n) background of an experiment.

The online neutron yield calculator [149] provides a convenient interface in which one can enter a material composition and select either the ^{232}Th , ^{238}U , or Sm decay chain, and then see the TALYS-evaluated (α , n) spectrum and yield as a function of the contaminant concentration. However, this tool does not allow the user much control over the input spectrum, and so the neutron yield from ^{235}U contamination cannot be determined, and decay chains cannot be broken out of secular equilibrium. Additionally, the calculation misinterprets TALYS data and therefore produces incorrect results. The details of this calculation and where [149] goes wrong will be discussed later. The end result of the error in [149] is that the errors in interpretation cancel out by luck in many cases, so that it often gives yields that reasonably match measurements, though it occasionally is off by bigger margins, and it consistently mis-represents the outgoing neutron spectra. The neutron spectra generated by [149]

tend to have sharper features and flatter shapes than those generated by SOURCES4A or seen in data.

In order to get a complete depiction of the (α, n) background of DarkSide-50, it was then necessary to develop our own TALYS-based tool to calculate (α, n) neutron yields and spectra.

The (α, n) Reaction

The underlying mechanism behind an (α, n) reaction starts with an α , typically from the decay of a radioisotope in the relevant material, overcoming the Coulomb barrier of a target nucleus. The total potential that a nucleus has to overcome to undergo this reaction is

$$V = k \frac{2Ze^2}{r_0 A^{1/3} + r_\alpha} \quad (4.15)$$

where $r_0 A^{1/3} \approx r_N$ is the nuclear radius, k is Coulomb’s constant, the nuclear charge is Ze , $r_\alpha \approx 2.3 \text{ fm}$ is the radius of an α , and the charge of the α is $2e$. Therefore, to first order, an incident α will need kinetic energy above this threshold to overcome this threshold and react with the nucleus. If the reaction is sufficiently exothermic, it may be possible for the α to tunnel through the potential barrier, but these effects will generally be small.

As a result of Equation 4.15, nuclei with higher charges tend to have higher Coulomb barriers, and incident α s therefore need higher kinetic energies to react with the nucleus. The opposite is true for isotopes with the same Z but different A ; increasing A tends to lower the Coulomb barrier, though the effects of increasing A grow more slowly than the effects of increasing Z . Generally, this means that lighter nuclei are much more likely to undergo (α, n) reactions than are heavier nuclei.

If the incident α has enough kinetic energy (generally more than $\sim 30 \text{ MeV}$ [147]), it can scatter directly off the nucleons in the target nucleus, and “knock” a neutron out this way. This process tends to produce a relatively sharp neutron spectrum, since the neutron is coming from a specific orbital in the ground state target nucleus. Since this process involves one step, it is typically very quick. However, this energy scale is far beyond the energy scale of α s emitted in the uranium and thorium decay chains, which are all below 10 MeV .

In this lower energy range, incident α s will instead form a compound nucleus with the target nucleus, creating an intermediate nucleus with $Z_C = Z + 2$ and $A_C = A + 4$. The compound nucleus is formed by the incident α “dissolving” inside of the target nucleus, and the energy of the nucleons quickly thermalizes. This thermalization process distributes the energy over all of the compound nucleus’s nucleons, and properties of the incident α nucleus, like its spin and parity, become irrelevant.

The compound nucleus is usually in an unstable excited state, and so it will naturally decay down, either to a lower energy level by emitting a γ -ray or to a lower energy nucleus by emitting a neutron, proton, β , or other particle. Generally, the energy levels in this excited nucleus are very broad, with a continuum of energies spanning the level. As a result, higher energy α s will excite the compound nucleus to a higher energy part of the nucleus’s energy level, while lower energy α s will reach lower

energy regions of these bands. However, due to the aforementioned thermalization process, the nuclear physics of the compound nucleus will be mostly independent of the particle that caused the compound nucleus to form. For example, a $^{64*}\text{Zn}$ formed by $^{63}\text{Cu} + p$ will decay the same way as if it were formed the $^{66}\text{Ni} + \alpha$ reaction, up to any energy differences from the incident proton and α .

When the compound nucleus decays, it will decay to a daughter nucleus with some probability, governed by the identity of the compound nuclear state and the structure of the daughter nucleus. For example, the probability of the compound nucleus decaying by neutron emission depends on the energy levels, spin, and parity of the compound nucleus's excited state as well as those of the energetically accessible states of the daughter nucleus. Often there will be many particles that can be emitted, and typically a compound nucleus will emit an unbound nucleon or set of nucleons with a probability distribution that can be predicted by the properties of the accessible states.

The relationship between emitted nuclides and the identities of the compound and daughter nuclei is well exemplified by the $^{nat}\text{C} + \alpha$ reactions. In particular, the nucleus ^{16}O has 8 protons and 8 neutrons, making it a doubly magic nuclei with all of its occupied proton and neutron orbitals filled. This property makes ^{16}O extremely stable with a negative mass excess. Therefore, the reaction



will happen with a very high probability. On the other hand, ^{17}O has one neutron too many to be in this incredibly stable doubly magic state. Therefore, when an α impinges upon a ^{13}C target, we instead see the reaction



where a neutron is emitted with very high probability. If the incident α has enough kinetic energy, the compound nucleus may instead decay to the first excited state of 6 MeV, but the neutron will still be emitted. As a result, ^{13}C has a significant (α, n) yield, and tends to produce a prominent ~ 5 MeV peak in the neutron energy spectrum whenever carbon is present, despite the relatively low natural abundance of the isotope.

At projectile energies above ~ 4 MeV, it becomes possible for the compound nucleus to decay and emit nuclides before the compound nucleus has fully thermalized [147]. This is referred to as a pre-equilibrium reaction, and tends to sharpen the emission spectrum, eventually interpolating between the compound and direct reactions. However, below 15 MeV, the pre-equilibrium cross section tends to be small compared to the total compound cross section, so the effect of this process on the total (α, n) spectrum in DarkSide-50 is small.

At higher energies (typically above ~ 8 MeV [147]), it becomes energetically possible for the compound nucleus to emitting multiple particles. In these cases, multiple neutrons or a deuteron may be emitted from the nucleus. However, in the energy

ranges involved with uranium and thorium α -decays, this process should be very rare.

The TALYS simulation code [147] simulates the relevant parent, compound, and daughter nuclei, combined with data, in order to determine all of the possible reactions that may occur and their associated cross sections. In the case of (α, n) reactions, we are primarily concerned with the case where the compound nucleus decays by emitting a neutron. However, in some cases the compound nucleus may also emit a γ -ray. Additionally, if the compound nucleus decays to an excited state of the daughter nucleus, the daughter nucleus can then decay as well, and emit a γ -ray in coincidence with the neutron. These associated γ -rays are important because they may provide an additional signal that can improve the efficiency with which we can reject neutrons in DarkSide-50.

While the (α, n) yield of a given nucleus depends on the complicated nuclear physics simulated by TALYS, the energy spectrum for the emitted neutron can be explained much more simply by classical kinematics. A very clear derivation of the energy spectrum is given in the SOURCES4A manual [4], which we will attempt to summarize here.

In general, the (α, n) reaction will have some associated Q -value. However, as mentioned above, the compound nucleus might decay to an excited state of the daughter nucleus. In this case, we define the reaction Q_m -value for producing a nuclide in the m^{th} excited state as

$$Q_m = Q - E_{ex}^m \quad (4.18)$$

where E_{ex}^m is the energy of the m^{th} excitation level of the daughter nucleus. For an α with kinetic energy T_α in the laboratory frame, the velocity of the center of mass v_c in the laboratory frame is given by

$$v_c = \left(\frac{m_\alpha}{m_\alpha + m_t} \right) \sqrt{\frac{2T_\alpha}{m_\alpha}} \quad (4.19)$$

where m_t is the mass of the target nucleus and m_α is the mass of the α , and $\sqrt{\frac{2T_\alpha}{m_\alpha}}$ is the velocity of the α in the laboratory frame.

Generally, we can assume that the compound nucleus will emit its neutron isotropically in the center of mass frame. We can therefore calculate the velocity of the emitted neutron V_n to be

$$V_n = \sqrt{\frac{Q_m}{m_n} \frac{2m_r}{m_r + m_n} + \frac{2T_\alpha}{m_n} \frac{m_t}{m_t + m_\alpha} \frac{m_r}{m_r + m_n}} \quad (4.20)$$

where m_r is the mass of the recoiling daughter nucleus, and m_n is the mass of the emitted neutron.

The velocity of the emitted neutron in the laboratory frame can then be found by adding $\vec{V}_n + \vec{v}_c$, and the kinetic energy of the emitted neutron is therefore given by

$$T_n^m = \frac{m_n}{2} \times \left[\left(\sqrt{\frac{2T_\alpha}{m_\alpha}} \left(\frac{m_\alpha}{m_\alpha + m_t} \right) + \sqrt{\frac{Q_m}{m_n} \frac{2m_r}{m_r + m_n} + \frac{2T_\alpha}{m_n} \frac{m_t}{m_t + m_\alpha} \frac{m_r}{m_r + m_n}} \cos \theta \right)^2 + \left(\sqrt{\frac{Q_m}{m_n} \frac{2m_r}{m_r + m_n} + \frac{2T_\alpha}{m_n} \frac{m_t}{m_t + m_\alpha} \frac{m_r}{m_r + m_n}} \sin \theta \right)^2 \right] \quad (4.21)$$

where θ is the angle between the incident α 's velocity and the velocity of the outgoing neutron.

Yield Calculations

For a given α energy T_α and target nucleus m , TALYS calculates the cross section of an α of that energy producing a neutron of a given kinetic energy T_n . If this cross section is given by $\sigma_m(T_\alpha, T_n)$, then as the neutron travels a distance dx , the yield for neutrons of this energy is given by $Y_m(T_\alpha, T_n)$ and can be calculated by

$$Y_m(T_\alpha, T_n) = \int \eta_m \sigma_m(T_\alpha, T_n) dx$$

where η_m is the number density of the nucleus m . As the α moves through the material it is traveling in, it steadily loses energy. This loss of energy can be described by the mass stopping power $S(T) = -\frac{1}{\rho} \frac{dT}{dx}$, where ρ is the mass density of the material. Since we want to integrate the cross section over the entire α track as the α slows down in order to get the total neutron yield, a simple change of variables allows us to integrate over the entire track length to obtain

$$\begin{aligned} Y_m^\alpha(T_n) &= \frac{\eta_m}{\rho} \int_0^{T_\alpha} \frac{\sigma_m(T'_\alpha, T_n)}{S(T'_\alpha)} dT'_\alpha \\ &= \frac{N_A C_m}{A_m} \int_0^{T_\alpha} \frac{\sigma_m(T'_\alpha, T_n)}{S(T'_\alpha)} dT'_\alpha \end{aligned} \quad (4.22)$$

where N_A is Avogadro's number, C_m is the mass fraction of the nucleus m in the material, and A_m is the mass number of the nucleus. To obtain the second line in Equation 4.22, we simply note that the number density and mass density can be related by $\rho_m = \frac{A_m}{N_A} \eta_m$ where $\frac{A_m}{N_A}$ is the mass of the nucleus and that $C_m = \rho_m / \rho$.

To obtain $S(T)$, we generated a library of mass stopping powers for every element (assuming natural abundance—though it should be noted that stopping powers change fairly little for different isotopes of the same element) using the SRIM code [96]. To determine the total stopping power, we iterate over every element in the material and look up the stopping power at that α energy, linearly interpolating when neces-

sary. We then add up the stopping powers for each element, multiplying the stopping power by the mass fraction of that element in the material.

The total neutron yield from an α of given energy in the material, $Y^\alpha(T_n)$ can then be found simply by summing over the neutron yield from each individual component

$$Y^\alpha(T_n) = \sum_m Y_m^\alpha \quad (4.23)$$

The calculator lets the user specify the list of target nuclei by defining a specific material. A material is defined in a text file with three columns: the chemical symbol for a given element, the mass number of the desired isotope, and mass fraction (in percent) of that isotope. The user may instead specify “0” as the mass number to use the naturally present isotopes of the specified element in their natural abundances. Values for natural abundances used are as specified by P. De Bièvre and P.D.P. Taylor in [150]. The natural abundance of each isotope is then multiplied by the mass fraction of the specified element to determine the mass fraction of each isotope present in the material.

Finally, we are generally interested in the (α, n) neutron yield of a decay chain (or of part of a decay chain, or, more generally, a set of α -decaying nuclei. To get this, we need a list of α energies.

The calculator we developed lets the user specify a list of α energies and their relative intensities (in percent), or a list of nuclei and their relative activities (defined as the probability in percent of the nucleus appearing in a decay of the chain times the branching ratio for the nucleus to α decay). The calculator accumulates a library of Evaluated Nuclear Structure Data Files (ENSDF) [151], containing a compilation of data on the decays that a given nucleus can undergo. For a specified nucleus, the calculator checks to see if the nucleus’s data is stored in the library. If not, the appropriate ENSDF datafile is downloaded and added to the user’s library. It then retrieves the α decay energies and branching ratios from the datafile. After looking up the appropriate data for each specified nucleus, the calculator assembles them into a list of α energies T_α and associated probabilities P_α .

To determine the total (α, n) neutron yield for a list of α energies producing neutrons of energy T_n , we calculate the yield $Y(T_n)$ as

$$Y(T_n) = \sum_\alpha P_\alpha Y^\alpha(T_n) \quad (4.24)$$

Since computational restraints mean that we can only run TALYS a countable number of times, it is necessary to discretize Equation 4.22. To do so, we replace the integral over dT'_α with a sum over α energies with a constant step size $\Delta T'_\alpha$. This makes the total yield become

$$Y(T_n) = \sum_\alpha P_\alpha \sum_m \frac{N_A C_m}{A_m} \sum_{T'_\alpha \in \{T_\alpha, T_\alpha - \Delta T'_\alpha, \dots, 0\}} \frac{\sigma_m(T'_\alpha, T_n)}{S(T'_\alpha)} \Delta T'_\alpha \quad (4.25)$$

For normal calculations, we typically run the code with $\Delta T'_\alpha$ set to 10 keV, though it can be set to 100 keV for faster calculations without much loss in precision.

TALYS typically produces neutron spectra $\sigma(T_\alpha, T_n)$ with irregularly sized bins. Since we sum over many spectra it is useful to re-bin the spectra to a standard size. We have chosen to re-bin spectra into 100 keV neutron energy bins, since the bin sizes produced by TALYS are consistently 100 keV or divisors of it, making the re-binning fairly straightforward.

Total (α, n) yields can be obtained either discretely integrating the spectrum produced by Equation 4.25 over the entire spectrum. However, we can generally skip this step, because in addition to providing the user with neutron spectra, it also gives the total neutron production cross section at a given α energy $\sigma(T_\alpha)$. We can therefore replace $\sigma(T_\alpha, T_n)$ with $\sigma(T_\alpha)$ in Equation 4.25. When doing these calculations, we keep track of both yields. Typically they agree to within a couple percent of each other, with the small deviation arising from the discretization of the neutron spectrum. We therefore use the total yields calculated directly from the sum over $\sigma(T_\alpha)$ rather than the integral whenever we cite total yields.

Validation

Before we use this tool for any calculations, it is first important to verify that it returns believable results. To do this, there are two kinds of comparisons we can do: we can compare the calculator’s yield calculations to measurements taken by various groups, and we can compare yield calculations and spectra with those produced by SOURCES4A.

Comparisons between neutron yield calculations and data are made easier by publications such as the one by R. Heaton et al. [31], which have aggregated many studies into one document. As reported in [31], difficulties in neutron detection compounded upon the rarity of these reactions and sensitivity to α energy make measurements of (α, n) yields quite difficult. As a result, different experiments measuring neutron yields on the same target at the same energy often show wide discrepancies, ranging up to $\sim 40\%$.

To estimate the accuracy of this (α, n) yield calculator, we compared its calculated yields for several elements to calculations made by SOURCES4A [4] with the extensions provided in [146] and measurements made by R. Heaton et al. [31], P. Stelson and F. McGowan [32], N. Roughton et al. [152], and J. Bair and J. Gomez del Campo [33]. In some cases, Heaton’s compilation included measurements from Roughton or Bair and Gomez del Campo. To avoid double-counting these results, we only included results from these two sources when they were not used by Heaton. Yield measurements in these studies typically report 15–20% uncertainties.

A comparison between all of these benchmark measurements and calculations is made in Appendix A, where we see an overall trend of this calculator systematically overestimating the yield when it does not agree with the measurements. For comparison, we see that SOURCES4A tends to agree with the data more frequently, though underestimates the yield when it disagrees with the data. Drawing strong conclusions based on these observations is difficult due to the high uncertainty in these




Figure 4.17: A scatter plot comparing the (α, n) yields calculated by the tool presented in this present work versus benchmark calculations in [31, 32, 33]. The straight black line is $x = y$, showing where this calculator agrees with a benchmark perfectly. Solid dots represent measurements presented in Heaton, squares measurements in Stelson and McGowan, triangles pointing up measurements in Bair and Gomez del Campo, hollow circles in Roughton et al., and triangles pointing down calculations done by SOURCES4A. Different colors correspond to measurements done on different target isotopes; progressions of points the same shape and color represent yields varying over energy, which increase monotonically. A breakdown of the yield calculations for each target isotope versus energy as measured by each benchmark and this calculator is given in Appendix A.

measurements and the frequent disagreement between measurements. However, from this comparison, it looks like SOURCES4A tends to be more accurate, deviating less from the data than the calculator presented here. This is largely because SOURCES4A uses (α, n) cross sections measured directly in data, and in many of these cases, the data used by SOURCES4A is the same as the measurements used for this comparison. While the calculator in this present work tends to be more conservative, generally overestimating the neutron yield, it is also more flexible than SOURCES4A, given that this tool can estimate (α, n) yields for any α energy incident on any isotope.

It is also worth noting that, as we saw in the discussion in Chapter 1.3.3, the upper bound on the background model in a low background experiment is more crucial to the experiment’s sensitivity than the lower limit is. This distinction means that when a trade-off is necessary, it is often more prudent to err on the side of overestimating the background in the model than underestimating.

For a more holistic picture of how this calculator compares to these benchmarks, we show a scatter plot of the yield calculated by this tool and each of the benchmarks




Figure 4.18: The distribution of the deviation of neutron yields calculated in this work compared to the various benchmark measurements and calculations for a range of isotopes with Z ranging from 3 to 49 and incident α energies ranging from 5 to 11 MeV. (*Top*) A scatter plot of the deviation versus the neutron yield. (*Bottom*) A histogram of all of the deviations, with a mean of 29.6% and an RMS of 31.9%.

in Figure 4.17. As can be seen in this figure, this calculator generally comes close to the benchmark yields, though it is systemically higher, and agrees best with the data presented by Stelson and McGowan, which focuses on mid-range masses with A between 27 and 120.

While the accuracy of this calculator varies depending on the incident α energy, the target isotope, and the experiment measuring the yield, we attempt to quantify how much these calculations deviate from experiment by defining the fraction deviation in the yield $\frac{\Delta Y}{Y} = \frac{Y_{\text{this}} - Y_i}{Y_{\text{this}}}$, where Y_{this} is the yield calculated by the present work, and Y_i is yield calculated by SOURCES4A or measured by one of the aforementioned measurement sets. In this uncertainty estimation, we included yield measurements between 5–11 MeV at 1 MeV intervals, with the exception of the data from Roughton, which reported measurements at irregular intervals.

The distribution of the deviations of these calculated yields to these benchmarks is shown in Figure 4.18. As evident in this figure, this calculator predicts yields that are, on average, too high by 29.6%, with a very wide spread on this deviation of 31.9%. The distribution shown in Figure 4.18 covers the range of Z from 3 to 49, which includes most materials likely to be found in high concentration in a low-background experiment, over an energy range of 5–11 MeV, over which most (α, n) -producing α s from the uranium and thorium decay chains will be found.

Assuming the deviations in Figure 4.18 are approximately described by a Gaussian distribution, a 1σ confidence interval around the mean gives a range of fractional deviations extending 2% above the predicted yield, and 61.5% below the predicted yield. Defining an 80% confidence interval around the mean fraction deviation gives a range extending from 11% above the predicted yield to 70% below the predicted yield.

For more specific calculations, the uncertainty can be calculated by computing the total (α, n) yield from the uranium and thorium chains on each material and seeing how the total yield predictions vary in each case.

Figures A.5, A.6, A.7, A.8, A.9, A.10, A.11, and A.12 in Appendix A also show the neutron energy spectra predicted by this calculator as compared to SOURCES4A for the primary elements found in neutron-emitting components of DarkSide-50. In summary, the spectra predicted by both calculators mostly agree, with the caveat that the spectra predicted by this calculator show more structure than those predicted by SOURCES4A, which is likely an artifact of the more detailed nuclear structure models employed by TALYS.

The biggest differences between both calculators become apparent in the metals, where certain key isotopes are missing from the SOURCES4A databases, causing SOURCES4A to massively under-predict the neutron yield and miss key spectral features. At lower energy, these missing isotopes are sometimes the only source of (α, n) neutrons, causing SOURCES4A to predict a yield of 0 neutrons per α where this calculator predicts a positive yield. This is seen most clearly in chromium, where SOURCES4A is missing data on ^{52}Cr , ^{53}Cr , and ^{54}Cr , and in nickel, where SOURCES4A is missing data on ^{58}Ni and ^{61}Ni .

Additionally, SOURCES4A lacks data on ^6Li ; this is likely a small effect, since the (α, n) yield of ^6Li tends to be more than two orders of magnitude lower than that of the more abundant ^7Li . More importantly, the SOURCES4A databases do not contain data on ^7Li at 8.2 MeV above, where this calculator predicts a comparatively high yield above $1 \times 10^{-4} \text{ n}/\alpha$.

Furthermore, SOURCES4A lacks data on ^9Be above 8 MeV, which we predict to have an even higher yield than ^7Li of $2.3 \times 10^{-4} \text{ n}/\alpha$ at 8 MeV. Since this isotope is only used in the electronics connectors which have a very low total mass, this is unlikely to yield a large difference in the final rates.

DarkSide-50 (α, n) Neutron Rates

In order to determine the (α, n) neutron production rate in an experiment we calculate the neutron yield for several materials commonly used in dark matter experiments, in terms of the expected number of neutrons produced per decay of the ^{232}Th and ^{235}U decay chains, as well as the upper and lower decay chains of ^{235}U (as defined in Figure 3.11) and the ^{210}Pb decay chain, included in the lower ^{238}U chain. It should be emphasized that the neutrons produced in the ^{210}Pb decay chain are also accounted for in the lower ^{238}U chain; however, since ^{210}Pb is sometimes out of secular equilibrium with the rest of the lower ^{238}U chain, we include the (α, n) yield for the ^{210}Pb decay

chain so that we can account for the neutrons produced by the excess ^{210}Pb that we have measured.

For each material, the total (α, n) yield was calculated in three different ways: using the TALYS-based tool discussed here, using data compiled from [31, 32, 33], and using SOURCES4A.

Calculating (α, n) yields for materials made of several different elements using neutron yields calculated for the individual elements involves combining different measurements. To do this, we follow the prescription suggested by Heaton [31]. In particular, we note that the neutron yield for a single isotope i is given by Y_i where

$$Y_i = \frac{N_0}{A_i} \int_0^{T_\alpha} \frac{\sigma_i(T'_\alpha)}{S_i(T'_\alpha)} dT'_\alpha$$

Where A_i is the mass number of the isotope, σ_i is the (α, n) cross section for an alpha of kinetic energy T_α , and $S_i = \frac{1}{\rho_i} \frac{dT_\alpha}{dx}$ is the mass stopping power for an α in the medium. For a composite material made of many different isotopes, we can write the yield Y_i^c as

$$\begin{aligned} Y_i^c &= \int_0^{T_\alpha} \frac{n_i \sigma_i(T'_\alpha)}{\sum_j \rho_j S_j(T'_\alpha)} dT'_\alpha \\ &= \int_0^{T_\alpha} \frac{\rho_i S_i(T'_\alpha)}{\sum_j \rho_j S_j(T'_\alpha)} \frac{n_i \sigma_i(T'_\alpha)}{\rho_i S_i(T'_\alpha)} dT'_\alpha \\ &= \frac{\rho_i S_i(T_\alpha)}{\sum_j \rho_j S_j(T_\alpha)} \int_0^{T_\alpha} \frac{\rho_i S_i(T'_\alpha)}{\sum_j \rho_j S_j(T'_\alpha)} dT'_\alpha \\ &= f_i \frac{S_i(T_\alpha)}{S_c(T_\alpha)} Y_i(T_\alpha) \end{aligned}$$

where f_i is the mass fraction of isotope i in the compound material c and S_c is the total mass stopping power of c .

For this calculation, we assume that stopping powers can be combined additively according to Bragg's law for this compound, which we expect to be a very good assumption [31]. In general, we assume that the mass stopping power is the same for all elements; effects from nuclear size on the total stopping power should vary little on the scales we are concerned about here. We also assume that the relative contribution of each element to the total stopping power is independent of the α energy—*i.e.* that the mass stopping power for each element follows the same functional form, differing only by a constant factor. Heaton et al. [31] estimate an uncertainty of $\sim 5\%$ due to variations in this ratio for most α s between 3 and 10 MeV, though this uncertainty may be higher for lighter elements, especially carbon and beryllium. Heaton et al. also estimate a $\sim 5\%$ uncertainty from the stopping power calculations performed by SRIM.




Figure 4.19: The fractional contribution to the total stopping power for each element in (*top*) borosilicate glass and (*bottom*) stainless steel.

The fractional contributions to the total stopping power, $f_i S_i(T_\alpha)/S_c(T_\alpha)$, of each element in borosilicate glass and stainless steel are shown in Figure 4.19. As can be seen in this figure, the assumption that the relative contribution to the total stopping power is independent of the α energy tends to hold to within a few percent, justifying the use of this assumption.

Calculations of the (α, n) yields from data and using SOURCES4A are limited by the availability of measurements at each energy. Yield calculations were compiled for each material used in DarkSide-50, as listed in Tables 2.1 and 2.2, as well as for xenon, acrylic (the composition of which is described by $(C_5O_2H_8)_n$), and NUV-HD silicon photomultipliers made by Fondazione Bruno Kessler (SiPMs, composed of 99.27% ^{nat}Si , 0.5% ^{nat}O , 0.2% ^{nat}Al , and ^{nat}Ti). These yield calculations are presented in Tables 4.10 and 4.11.

For some of the yield calculations made using either SOURCES4A or data, some isotopes and or entire elements were excluded due to insufficient data. In cases where the present work found that the excluded isotope contributed to the overall neutron yield, we list the missing elements and isotopes in Table 4.12. For the most part, the isotopes missing from SOURCES4A and the compiled data tend to contribute relatively small amounts to the total neutron yield of the material, either due to the isotope being present in small amounts or due to isotope having a low (α, n) yield compared to other present isotopes. The biggest exception to this trend is beryllium, which has a very high neutron yield, yet is missing from the SOURCES4A dataset at the relevant energy ranges. As a result, the neutron yield predicted by SOURCES4A for the Be-Cu alloy used in some of the electronics connectors in DarkSide-50 is severely underestimated.

In addition to the uncertainty inherent in each (α, n) calculation, either due to inaccuracies in the data and computational tools used, some materials have additional uncertainty across all three calculations due to uncertainties in their composition.

The main contributor to the Be-Cu Alloy (α, n) yield is from the beryllium. Materion [153] cites a 50% uncertainty ($0.4 \pm 0.2\%$) in the percent mass composition of beryllium. They also cite a 22% uncertainty ($1.8 \pm 0.4\%$) in the percent mass composition of nickel. Combining these uncertainties according to their relative contributions as predicted by this work gives a 47% systematic error in the neutron yield of the Be-Cu alloy used in the PMT connectors.

The main contributor to the Borosilicate Glass (α, n) yield comes from boron. The composition used for this computation comes from personal communications with Hamamatsu, which reported the percent weight of SiO_2 , Al_2O_3 , B_2O_3 , Li_2O , Na_2O , and BaO . This report did not provide uncertainties for these concentrations, and so we will assume an uncertainty of one in the least significant digit. This report says that B_2O_3 constitutes 18%, so we will assume an uncertainty of 1%. This translates to a relative 5% uncertainty in the contribution of boron to the total neutron yield. The next leading contributions comes from aluminum and lithium, which we estimate have a 17% and 100% uncertainty in their relative contributions to the neutron yield, respectively. Together, we estimate a total uncertainty of 5.7% to the total neutron yield of borosilicate glass.

In brass, there was significant uncertainty in the type of brass used in the electronics parts. In Table 2.2, we assumed the relevant alloy is common brass, which is used for these applications. However, other plausible alloys include wire brass, red brass, yellow brass, and 70/30 brass. In general, based on the tolerances in the composition for each of these alloys and the different compositions common for these alloys, we estimate a 5% uncertainty in the percent mass of copper, which is the dominant (α , n) source. This gives a total uncertainty of 6.3% for the (α , n) yield of brass due to uncertainties in its composition.

The neutron yield of Kovar is dominated by iron and cobalt. The composition of Kovar is reported by the Carpenter Technology Corporation [154], which provides the concentrations of each component other than iron to 0.01%, and attributes the remainder of the mass to iron. We therefore assume an uncertainty of 0.01% for each listed component and a 0.02% uncertainty for iron. This implies an uncertainty of 0.01% in the (α , n) yield of Kovar as a result in the uncertainty of the composition.

The PMT electronics use Indium3.2 lead-free solder produced by the Indium Corporation[®] [155], using the SAC305 ALPHA-Lo[®] alloy from Pure Technologies [156]. The chemical composition of this alloy is reflected in the solder composition given in Table 2.2. The primary (α , n) target is copper, which is quoted as having a range of 0.04% by the manufacturers [155]. This contributes an uncertainty of 0.02% to the (α , n) yield of solder.

The biggest contributions to the (α , n) yield of 304L stainless steel come from iron and chromium. The specifications by [87] give a range of values for each element in 304L stainless steel other than iron, and attributes the remainder to iron; this includes elements such as carbon and silicon for which the specifications cite a maximum mass fraction. The concentration of chromium is listed as having a range of 2%, while the concentration of iron can vary by 11%. Overall, this leads to an uncertainty in the (α , n) yield of stainless steel of 3% due to uncertainty in the composition.

Using the masses in Tables 2.1 and 2.2, the counting results in Tables 3.2, 3.3, and 3.4, and the (α , n) yields in Tables 4.10 and 4.11, we can predict the rate of (α , n) neutron production in DarkSide-50.

The upper limits in Tables 3.2, 3.3, and 3.4 are all 90% confidence limits. To combine these with measurements that are quoted with error bars of one standard deviation, we note that the 90% confidence level corresponds to 1.645σ , and treat the limit as a measurement of $0_{-0}^{+\sigma}$. Upper and lower error bars were then added in quadrature for each measurement to determine the total neutron yield of each component.

Table 4.13 gives the (α , n) yields for each component of the TPC listed in Tables 2.1 and 2.2, as predicted by the calculator presented in this work, data presented in [31, 32, 33], and by SOURCES4A. The uncertainties in this table come from the uncertainties in the masses of each component and from the uncertainties in the counting results from the various measurements of the uranium and thorium decay chains. This table does not include systematic uncertainties in the (α , n) yields for the three ways used to calculate the yields.

In general, we see that the calculator here is the most conservative, estimating a total yield greater than that predicted by the data by a factor of 1.6. We also see

Table 4.10: (α , n) yield calculations determined by the calculator presented in this work, the presented data compilations, and calculations performed using SOURCES4A. Entries marked with a dash were not possible to calculate due to limitations in the available data.

		^{232}Th	^{235}U	$^{238}\text{U}_{\text{upper}}$	$^{238}\text{U}_{\text{lower}}$	^{210}Pb
		Yield (neutrons/decay of chain)				
Acrylic	This Work	1.35×10^{-6}	1.44×10^{-6}	2.20×10^{-7}	9.86×10^{-7}	1.19×10^{-7}
	Data	9.11×10^{-7}	9.28×10^{-7}	1.07×10^{-7}	6.32×10^{-7}	6.46×10^{-8}
	SOURCES	2.56×10^{-6}	2.81×10^{-6}	3.77×10^{-7}	1.10×10^{-6}	2.23×10^{-7}
Al_2O_3	This Work	2.21×10^{-5}	2.03×10^{-5}	5.19×10^{-7}	1.38×10^{-5}	7.28×10^{-7}
	Data	8.54×10^{-6}	8.54×10^{-6}	2.11×10^{-7}	5.96×10^{-6}	2.82×10^{-7}
	SOURCES	9.89×10^{-6}	8.59×10^{-6}	2.10×10^{-7}	2.00×10^{-6}	2.76×10^{-7}
$^{\text{nat}}\text{Al}$	This Work	4.54×10^{-5}	4.16×10^{-5}	1.01×10^{-6}	2.83×10^{-5}	1.47×10^{-6}
	Data	1.99×10^{-5}	1.19×10^{-5}	3.37×10^{-7}	1.19×10^{-5}	5.26×10^{-7}
	SOURCES	1.96×10^{-5}	1.68×10^{-5}	3.21×10^{-7}	3.78×10^{-6}	4.99×10^{-7}
$^{\text{nat}}\text{Ar}$	This Work	2.16×10^{-5}	1.65×10^{-5}	5.54×10^{-8}	1.39×10^{-5}	2.22×10^{-7}
	Data	—	—	—	—	—
	SOURCES	2.22×10^{-5}	1.48×10^{-5}	1.77×10^{-7}	2.65×10^{-6}	2.32×10^{-7}
Be-Cu Alloy	This Work	6.23×10^{-6}	6.74×10^{-6}	4.67×10^{-6}	4.67×10^{-6}	7.57×10^{-7}
	Data	5.01×10^{-6}	5.26×10^{-6}	8.29×10^{-7}	3.61×10^{-6}	4.86×10^{-7}
	SOURCES	2.48×10^{-7}	1.47×10^{-8}	1.10×10^{-6}	2.84×10^{-9}	4.89×10^{-7}
Boro-silicate Glass	This Work	2.43×10^{-5}	2.57×10^{-5}	3.94×10^{-6}	1.77×10^{-5}	2.24×10^{-6}
	Data	1.33×10^{-5}	1.41×10^{-5}	2.25×10^{-6}	9.79×10^{-6}	1.32×10^{-6}
	SOURCES	1.28×10^{-5}	1.40×10^{-5}	2.16×10^{-6}	5.99×10^{-6}	1.30×10^{-6}
Brass	This Work	3.03×10^{-7}	1.45×10^{-8}	6.51×10^{-14}	2.50×10^{-8}	5.58×10^{-13}
	Data	1.81×10^{-7}	1.19×10^{-8}	0	2.31×10^{-8}	0
	SOURCES	1.59×10^{-7}	9.44×10^{-9}	0	1.82×10^{-9}	0
Cirlex	This Work	3.12×10^{-6}	2.62×10^{-6}	2.67×10^{-7}	2.01×10^{-6}	1.43×10^{-7}
	Data	1.64×10^{-6}	1.41×10^{-6}	1.20×10^{-7}	1.04×10^{-6}	7.26×10^{-8}
	SOURCES	1.94×10^{-6}	1.77×10^{-6}	1.65×10^{-7}	5.32×10^{-7}	9.43×10^{-8}
$^{\text{nat}}\text{Cu}$	This Work	3.82×10^{-7}	1.75×10^{-8}	0	3.07×10^{-8}	0
	Data	2.13×10^{-7}	1.47×10^{-8}	0	2.88×10^{-8}	0
	SOURCES	2.53×10^{-7}	1.50×10^{-8}	0	2.90×10^{-9}	0

Table 4.11: (α, n) yield calculations determined by the calculator presented in this work, the presented data compilations, and calculations performed using SOURCES4A. Entries marked with a dash were not possible to calculate due to limitations in the available data.

		^{232}Th	^{235}U	$^{238}\text{U}_{\text{upper}}$	$^{238}\text{U}_{\text{lower}}$	^{210}Pb
		Yield (neutrons/decay of chain)				
Fused Silica	This Work	1.81×10^{-6}	1.66×10^{-6}	7.63×10^{-8}	1.15×10^{-6}	7.84×10^{-8}
	Data	1.47×10^{-6}	1.37×10^{-6}	8.27×10^{-8}	9.41×10^{-7}	7.23×10^{-8}
	SOURCES	1.73×10^{-6}	1.54×10^{-6}	1.06×10^{-7}	4.85×10^{-7}	8.55×10^{-8}
Kovar	This Work	1.21×10^{-6}	2.85×10^{-7}	3.24×10^{-9}	3.22×10^{-7}	4.01×10^{-9}
	Data	1.14×10^{-6}	2.62×10^{-7}	2.19×10^{-9}	3.31×10^{-7}	2.56×10^{-9}
	SOURCES	1.16×10^{-6}	4.10×10^{-7}	2.28×10^{-8}	1.03×10^{-7}	1.38×10^{-8}
Poly-ethyl-ene	This Work	1.50×10^{-6}	2.46×10^{-7}	1.10×10^{-6}	1.33×10^{-7}	1.61×10^{-6}
	Data	1.01×10^{-6}	1.03×10^{-6}	1.19×10^{-7}	7.01×10^{-7}	7.18×10^{-8}
	SOURCES	1.29×10^{-6}	1.35×10^{-6}	1.60×10^{-7}	4.82×10^{-7}	9.21×10^{-8}
Mylar	This Work	1.50×10^{-6}	2.46×10^{-7}	1.10×10^{-6}	1.33×10^{-7}	1.61×10^{-6}
	Data	1.01×10^{-6}	1.03×10^{-6}	1.19×10^{-7}	7.01×10^{-7}	7.18×10^{-8}
	SOURCES	1.29×10^{-6}	1.35×10^{-6}	1.60×10^{-7}	4.82×10^{-7}	9.21×10^{-8}
SiPM	This Work	3.52×10^{-6}	3.10×10^{-6}	7.85×10^{-8}	2.15×10^{-6}	1.18×10^{-7}
	Data	2.63×10^{-6}	2.31×10^{-6}	6.45×10^{-8}	1.59×10^{-6}	9.24×10^{-8}
	SOURCES	3.21×10^{-6}	2.67×10^{-6}	1.12×10^{-7}	7.46×10^{-7}	1.18×10^{-7}
Solder	This Work	2.51×10^{-9}	1.29×10^{-10}	0	2.19×10^{-10}	0
	Data	1.37×10^{-9}	9.53×10^{-11}	0	1.87×10^{-10}	0
	SOURCES	1.64×10^{-9}	9.68×10^{-11}	0	1.87×10^{-11}	0
Stain-less Steel	This Work	1.95×10^{-6}	4.49×10^{-7}	1.32×10^{-9}	5.40×10^{-7}	2.12×10^{-9}
	Data	1.25×10^{-6}	2.85×10^{-7}	1.04×10^{-9}	3.73×10^{-7}	1.43×10^{-9}
	SOURCES	1.19×10^{-6}	2.78×10^{-7}	1.64×10^{-9}	4.60×10^{-8}	1.71×10^{-9}
Teflon	This Work	1.27×10^{-4}	1.32×10^{-4}	1.17×10^{-5}	8.87×10^{-5}	9.32×10^{-6}
	Data	9.26×10^{-5}	9.56×10^{-5}	8.15×10^{-6}	6.46×10^{-5}	6.33×10^{-6}
	SOURCES	6.91×10^{-5}	7.35×10^{-5}	5.96×10^{-6}	2.54×10^{-5}	4.83×10^{-6}
$^{\text{nat}}\text{Ti}$	This Work	7.31×10^{-6}	2.62×10^{-6}	2.93×10^{-9}	2.76×10^{-6}	1.15×10^{-8}
	Data	6.30×10^{-6}	2.75×10^{-6}	0	2.60×10^{-6}	2.13×10^{-8}
	SOURCES	6.08×10^{-6}	2.47×10^{-6}	5.86×10^{-9}	3.74×10^{-7}	1.56×10^{-8}
Viton	This Work	1.16×10^{-4}	1.16×10^{-4}	1.07×10^{-5}	8.08×10^{-5}	8.49×10^{-6}
	Data	8.44×10^{-5}	8.72×10^{-5}	7.43×10^{-6}	5.89×10^{-5}	5.77×10^{-6}
	SOURCES	6.26×10^{-5}	6.66×10^{-5}	5.39×10^{-6}	2.30×10^{-5}	4.37×10^{-6}
$^{\text{nat}}\text{Xe}$	This Work	6.10×10^{-12}	3.17×10^{-14}	0	1.76×10^{-13}	0
	Data	—	—	—	—	—
	SOURCES	—	—	—	—	—

Table 4.12: Isotopes missing from the SOURCES4A and data neutron yield calculations presented in Tables 4.10 and 4.11. Where only a chemical symbol is given, all isotopes of that element are missing from the data. For borosilicate glass, the contribution of lithium to the total yield was only available at low energy in SOURCES4A, and therefore was included for ^{210}Pb and the lower ^{238}U chain.

Be-Cu Alloy	Data	^{61}Ni , ^{64}Ni
	SOURCES	Be, ^{61}Ni
Borosilicate Glass	Data	Ba
	SOURCES	Li, Ba
Brass	Data	^{67}Zn
	SOURCES	Zn
Cirlex	SOURCES	^{15}N
Kovar	Data	Cr, ^{47}Ti , ^{61}Ni , ^{64}Ni
	SOURCES	Zr, Mo, ^{52}Cr , ^{53}Cr , ^{54}Cr
Solder	Data	Sn
	SOURCES	Ag, Sn
Stainless Steel	Data	Cr, P, S, ^{61}Ni , ^{64}Ni
	SOURCES	S, ^{15}N , ^{52}Cr , ^{53}Cr , ^{54}Cr , ^{61}Ni
Titanium	Data	^{47}Ti
Xenon	Data	Xe
	SOURCES	Xe

that the predictions made by SOURCES4A and the data tend to agree more closely, though SOURCES4A is lower than the data by a factor of 1.4. The closer agreement between data and SOURCES4A than between data and this calculator is to be expected, because SOURCES4A is more data-driven and directly uses yields from many of the same measurements included in the data compilations used here.

We also group together each of the individual components into larger-scale components to provide a more holistic view of where the neutrons are coming from. The groups include

- (A) Cryostats: the lateral walls and upper and lower domes of the inner and outer cryostats (*i.e.* the bodies), all made of 5 mm 304L stainless steel
- (B) Cryostat Flanges: the 25 mm stainless steel around the flanges that join the upper cryostat domes to the bodies, the Viton o-ring inside the outer cryostat flange, and the hardware that held the flanges shut
- (C) Extraction Grid: the stainless steel grid that supplies the TPC's extraction field as well as the stainless steel support ring that holds the grid in place

Table 4.13: Total (α , n) neutron production rates per year as predicted by the calculator presented in this work, published data, and SOURCES4A calculations. Estimates are made for each component of the DarkSide-50 TPC and are arranged into 10 groups. (A) Cryostats, (B) Cryostat Flanges, (C) Extraction Grid, (D) Teflon Structures, (E) TPC Windows, (F) Liquid Argon, (G) Electric Field Cage, (H) PMTs, (I) PMTHardware, (J) Multilayer Insulation.

		Total (neutrons/1 year)					
		This work	Data	SOURCES	This work	Data	SOURCES
A	Bodies	8.2 ± 2.7	5.3 ± 1.7	4.5 ± 1.6	8 ± 3	5 ± 2	5 ± 2
B	Flanges	7.8 ± 6.2	5.3 ± 4.3	1.6 ± 0.7	805 ± 33	584 ± 24	296 ± 11
	Bolts	3.2 ± 0.3	2.0 ± 0.2	1.9 ± 0.2			
	Nuts	2.2 ± 0.3	1.4 ± 0.2	1.1 ± 0.1			
	Washers	0.54 ± 0.05	0.35 ± 0.03	0.31 ± 0.03			
	Tubes	$0.06^{+0.23}_{-0.03}$	$0.04^{+0.15}_{-0.02}$	$0.01^{+0.14}_{-0.01}$			
	O-ring	791 ± 33	574 ± 24	292 ± 11			
C	Grid(10^{-4})	1.9 ± 0.5	1.2 ± 0.3	1.2 ± 0.3	0	0	0
	Support	0.03 ± 0.01	0.02 ± 0.01	0.02 ± 0.01			
D	Reflector	$0^{+1.1}_{-0}$	$0^{+0.8}_{-0}$	$0^{+0.4}_{-0}$	0^{+5}_{-0}	0^{+3}_{-0}	0^1_{-0}
	Mounts	$0^{+4.5}_{-0}$	$0^{+3.3}_{-0}$	$0^{+1.3}_{-0}$			
E	Windows	$0^{+0.004}_{-0}$	$0^{+0.003}_{-0}$	$0^{+0.002}_{-0}$	0	0	0
F	Argon	0.12 ± 0.01	—	0.05 ± 0.01	0	—	0
G	Rings	$0^{+0.05}_{-0}$	$0^{+0.03}_{-0}$	$0^{+0.03}_{-0}$	1.2 ± 0.1	0	0
	Resistors	1.2 ± 0.1	0.50 ± 0.03	0.37 ± 0.03			
H	Stem	1490 ± 110	833 ± 64	730 ± 60	1490 ± 110	833 ± 64	730 ± 61
	Shell	$0^{+0.01}_{-0}$	$0^{+0.01}_{-0}$	$0^{+0.01}_{-0}$			
I	Screws	0^{+3E-6}_{-0}	0^{+2E-6}_{-0}	0^{+2E-6}_{-0}	35 ± 3	18 ± 1	15 ± 1
	Board	$0.08^{+0.05}_{-0.01}$	$0.04^{+0.02}_{-0.01}$	$0.022^{+0.03}_{-0.003}$			
	Resistors	$2.5^{+0.2}_{-0.1}$	1.0 ± 0.1	0.9 ± 0.1			
	Capacitors	8.1 ± 2.9	3.5 ± 1.2	1.9 ± 0.4			
	Conn.	$0^{+0.0001}_{-0}$	$0^{+0.0001}_{-0}$	$0^{+0.0001}_{-0}$			
	Solder	$0^{+0.00001}_{-0}$	$0^{+0.00001}_{-0}$	$0^{+0.00001}_{-0}$			
	C.A. Board	$0.04^{+0.02}_{-0.01}$	$0.020^{+0.01}_{-0.002}$	$0.010^{+0.01}_{-0.001}$			
	C.A. Parts	4.5 ± 0.2	1.8 ± 0.1	1.5 ± 0.1			
	RF Conn.	12 ± 1	7.3 ± 0.6	7.1 ± 0.6			
	MCX Conn.	7.5 ± 0.9	4.1 ± 0.5	4.0 ± 0.5			
J	Mylar	3.2 ± 0.9	2.0 ± 0.6	1.6 ± 0.4	5 ± 1	3 ± 1	2 ± 1
	Al Layer	1.3 ± 0.8	0.6 ± 0.4	$0.2^{+0.2}_{-0.1}$			
	Netting	0.04 ± 0.01	0.03 ± 0.01	0.02 ± 0.01			
Total					2340 ± 120	1440 ± 70	1050 ± 60

- (D) Teflon Structures: the Teflon reflector that surrounds the TPC as well as the support structures around the TPC and the mounts that hold the PMTs in place
- (E) TPC Windows: the fused silica diving bell and cathode that separate the PMTs from the volume of active argon
- (F) Liquid Argon: both the active and inactive liquid argon volumes
- (G) Electric Field Cage: the copper field cage rings that supply the drift field in the TPC as well as the alumina resistors that hold the rings at different voltages
- (H) PMTs: the physical PMT structures, including the borosilicate glass stem and the Kovar shell
- (I) PMTHardware: the copper support screws that hold the PMTs to the Teflon mount, the electronic components of the PMTs (including the Cirlex circuit board and alumina resistors and capacitors), the electronic components of the cold amplifiers (including the Cirlex board and the alumina circuit elements (labeled as “C.A. Parts” in Table 4.13), and the various copper and Be-Cu alloy connectors
- (J) Multilayer Insulation: the Mylar sheets, aluminum coating, and polyethylene netting that separates the inner and outer cryostats

In the last group of columns in Table 4.13, we round the neutron production rates to the nearest integer, and larger-scale components that are not within their uncertainty of producing 1 neutron per year are rounded down to 0.

As can be seen in Table 4.13, the two dominant sources of background are in the cryostat flanges, mostly due to the Viton o-ring, and the borosilicate glass stems in the PMTs. These relatively high rates are due to the high (α, n) yields of fluorine in the Viton o-ring and boron in the borosilicate glass PMT stems and the high relatively high levels of radioactive contaminants in each of these components. We also see that the PMT hardware, cryostats, and multilayer insulation also produce neutrons at a sub-dominant but not necessarily negligible rate. The other components can mostly be neglected for the rest of this analysis.

The energy spectra of the neutrons from these sources are shown in Figures 4.20–4.24.

Notably, a future detector could drastically reduce by around 2000 neutrons/year the (α, n) neutron background by avoiding Viton o-rings (for example, by using titanium o-rings for the inner and outer cryostat) and either by using PMTs with ceramic stems—which may be cleaner and have a lower neutron yield—or by using silicon photomultipliers, which can avoid borosilicate glass altogether.

Now that we have characterized the neutron backgrounds typical in a dark matter experiment and what will be present in DarkSide-50, the next step is to establish how these backgrounds will be eliminated.




Figure 4.20: (α, n) neutron energy spectrum from the cryostat body.




Figure 4.21: (α, n) neutron energy spectrum from the cryostat flanges.




Figure 4.22: (α, n) neutron energy spectrum from the multilayer insulation.




Figure 4.23: (α, n) neutron energy spectrum from the TPC PMTs.




Figure 4.24: (α , n) neutron energy spectrum from the TPC PMT hardware.

Chapter 5

Neutron Veto Concept

The basic idea behind the neutron veto is simple: since neutrons can provide signals in the LAr that look identical to WIMPs, we put the LAr TPC in two nested neutron detectors. If we see a coincidence between the LAr TPC and either of the vetoes, we discard the event as being a neutron.

The goal of this anti-coincidence system is to allow radiogenic and cosmogenic neutrons to be vetoed. Of the two nested outer detectors, the outermost one is the 11 m diameter by 10 m high cylindrical Water Cherenkov Veto, or WCV. This detector provides passive shielding to the two detectors inside of it, and provides a veto for eliminating cosmogenic neutrons based on the Cherenkov light produced by the muons that produce them or by the products of the electromagnetic shower also produced by the muon.

Inside the WCV is the Liquid Scintillator Veto, or LSV. This detector contains a boron-loaded organic liquid scintillator that can detect coincidences between neutrons that scatter in the LAr TPC and itself. Both the prompt coincidence signal caused by the neutron thermalizing in the LSV, as well as the delayed coincidence signal caused by the neutron capture reaction in the scintillator, allow an anti-coincidence cut to be made in DarkSide-50 to efficiently remove neutron backgrounds.

While much of the foundational work developing the idea behind such a vetoing system was originally explored by Wright et al. [157], this chapter will discuss the underlying physics behind how the veto system works.

5.1 Water Cherenkov Veto

The WCV serves two primary purposes. The first is to provide passive shielding to the TPC and LSV. The rocks surrounding the laboratory and the rest of the environment surrounding the experiment have some natural radioactivity. The second function is to veto signal produced by muons that may be in coincidence with a cosmogenic neutron by detecting the associated Cherenkov light produced by these particles.




Figure 5.1: (*Left*) Total interaction cross section for neutrons on (*blue*) ^1H and (*green*) ^{16}O . Data from [7]. (*Right*) The neutron mean free path in water.

5.1.1 Passive Shielding

The WCV used in DarkSide-50 contains ~ 100 ktonnes of water, which blocks most of the external γ -ray and neutrons produced by the surrounding environment from reaching the LSV.

Figure 5.1 shows the mean free path of neutrons in water. Since there is a minimum of 3 meters of water between the inner wall of the WCV and the outer wall of the LSV in DarkSide-50, virtually all external radiogenic neutrons will be blocked by the WCV before they even reach the LSV.

According to [158], the absorption length of 2 MeV γ -rays in water is ~ 38 cm. Since the highest energy naturally present radiogenic γ -rays come from ^{208}Tl and have an energy of 2.2 MeV, we expect the WCV to suppress the rate of external radiogenic γ -rays from reaching the LSV by over four orders of magnitude.

This shielding from external γ -rays is crucial to the successful operation of the LSV. In order to achieve high radiogenic neutron vetoing efficiency, the LSV must be sensitive to relatively low energy signals. Too high of a background rate in the LSV may hide these lower energy signals and prevent the LSV from obtaining its target efficiency. The passive shielding provided by the WCV is therefore instrumental to the performance of the LSV.

5.1.2 The Cherenkov Effect

When a charged particle moves through a medium, it will electromagnetically perturb the molecules that it passes, creating a time-varying perturbation in the electromagnetic field. Typically, since the electromagnetic force is conservative, there will be no net change in the field felt by a molecule long before and long after the charged particle has passed it. However, if the medium the particle is traveling in has an




Figure 5.2: Illustration of the Cherenkov effect from a charged particle moving through some medium with a velocity greater than the speed of light in that medium.

index of refraction n , the speed of light in that medium will be $c_n = c/n$, where c is the speed of light in a vacuum.

If the charged particle has velocity $v > c_n$, it travels faster than light in the medium and therefore out-paces the perturbations it creates in the electromagnetic field. Perturbations with a wavelength λ can then add constructively, as illustrated in Figure 5.2 and radiate energy in a cone of light opening at some characteristic angle.

This process is known as the Cherenkov effect, and tends to take place on very short timescales when charged particles pass through a medium. The number of photons N emitted by a particle with a kinetic energy T over a track of length x with a stopping power $S(T)$ per unit energy lost at a wavelength λ is given by

$$\begin{aligned} \frac{d^2 N}{dx d\lambda} &= \frac{2\pi\alpha}{\lambda^2} \left(1 - \frac{1}{\beta^2 n(\lambda)^2} \right) \\ \frac{d^2 N}{dT d\lambda} &= \frac{1}{S(T)} \frac{2\pi\alpha}{\lambda^2} \left(1 - \frac{1}{\beta^2 n(\lambda)^2} \right) \end{aligned} \tag{5.1}$$

where $\beta = v/c$ is the speed of the particle and α is the fine structure constant.

By integrating Equation 5.1 over the range of wavelengths the detector is sensitive to (roughly 200–500 nm) over a particle’s track, we can predict the number of photons produced by a charged particle traveling through the WCV. As shown in Figure 4.13, charged particles produced by a muon typically have kinetic energies on the \sim GeV scale.

The index of refraction of water is ~ 1.33 in the wavelength range we are interested in, and it varies slowly over the range of wavelengths that the WCV PMTs are sensitive to. Given the energy scale of the charged particles and the size of the WCV, we can therefore expect within an order of magnitude of 10,000 photons to be produced in the WCV by the various charged cosmogenic particles.

Therefore, as long as these photons can be detected, a large water tank outfitted with PMTs, such as the DarkSide-50 WCV, should be able to veto cosmogenic events very efficiently.

5.2 Liquid Scintillator Veto

The LSV completely surrounds the cryostat containing the LAr TPC. This 4π coverage allows the LSV to detect γ -rays and neutrons that scatter in the TPC and in the LSV and produce coincident signals, making it possible to veto and make *in situ* measurements of backgrounds in DarkSide-50.

The DarkSide-50 LSV is a 4 m diameter stainless steel sphere outfitted with 110 Hamamatsu R5912 PMTs. It is filled with a scintillator cocktail consisting of pseudocumene as the primary scintillator, trimethyl borate for boron-loading, and the wavelength shifter PPO.

5.2.1 The Scintillator

The primary scintillator in the DarkSide-50 LSV is 1,2,4-Trimethylbenzene, also known as pseudocumene or PC, and is described by the chemical formula $C_6H_3(CH_3)_3$. Figure 5.3 shows a diagram of the PC molecule.

This is the same primary scintillator that was used in Borexino. Since the PC used by Borexino was derived from a petroleum base, the ^{14}C background in the PC is very low, as discussed in Chapter 3.1.6 and measured by [109]. Using PC from the same source, DarkSide-50 was able to obtain highly radiopure PC with a good scintillation efficiency.

Scintillation Mechanism

Since PC is a benzenoid structure, it scintillates through the same mechanism illustrated in Figure 4.7 as was discussed in the context of TPB and pTP. When energy is deposited in the scintillator, it produces some number of excitons. These excitons may eventually be captured by a scintillator molecule and can decay radiatively, producing a scintillation photon, or non-radiatively, releasing the exciton’s energy as heat.




Figure 5.3: Molecular diagrams of (*top left*) PC, (*top right*) TMB, and (*bottom*) PPO.

The number of excitons produced when some amount of energy is deposited into a scintillator depends on the excitation energy of the scintillator as well as the amount of energy deposited. Since there is a fixed probability of an exciton being dissipated radiatively by the molecule, we expect the number of photons produced in a scintillation event to be proportional to the initial energy deposited.

However, when the scintillator is exposed to a given form of ionizing radiation, some of the energy will end up ionizing the scintillator molecules, or otherwise excite molecules into configurations where they are unlikely to decay radiatively. These “damaged” species may absorb excitons and dissipate their energy non-radiatively. While these species tend to return to normal on the scale of ~ 1 ns (ionized species will recombine with the molecules they were freed from, and excited molecular states will return to their ground state configurations, for example), they can locally suppress the amount of scintillation light produced in an event.

Due to the short recovery timescale, this effect tends to suppress fast scintillation light more than short scintillation light; as a result, the fast component of the scintillation (from the singlet state decay) tends to get suppressed more heavily than




Figure 5.4: The (*black*) attenuation length, (*blue*) absorption spectrum, and (*red*) emission spectrum of PC, as measured by [34].

the long component from the triplet state decay. This behavior is in contrast to the effects of impurity quenching. Since impurities introduced to the scintillator do not necessarily recover over time, they effect both the fast and slow components. Since scintillator molecules in the triplet state have more time to interact with these impurities, the slow component tends to be more suppressed by impurities than the fast component is.

The effect of suppressing scintillation light through this mechanism is referred to as ionization quenching, and it is described by Birks in [159]. The model presented by Birks assumes that the number of “damaged” molecules produced is proportional to the stopping power $\frac{dE}{dx}$ of the incident radiation and is equal to $B\frac{dE}{dx}$ where B is a proportionality constant. The model then assigns a probability k of a damaged molecule dissipating the exciton relative to the probability of an undamaged molecule doing so.

Birks then describes the number of fluorescent photons F produced per unit track length x of the incident radiation as

$$\frac{dF}{dx} = \frac{Y \frac{dE}{dx}}{1 + kB \frac{dE}{dx}} \quad (5.2)$$

where Y is the number of excitons produced by the radiation per unit stopping power. Since k and B appear as a single product in this equation and cannot be disentangled, they are generally counted as a single parameter called the “Birks’ constant”, kB .

Figure 5.4 shows the emission and absorption spectrum of photons in pure PC, along with their attenuation length. Since the LSV is a large volume of liquid scintillator and photons typically have to travel several meters before they are detected,

the probability of the scintillator absorbing its own light is particularly important. Since the emission and absorption spectra significantly overlap, PC may absorb much of its own scintillation light. This attenuation can drastically reduce the amount of light that is ultimately collected and reduce the sensitivity of the LSV.

5.2.2 Wavelength Shifter

In order to decrease the amount of self-absorption, we add a wavelength shifter to the scintillator. For the DarkSide-50 LSV, we use 2,5-Diphenyloxazole, or PPO, as shown in Figure 5.3. PPO is described by the chemical formula $C_{15}H_{11}NO$.

PPO is a very efficient scintillator with a large Stokes shift and a short scintillation time. Like PC, it is primarily composed of benzenoid rings. This similar structure means that energy deposited in PC can be thermally transferred non-radiatively to PPO molecules very efficiently. As a result, only a low concentration of PPO is generally needed for the scintillation energy to be transferred to PPO. Since PPO is a more efficient scintillator than PC and has a shorter decay time, adding PPO to the scintillator can increase the scintillation efficiency and decrease the response time. Additionally, since the PPO can radiatively absorb excitons instead of “damaged” molecules, PPO may also be able to decrease kB for the scintillator.

PPO’s large Stokes shift means that it shifts PC scintillation light to much longer wavelengths, where the PMTs are more sensitive and far from the absorption peak of PC. In general, the probability of light at a given wavelength λ being absorbed by a material such as PPO is given by the decadic molar extinction coefficient ϵ . If the PPO has a molar concentration c , then the decadic molar extinction coefficient can be related to the attenuation length L by the Beer-Lambert law, given by

$$L(\lambda) = \frac{\log_{10} e}{c\epsilon(\lambda)} \quad (5.3)$$

where $e \approx 2.72$ is Euler’s number; the factor of $\log_{10} e$ converts the decadic molar extinction coefficient to the Napierian (natural) base of the attenuation length.

Equation 5.3 was used to calculate the absorption length in Figure 5.4. Since PC and PPO may both absorb photons at a given wavelength, the attenuation length of the mixture can be estimated by adding the attenuation lengths from PC and from PPO in reciprocal. However, while photons absorbed by PC are typically lost, PC re-emits photons at a longer wavelength with an efficiency of $\sim 82\%$ [160].

Figure 5.5 shows the emission spectrum of PC, as well as the absorbance— $c\epsilon$ in Equation 5.3—of three different scintillator cocktails containing various mixtures of PC, PPO, and TMB. Since energy deposited in PC is transferred non-radiatively to PPO on very short timescales, we expect the primary scintillation light of a PC-PPO mixture to resemble the emission spectrum of PPO more closely than that of PC.

The absorbances shown in Figure 5.5 were measured by A. Ianni and Y. Suvorov using an optical test bench in Perugia. The measurements made for PC+1.5 g/L of PPO and for PC+2.5 g/L of PPO were made by measuring the transmittance of






Figure 5.5: (*Top*) The emission spectrum of PPO, as measured by [35]. (*Bottom*) The absorbance of (*black, dashed*) PC+ 1.5 g/L of PPO, (*red, solid*) PC+ 2.5 g/L of PPO, and (*blue, dotted*) an equal mixture of PC and TMB with 2.5 g/L of PPO, measured by A. Ianni and Y. Suvorov.

light through a 10 cm cell filled with the scintillator cocktail being studied, while the measurements of PC+TMB+2.5 g/L of PPO were made using a 1 cm cell. As a result of the smaller size, the measurement of the absorbance of this latter cocktail becomes statistically limited at longer wavelengths where the absorbance is very low. The absorbance measurements of the PC+TMB+2.5 g/L of PPO cocktail at wavelengths greater than ~ 380 nm are therefore dominated by noise; the true absorbance of this cocktail at these wavelengths is likely much lower than is shown in this figure. Measurements taken by A. Ianni indicate the TMB is largely transparent to PPO scintillation light, indicating that the absorbance of this cocktail is likely much lower at these wavelengths.

5.2.3 Boron-loading Agent

While the scintillator cocktail described thus far can detect scintillation light from particles scattering in it, sufficiently low energy neutrons may produce too little light to be detected. To avoid losing these signals, we load the scintillator with boron by adding trimethyl borate, or TMB, shown in Figure 5.3 and described by the chemical formula $\text{BO}_3(\text{CH}_3)_3$.

^{10}B has a natural abundance of 19.9% [150] and has a thermal neutron capture cross section of 3838 b [157]. When a neutron captures on ^{10}B it can decay through two channels:



In the former case, the compound ^{11}B nucleus decays directly to the ground state of ^7Li . While the α and ^7Li nucleus each have over 1 MeV of kinetic energy, their relatively large masses and consequently high stopping powers means that their scintillation light will be heavily suppressed by ionization quenching, as described in Equation 5.2. In the latter case a 478 keV γ -ray is also produced. This high energy γ -ray will produce electron recoils, which will be quenched very little. However, if the γ -ray goes back into the cryostat, it may be lost entirely.

Nevertheless, in both cases described in Equation 5.4, the α and ^7Li nucleus will only travel $\sim 1 \mu\text{m}$ before stopping, meaning that they will never escape the LSV. As long as these heavily quenched reaction products can be detected with a very high efficiency, these neutron captures can reliably be detected and may produce a veto signal independent of the energy of the neutron when it first enters the LSV. Understanding the optics and the scintillator response to backgrounds is therefore essential for designing a highly efficient neutron veto.

In addition to capturing on ^{10}B , fast neutrons above ~ 7.2 MeV can capture on ^{11}B . These capture reactions are summarized in Figure 5.6. However, these captures tend to be sub-dominant to the captures on ^{10}B .




Figure 5.6: Fast neutron capture reaction cross sections on ^{11}B for (*blue*) the $^{11}\text{B}(\text{n},\alpha)^8\text{Li}$ reaction, (*green*) the $^{11}\text{B}(\text{n},\text{t})^9\text{Be}$ reaction, and (*purple*) the $^{11}\text{B}(\text{n},\text{p})^{11}\text{Be}$ reaction.

5.2.4 Prompt Coincidence Signals

Neutrons and γ -rays may produce prompt coincidences between the LSV and the TPC by scattering once in the LAr and a second time in the LSV. Since neither of these particles are charged, they will not directly scintillate in the LSV. Instead, γ -rays will primarily scatter off of electrons, and neutrons will scatter off of nuclei. These recoiling charged particles will then scintillate as described in Equation 5.2.

Since lower energy particles tend to have higher stopping powers and neutrons and γ -rays distribute their energy among multiple particles, we would expect, for example, a γ -ray of a given energy to be quenched more heavily than an electron of the same initial energy. This effect is more pronounced for neutrons, since nuclear recoils are generally more quenched than electron recoils.

The maximum energy, T_{max} that a neutron can transfer to a nucleus of mass number A is kinematically limited to

$$T_{\text{max}} = \frac{4A}{(1+A)^2} \quad (5.5)$$

For fast neutrons scattering on a given nucleus, the energy distribution is typically uniform ranging from 0 to T_{max} . As a result, neutrons will most efficiently transfer their energy to protons as they thermalize, and they will produce a broad signal in the LSV equal to the sum of the quenched energy from each collision.

Measurements made by [39] found that proton recoils in an organic liquid scintillator are quenched by a factor ~ 10 . Recoils off of heavier nuclei will be quenched

even more. We can therefore expect the signal made by the neutron as it thermalizes in the scintillator to be suppressed by such a factor. Since this signal is inherently dependent on the energy of the neutron when it reaches the LSV, lower energy neutrons may produce too little light to be detected as they slow down. This is especially true for neutrons that thermalize in the cryostat and are already thermal when they reach the scintillator. In order to detect these neutrons, we must rely on the signal the neutron produces when it captures, which we will discuss later.

Inelastic Scatters

In addition to the signals that a neutron makes as it slows down in the LSV, the neutrons may also inelastically scatter in the various detector components and produce a γ -ray that may be seen in the LSV. While many inelastic scatters are possible on the different nuclei that appear in the TPC and cryostat materials, two important lines appear around 1.43 MeV and 1.46 MeV for ^{52}Cr and ^{40}Ar , respectively, with cross sections ranging from approximately 0.5–1.5 b for neutrons with kinetic energy between 1 and 10 MeV. Chromium makes up $\sim 18\%$ of the mass of the stainless steel of the cryostat, and ^{52}Cr has a natural abundance of 83.8% [150], while ^{40}Ar comprises nearly all of the mass of the LAr, making these signal important for understanding the prompt neutron spectrum.

Additionally, ^{19}F , which is found in the Teflon reflector and support structures of the TPC, has two excited states: one state at 110 keV, and a metastable state at 197 keV. The ground state of ^{19}F has angular momentum and parity of $J_\pi = 1/2+$. Since the first excited has angular momentum and parity $J_\pi = 1/2-$, transitions to this state in inelastic neutron scatters are suppressed due to the change in parity. However, the second excited state has $J_\pi = 5/2+$, making it more accessible. This state decays to the ground state with a branching ratio close to 100%, producing a 197 keV γ -ray with a lifetime of 87.3 ns. The relatively long lifetime of this state means that it may produce a γ -ray with a ~ 100 ns delay from the neutron scatter in the TPC.

A summary of the different inelastic scatters and their associated γ -rays is given in Table 5.1. While this is not a complete list of all states that may be excited by inelastic neutron scatters on the various components of the TPC and cryostat, it covers many of the primary γ -ray lines we may expect to see in prompt coincidence with neutron scatters in the TPC, which may compliment the neutron thermalization signal.

5.2.5 Delayed Coincidence Signals

The primary delayed coincidence neutron signal comes from the neutron capture reaction on ^{10}B , as shown in Equation 5.4. In addition to capturing on ^{10}B , a thermal

Table 5.1: Dominant inelastic scatter signals that may be seen in the LSV in prompt coincidence with neutron scatters in the TPC. The excited nucleus, its ground state angular momentum and parity J_π^0 , their excited state energy levels E' and their angular momenta and parity J'_π , the excited states' lifetime $T_{1/2}$, the cross section for a 3 MeV neutron incident on each nucleus exciting the nucleus to each given state σ , the final state energy levels that each excited state can decay to E^f , the γ -ray energy released by this transition E_γ , and the branching ratio for the excited state to decay to each final state BR. Energy level structure information from [3, 5, 6], and cross sections from [7].

Nucleus	J_π^0	E' [keV]	J'_π	$T_{1/2}$	σ [b]	E^f [keV]	E_γ	BR [%]
^{40}Ar	0+	1461	2+	1.12 ps	0.6	0	1461	100
		2121	0+	90 ps	0.08	1461	660	100
		2524	2+	0.22 ps	0.9	1461	1063	57
						0	2524	43
^{52}Cr	0+	1434	2+	0.78 ps	0.74	0	1434	100
		2370	4+	6.7 ps	0.12	1434	936	100
		2646	0+	?	0.06	1434	1213	100
		2768	4+	1.9 ps	0.03	2369	398	2
						1434	1334	98
^{19}F	1/2+	110	1/2-	0.6 ns	0.09	0	110	100
		197	5/2+	89.3 ns	0.33	110	87	<0.06
						0	197	~100
		1346	5/2-	2.86 ps	0.11	197	1149	3
						110	1236	97
		1459	3/2-	62 fs	0.12	1346	113	<0.2
						197	1262	11
						110	1349	69
						0	1459	20
		1554	3/2+	3.5 fs	0.17	1459	95	<0.14
						1346	208	<0.01
						197	1357	93
						110	1444	5
						0	1554	2

neutron may also capture on ^1H , ^{12}C , ^{16}O , ^{17}O , or ^{18}O through the following reactions

$$\begin{aligned}
&^1\text{H} + n \rightarrow ^2\text{H} + \gamma \text{ (2223 keV)} \quad I_\gamma/I_\gamma(\text{max}) = 100\% \quad \sigma = 0.33 \text{ b} \\
&^{12}\text{C} + n \rightarrow \begin{cases} ^{13}\text{C} + \gamma \text{ (3090 keV)} & I_\gamma/I_\gamma(\text{max}) = 100\% \\ ^{13}\text{C} + \gamma \text{ (4945 keV)} & I_\gamma/I_\gamma(\text{max}) = 67\% \\ ^{13}\text{C} + \gamma \text{ (1860 keV)} & I_\gamma/I_\gamma(\text{max}) = 57\% \end{cases} \quad \sigma = 0.0034 \text{ b} \\
&^{16}\text{O} + n \rightarrow ^{17}\text{O} + \gamma \text{ (4145 keV)} \quad I_\gamma/I_\gamma(\text{max}) = 100\% \quad \sigma = 0.0002 \text{ b} \\
&^{17}\text{O} + n \rightarrow \begin{cases} ^{18}\text{O} + \gamma \text{ (6066 keV)} & I_\gamma/I_\gamma(\text{max}) = 100\% \\ ^{18}\text{O} + \gamma \text{ (8047 keV)} & I_\gamma/I_\gamma(\text{max}) = 23\% \\ ^{18}\text{O} + \gamma \text{ (4413 keV)} & I_\gamma/I_\gamma(\text{max}) = 14\% \end{cases} \quad \sigma = 0.0039 \text{ b} \\
&^{18}\text{O} + n \rightarrow \begin{cases} ^{19}\text{O} + \gamma \text{ (3859 keV)} & I_\gamma/I_\gamma(\text{max}) = 100\% \\ ^{19}\text{O} + \gamma \text{ (3955 keV)} & I_\gamma/I_\gamma(\text{max}) = 85\% \\ ^{19}\text{O} + \gamma \text{ (2484 keV)} & I_\gamma/I_\gamma(\text{max}) = 72\% \end{cases} \quad \sigma = 0.0002 \text{ b}
\end{aligned} \tag{5.6}$$

where σ is the thermal neutron capture cross section, $I_\gamma/I_\gamma(\text{max})$ is the intensity of the γ -ray, relative to the maximum intensity γ -ray [7]. For ^{12}C only the three dominant γ -rays are shown (notably, ^{12}C will often produce multiple γ -rays after capturing a neutron).

The capture time for neutrons in the LSV and the fraction that catch on each of these isotopes depends on the relative concentrations of PC and TMB; for a scintillator cocktail with equal parts PC and TMB, we would expect a capture time $\sim 2.2 \mu\text{s}$ with $\sim 0.8\%$ of neutrons capturing on ^1H rather than ^{10}B , while a cocktail with 5% TMB and 95% PC would have a capture time of $\sim 22 \mu\text{s}$, with $\sim 7.7\%$ of neutrons capturing on ^1H rather than ^{10}B .

While the γ -rays produced by the neutron captures on ^1H and ^{12}C will be quenched relatively little, their mean free path in the scintillator $\sim 20 \text{ cm}$ means that some of these γ -rays may go back into the cryostat without producing any visible signal in the LSV. This risk is notably not an issue for neutrons that capture on ^{10}B , where the α and recoiling ^7Li nucleus will deposit all of their energy into the scintillator.

Other Neutron Captures

There is also some probability that neutrons will not make it into the LSV after they scatter in the TPC. In these cases, they may thermalize and capture on the TPC and cryostat materials. In the cryostat, the two most likely isotopes to capture a neutron are ^{56}Fe and ^{52}Cr , described by the reactions

$$\begin{aligned}
&^{56}\text{Fe} + n \rightarrow ^{57}\text{Fe} + \gamma \text{ (7.6 MeV)} \quad \sigma = 2.6 \text{ b} \\
&^{52}\text{Cr} + n \rightarrow ^{53}\text{Cr} + \gamma \text{ (7.9 MeV)} \quad \sigma = 0.9 \text{ b}
\end{aligned} \tag{5.7}$$

In the TPC, neutrons may capture on ^{19}F and ^{12}C in the Teflon or on ^{28}Si in the fused silica. The neutron capture reaction on ^{12}C is described in Equation 5.6, while

the captures on ^{19}F and ^{28}Si are described by

$$\begin{aligned}
^{28}\text{Si} + n \rightarrow & \begin{cases} ^{29}\text{Si} + \gamma \text{ (4934 keV)} & I_\gamma/I_\gamma(\text{max}) = 62\% \\ ^{29}\text{Si} + \gamma \text{ (3539 keV)} & I_\gamma/I_\gamma(\text{max}) = 69\% \\ ^{29}\text{Si} + \gamma \text{ (2093 keV)} & I_\gamma/I_\gamma(\text{max}) = 20\% \\ ^{29}\text{Si} + \gamma \text{ (6380 keV)} & I_\gamma/I_\gamma(\text{max}) = 11\% \end{cases} \quad \sigma = 0.17\text{b} \\
^{19}\text{F} + n \rightarrow & ^{20}\text{F} + \gamma \text{ (6.6 MeV)} \quad I_\gamma/I_\gamma(\text{max}) = 100\% \quad \sigma = 0.01\text{b}
\end{aligned} \tag{5.8}$$

Cross sections in Equations 5.6, 5.7, and 5.8 were reported by [7]. In Equation 5.8, only the four γ -rays with the highest intensities are listed for the $n(^{28}\text{Si}, ^{29}\text{Si})\gamma$ reaction [161].

All of these materials have relatively low neutron capture cross sections compared to the LSV. Additionally, since the cryostat and TPC are primarily made of mid-to-high mass isotopes, it is relatively unlikely that neutrons will fully thermalize and capture before reaching the LSV. However, in cases where this does happen, [157] shows that many neutron captures in the cryostat and TPC materials can be detected by their resulting γ -rays with a time window of $60\text{ }\mu\text{s}$.

5.2.6 Muon and Cosmogenic Neutron Signals

While we expect muons and their resulting electromagnetic showers to produce easily visible signals in the WCV, as discussed in Section 5.1, their signal can be bolstered by the scintillation and Cherenkov light that they produce in the LSV.

Since PC has a refractive index ~ 1.5 and TMB has a refractive index of ~ 1.3 , we can expect a cocktail containing both of them to have a refractive index between these two values. This refractive index is higher than that of water, meaning that high energy charged particles passing through the LSV will produce even more Cherenkov light than was produced in the WCV.

Furthermore, photonuclear reactions and collisions between muons and coincident pions and electrons with the liquid scintillator will produce large amounts of scintillation light. In general, due to the high energy of these events, we expect passing muons to produce far more light in the LSV than we expect to see from any other particle, making these signals very easy to identify.

While muons are easy to detect when they pass through the LSV, the smaller size of the LSV means that muons are more likely to enter the WCV than the LSV. Nevertheless, much of the muon vetoing power of the LSV comes from its ability to detect cosmogenic neutrons as they pass through the WCV.

While fast neutrons coming from the rocks of the laboratory are highly unlikely to ever penetrate the WCV due to the relatively high elastic scattering cross section of fast neutrons on hydrogen, neutrons with energy on the GeV scale tend to have elastic scattering cross sections around $0.1\text{--}1\text{ b}$. This cross section range is consistent with mean free paths between 30 cm and 3 m ; higher energy neutrons will have even longer mean free paths.

These long mean free paths means that many cosmogenic neutrons may penetrate the shielding provided by the WCV and LSV and eventually scatter in the LAr. Since these neutrons are generally accompanied by the muons and their associated electromagnetic shower, we expect to veto them efficiently with these particles' signals in the WCV and LSV.

However, if a neutron is produced by a muon in the rock at an oblique angle to the muon and its shower, it is possible that the WCV and LSV will not see such a clear signal. Furthermore, such a signal may also be missing if the electromagnetic shower producing the neutron happens deep within the rocks around the laboratory; if charged particles produced by the shower are stopped by the rocks while the neutron escapes, it may not be accompanied by a clear signal in the WCV. These two cases are likely the most dangerous classes of cosmogenic neutron background.

In these case, the LSV can detect the neutrons if they scatter in the scintillator before or after they enter the TPC. As a result, as long as the neutron scatters at least once in the LSV while it is either going towards or away from the TPC, the LSV may be able to veto these neutrons, even if they do not capture. This extra layer of vetoing power will help the outer detectors to eliminate these cosmogenic backgrounds.

Chapter 6

Designing and Prototyping

In Chapter 5 we proposed an outer detector design concept that could, in theory, efficiently veto cosmogenic and radiogenic neutrons, allowing a dark matter experiment such as DarkSide-50 to run background free.

However, as discussed in this section, designing such a veto system requires being sensitive to low energy signals and understanding the optics of the detector well enough to be confident that it is unlikely to miss a neutron signal that might otherwise produce a false WIMP event in the TPC.

In order to ensure that the outer detectors' design can achieve the requisite sensitivity, we performed a series of tests on prototype detectors and then ran optical simulations scaled up to the full DarkSide-50 geometry. In this chapter, we will discuss the prototype studies that were done to understand the optical properties of the outer detectors and guide the design.

6.1 Material Choices

One of the primary reasons for doing these prototype tests was to help choose reflectors for the WCV and LSV as well as a scintillator cocktail for the LSV that would give us the sensitivity we need.

Much of these early tests were done in close collaboration with E. Shields and are discussed in [13], and many aspects of the design drew heavily upon experience gained and passed on through collaborators working on Borexino [162].

6.1.1 WCV

The WCV design is based upon the Borexino CTF detector described in [162].

It is filled with $\sim 1,000$ tonnes of water, purified by Borexino water purification plant to achieve very low contamination levels.

The WCV is equipped with 80 ETL 9351 8" PMTs, lining the inside of the cylindrical vessel and the outside of the stainless steel sphere containing the LSV. These PMTs are equipped with an optical fiber used for laser calibration, as will be discussed in Chapter 8.2, as well as an acrylic canonical light collector.

Since the 80 8" PMTs provide fairly little phototube coverage, we lined the inner surfaces of the WCV and the outside of the LSV stainless steel sphere and support structures with a Tyvek-based reflector. The reflector consists of two layers of DuPont Tyvek 1082D bonded together by a layer of polyethylene. Tyvek is a void-based reflector which relies on the different indices of refraction between a matrix of microscopic voids and the material itself to reflect light with wavelengths between 300–800 nm diffusely with an efficiency $\sim 96\%$ in air. Since Tyvek is mildly hydrophobic, water cannot fill the voids when it is submerged in water; instead the higher refractive index of water means that Tyvek has an efficiency $\sim 99\%$ when submerged in water [163].

The signal produced by muons in the WCV is expected to be far above the background, meaning that optimizing the light collection efficiency of the WCV is not as crucial as it is for the LSV. Instead, we were able to save money by using Tyvek reflector supplied by the Daya Bay collaboration [163], and PMTs and light collectors supplied by the CTF experiment [162].

6.1.2 LSV Scintillator

The success of the LSV at vetoing radiogenic neutrons depends on its ability to detect the highly quenched signals that may be produced by neutrons thermalizing in the detector and the $\alpha+{}^7\text{Li}$ produced by 6.4% of all neutron captures on ${}^{10}\text{B}$, as discussed in Equation 5.4. Since these signals may produce relatively little light, it is crucial that the LSV have a high light yield.

As discussed in Chapter 5.2.1, the primary solvent in the scintillator cocktail will be PC. The experience that the DarkSide collaboration has inherited from Borexino has given us a PC purification and distillation plant we know to be reliable and shown that a PC-based liquid scintillator can be used to make a high-light yield detector. Additionally, the same petroleum-derived, low- ${}^{14}\text{C}$ PC used by Borexino can be used in DarkSide-50 to achieve a low ${}^{14}\text{C}$ activity.

TMB will also be added to the scintillator as a source of boron. One important feature of TMB, as we learned from our studies, is that it mixes well with PC. When TMB is first added to PC, their different densities cause them to separate; since they have different indices of refraction, a visible line forms at their interface showing where the TMB layer ends and the PC layer begins. However, once the two fluids have been thoroughly mixed together, no such boundary appears; we have found that PC and TMB remain evenly mixed without visibly separating again for over three years.

Additionally, studies done by A. Ianni have shown that scintillator cocktails containing equal concentrations of PC and TMB and 10 g/L of PMP (a wavelength shifter similar to PPO) has $\sim 85\%$ of the light yield of a cocktail containing only PC and PMP. Similar effects have been reported by Birks [160] showing that a cocktail consisting of equal parts PC and xylene has a light yield $\sim 80\%$ of that of a pure xylene cocktail. This effect is likely due to the ability of TMB to non-radiatively transfer energy deposited into it to PC or wavelength shifter molecules. Furthermore, A. Ianni has found that TMB is largely optically transparent to the scintillation light of PC and PPO, meaning that adding TMB to a PC+PPOscintillator may increase




Figure 6.1: The relative light yield of a PPO+linear alkylbenzene solution, for various PPO concentrations taken inside a 5 cm diameter by 5 cm tall cell, from [36].

the attenuation length of light. These properties mean that a high concentration of TMB can be used without severely compromising the light yield of the detector.

We considered the use of PPO and PMP as a secondary wavelength shifter at various concentrations. PMP (1-phenyl-e-mesityl-2-pyrazoline) is a wavelength shifter similar to PPO with a larger Stokes shift and a higher scintillation efficiency. While we did perform tests with PMP and found that it can perform well, we ultimately ended up deciding to use PPO due to constraints on the availability of PMP.

The ideal concentration of wavelength shifter is ultimately a balance of the efficiency with which energy will be transferred to it, which increases with concentration, and the amount of light lost due to PPO re-absorbing and non-radiatively dissipating its own scintillation light, which also increases with concentration. Figure 6.1 shows the change relative light yield of a small vial containing the linear alkylbenzene scintillator with various concentrations of PPO, measured by [36]. At low PPO concentrations, the light yield increases approximately linearly, while at higher concentrations it appears to level off. The maximum light yield is achieved at PPO concentrations near 0.3% mass fraction, corresponding to a concentration of ~ 3 g/L.

Secondary Wavelength Shifter

We also considered the use of two secondary wavelength shifters: POPOP (1,4-bis(5-phenyloxazol-2-yl) benzene, described by the chemical formula $C_{24}H_{16}N_2O_2$) and bis-MSB (1,4-Bis(2-methylstyryl)benzol, described by the chemical formula $C_{24}H_{22}$).

Secondary wavelength shifters work very similarly to primary wavelength shifters, except that they primarily work by absorbing scintillation photons and radiatively re-emitting them. The use of a secondary wavelength shifter can potentially improve the performance of a detector by shifting scintillation light to longer wavelengths, further from the absorption peaks of the primary solvent and wavelength shifter, where it will have a longer attenuation length. It can also shift the photons to wavelengths




Figure 6.2: The (*solid*) absorption and (*dashed*) emission spectra for (*black*) POPOP dissolved in cyclohexane [35], and (*red*) bis-MSB dissolved in PC [34].




Figure 6.3: The relative light yield of a linear alkylbenzene solution with 3 g/L PPO in a 5 cm diameter by 5 cm tall cell with varying concentrations of either (*left*) POPOP or (*right*) bis-MSB dissolved in it, from [36].

where the PMTs may be more sensitive, or the reflector may be more reflective. The absorption and emission spectra for POPOP and bis-MSB are shown in Figure 6.2.

Figure 6.3 shows the relative light yield of a linear alkylbenzene solution containing 3 g/L and varying concentrations of POPOP and bis-MSB. Here, we find that the light yield stops increasing at a relative mass fraction of $\sim 0.0025\%$ for POPOP and $\sim 0.001\%$ for bis-MSB, corresponding to concentrations of 0.25 g/L and 0.15 g/L respectively. These measurements were made using a 5 cm diameter by 5 cm tall Teflon cell with a UV-transparent window on one face, and do not reflect potential




Figure 6.4: The degradation of the attenuation length of PC over time when exposed to a surface-to-volume ratio of 20 m^{-1} of electropolished stainless steel compared to PC left unexposed. In order to speed up the interactions, the sample in stainless steel was heated by 50°C for the first 300 hours and then by 100°C afterwards, from [13].

changes in the light yield of a larger detector due to the attenuation length of light in the scintillator.

Impurity Effects

Introducing impurities to a liquid scintillator can decrease its efficiency by absorbing scintillation light, decreasing the scintillator's attenuation length, or by providing modes by which excitation energy can be dissipated non-radiatively, decreasing the yield of the scintillator. This latter effect is referred to as impurity quenching.

Oxygen Contamination Oxygen is a particularly strong quenching agent, and even a small contamination from oxygen can have a dramatic effect on its scintillation yield [164]. Because oxygen is paramagnetic, it can realign its magnetic field and dissipate excitation energy non-radiatively. Furthermore, large amounts of oxygen can turn PC yellow and make it less transparent.

Even if the scintillator is stored in an air-tight container, oxygen can still enter the system by outgassing from the container materials. In order to remove oxygen that may have been dissolved in the PC, one can distill the PC in a vacuum, or bubble it with nitrogen to remove the oxygen gas.

Containers made of metals that contain metal oxides may provide another source of oxygen. Some information regarding the long-term chemical compatibility of PC with various metals can be found in [165].

Effects of PC interacting with stainless steel have been observed to cause the attenuation length of light in PC to decrease over time [13]. Figure 6.4 shows the degradation of PC when exposed to a surface-to-volume ratio of 20 m^{-1} of electropol-

ished stainless steel for 24 days. After 400 hours, the attenuation length of 310 nm light drastically decreases.

Large vessels with lower surface-area-to-volume ratios are expected to show less degradation over time. For experiments that last a long time, this effect may still be significant and require either coating the stainless steel to prevent the reaction or re-purifying the scintillation mixture multiple times during operation.

Water Contamination TMB is highly hygroscopic and will form boric acid when brought into contact with moisture through the reaction,



This reaction is exothermic, so if it happens at a high rate, it may pose a safety risk. Even in small amounts, however, this reaction precipitates out boric acid, which is a fine white powder that may serve as a contaminant in the scintillator. The presence of boric acid may therefore decrease the attenuation length or efficiency of the scintillator.

It is therefore important that TMB be distilled and stored in a moisture-free environment, and that all surfaces coming into contact with TMB be as dry as possible.

6.1.3 LSV Reflector

Because 100% phototube coverage would be prohibitively expensive, the light collection efficiency of the LSV can be improved by lining the inner surfaces of the stainless steel sphere containing the scintillator with a reflector. That way, light that does not immediately get detected by a phototube can be reflected by the walls until it is detected.

Several reflector candidates were considered, including types of Tyvek (a paper-like void-based reflector made of polyethylene chains), Duraflect (a void-based reflector made of compacted Teflon powder), Crystal Wrap (a thin Teflon-based film with a high void fraction), high purity aluminum foil (a metallic surface reflector), 3M Foil (a highly reflective specular surface reflector), and Lumirror (a polyethylene terephthalate-based—PET-based—multi-layer film used in LCD screens).

Most of these void-based reflectors were discarded from consideration because they lose reflectivity when submerged in the scintillator—liquid scintillator can fill the voids, decreasing the difference in the refractive index between the voids and the reflector material, causing the reflector to become translucent. While many of these reflectors are hydrophobic so water cannot fill the voids, this property does not protect the reflectors from the scintillator. 3M foil, despite its high reflectivity, was discarded because it chemically reacts with TMB, and aluminum foil was discarded due to its comparatively low reflectivity.

The two remaining reflector candidates were Tyvek 4077D and Lumirror 188 E6SR. Tyvek 4077D has a titanium oxide component that keeps the scintillator from filling the voids.




Figure 6.5: Reflectance measurements of Tyvek reflectors. The black curve shows the reflectance of Tyvek 4077D with one layer, while blue shows the same with 6 layers. Additional curves compare two different types of Tyvek after soaking in a PC+TMB cocktail for various lengths of time: (*red*) 4077D after 3 hours, (*orange*) 4077D after 3 weeks, (*green*) 1070D before soaking, and (*cyan*) 1070D after 3 weeks. (*Not shown*) Tyvek 1082D exhibited very similar behavior to 1070D.

Reflectance measurements were made using a Perkins Elmer Lambda-650 spectrophotometer, which measures the diffuse reflectance of a material between 200–860 nm at an 8° angle from the reflector surface. Reflectance measurements were normalized to measurements made of a Spectralon plug, which is nearly 100% reflective in this wavelength range [166]. Multiple measurements of the same sample with this spectrophotometer showed variations in the reflectance $\sim 0.1\%$.

Reflectance measurements of Tyvek 4077D and 1070D, respectively with and without the titanium-oxide component, are shown in Figure 6.5, and measurements of Lumirror 188 E6SR are shown in Figure 6.6. These figures show the reflectance of different numbers of layers of these reflectors after soaking in a PC+TMB solution for various lengths of time. Key features of these measurements are summarized in Table 6.1.

Tyvek 1082D, used in the WCV, exhibits a very similar reflectance to Tyvek 1070D before soaking in scintillator, as shown in Figure 6.5. While we see that the reflectance greatly diminishes after soaking in scintillator, we have found that it does not do so in water, and so this reflector candidate is viable in the WCV, but not in the LSV.

We also see that Tyvek 4077D maintains a high reflectance when submerged in scintillator for a long time; however the protective titanium oxide component introduces a low-wavelength cutoff below ~ 400 nm.

Figure 6.6 shows that a single layer of Lumirror can have a high reflectance, reaching a maximum of $\sim 98\%$ when dry. However, it also exhibits a low-wavelength




Figure 6.6: Reflectance measurements of Lumirror 188 E6SR, after soaking in a PC+TMB cocktail for varying lengths of time: (*black*) before soaking, (*blue*) after 3 days, (*orange*) after 3 weeks, (*green*) after 4 months, and (*cyan*) after 10 months. (*Red*) The improvement of a dry sample while using 4 layers. The bump around 240 nm is due to the specular component of the reflectivity.




Figure 6.7: A small sample of Lumirror 188 E6SR that has been soaked in scintillator for nine months. Some degradation of the reflectance at the edges has occurred due to the creeping of scintillator through the sides. The reflectance of the central part of this sample remains unaffected, however (see Figure 6.6).

Table 6.1: Summary of the Tyvek and Lumirror reflectances shown in Figs. 6.5 and 6.6. R_{250} and R_{500} are the reflectances measured at 250 nm and 500 nm, respectively. $\lambda_{50\%}$ and $\lambda_{90\%}$ are the wavelengths at which each reflector had 50% or 90% reflectance, respectively. Notably, the six layers of Tyvek 1070D was above 50% for the entire range of wavelengths measured while dry, and the second time it passes the 90% reflectance mark is reported. While it is dry, this sample is below 50% reflectance for the entire range of wavelengths measured.

Description	R_{250} [%]	R_{500} [%]	$\lambda_{50\%}$ [nm]	$\lambda_{90\%}$ [nm]
Tyvek 4077D				
1 layer, dry	25.5	96.0	392	411
6 layers, dry	23.1	99.2	392	411
6 layers, soaked 3 hours	10.4	98.7	401	416
6 layers, soaked 3 weeks	8.1	93.6	403	422
Tyvek 1070D				
6 layers, dry	85.2	93.2	—	335
6 layers, soaked 3 weeks	9.5	41.5	—	—
Lumirror 188 E6SR				
1 layer, dry	18.0	97.8	315	344
4 layers, dry	17.3	99.4	315	345
1 layer, soaked 3 days	13.2	97.1	316	349
1 layer, soaked 3 weeks	10.1	97.8	316	348
1 layer, soaked 4 months	10.0	97.6	317	349
1 layer, soaked 10 months	13.2	97.1	318	352

cutoff below ~ 325 nm. Because Lumirror has two protective layers surrounding the reflective void-based region inside, it is minimally affected by being submerged in scintillator, with the exception of a slow, creeping degradation that grows along the sides. This degradation is shown in Figure 6.7, where we can see the Lumirror becoming somewhat translucent. This effect takes place at a rate of approximately 1 cm over the course of 9 months. Measurements of the degraded edges show that their peak reflectance drops to $\sim 83\%$ at 380 nm; however, even when the edges lose reflectance, the bulk remains unaffected. A large detector running for several years can avoid a decreasing light yield due to this effect by overlapping edges of the Lumirror by a several centimeters so that the translucent edges will be in front of reflective bulk.

It has also been shown that Lumirror fluoresces with primary absorption between 320–420 nm and emits around 440 nm [167]. This fluorescence may act as an additional wavelength shifter and may increase our light collection. However, the efficiency of this fluorescence is unreported.

As will be discussed in Section 6.2, we ultimately decided to use Lumirror 188 E6SR in the DarkSide-50 LSV. In order to be sure that the natural radioactivity present in Lumirror would not produce too high of a background rate in the LSV, we sent a 1.7 kg sample to M. Laubenstein to measure using the Germanium Counting

Table 6.2: Radionuclide measurements made of a 1.7 kg sample of Lumirror E6SR (188 μm thick). Measurements were made by M. Laubenstein using the Germanium Counting technique. Nuclides that are part of the same decay chain are grouped together.

Decay Chain	Nuclide	Concentration [mBq/kg]
^{232}Th	^{228}Ra	72 ± 6
	^{228}Th	31 ± 3
^{238}U	^{226}Ra	657 ± 26
	^{234}Th	< 330
	^{234m}Pa	150
^{235}U	^{235}U	9 ± 4
	^{40}K	430 ± 50
	^{137}Cs	8 ± 1
	^{60}Co	< 1.8

technique discussed in Chapter 3.2. The results of these measurements are shown in Table 6.2.

The DarkSide-50 LSV, with a 4 m inner diameter, has ~ 7 kg of Lumirror inside of it. Given the counting results in Table 6.2, the radioactivity of the Lumirror should be low enough to make it a viable choice.

6.2 Full Prototype Detector

In order to test the viability of the LSV design and to choose the best scintillator cocktail and reflector combination, we re-purposed a bell jar from an evaporator to construct a prototype detector.

6.2.1 Apparatus

The prototype detector was made from a roughly cylindrical bell jar with a 30.5 cm diameter and a height of 30.5 cm, with a rounded section on the top. The bell jar could be opened by removing the lid, which clamped shut over a Viton o-ring.

A Conflat port on the side allowed the insertion of a Hamamatsu R11065 PMT, similar to those used by the DarkSide-50 TPC. The PMT was contained inside of a glass jar to separate it from inside of the bell jar, and optical coupling grease was used between the face of the PMT and the inside of the glass jar.

A series of small screws were welded to the inside of the bell jar so that the various reflector candidates could be attached to the inner surface of the bell jar, and a cylindrical sheath of reflector was placed around the sides of the PMT housing. A diagram of the bell jar is shown in Figure 6.8.

Figure 6.9 shows the reflectance of Lumirror as compared to the absorption and emission spectra of the various primary and secondary wavelength shifters discussed




Figure 6.8: Cross section of bell-jar setup for each experiment. (*Left*) The cell setup shows a cell filled with scintillator in the otherwise dry bell jar. (*Right*) The filling setup shows the bell jar, which has been modified to be filled with scintillator.




Figure 6.9: Plot of a typical quantum efficiency distribution for an R11065 PMT, provided by Hamamatsu [14], shown alongside the reflectance spectrum for Lumirror and the emission spectra of PPO [35], POPOP [35], and bis-MSB [37], all scaled to arbitrary units.

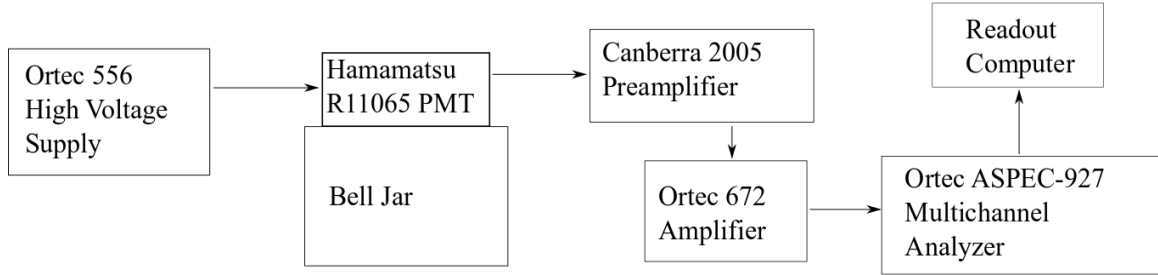


Figure 6.10: The data acquisition electronics setup for the bell jar measurements.

in Section 6.1 as well as the quantum efficiency of the R11065 PMT used for these studies.

6.2.2 Modes of Operation

As shown in Figure 6.8, the bell jar was operated in two modes of operation. First, it was tested dry with an isolated scintillator cell placed on the floor inside. This setup was used as a preliminary test to determine the optimal scintillator and reflector combination to be used in the LSV. The cells were fused silica cylinders, 7.6 cm in diameter and 7.6 cm tall. The cells were filled with various scintillator cocktails through a stem that was approximately 0.5 cm in diameter and 6 cm tall after being sealed shut once the cell was filled. In this configuration, the reflector lined the curved face of the bell jar, as well as all of the other internal surfaces.

The second series of tests were performed with the bell jar filled with scintillator. A mesh grid separated the cylindrical portion of the bell jar from the rounded end, as shown in Figure 6.8, and Lumirror covered the mesh grid instead of the rounded face. Fittings were added to the lid of the bell jar through which the bell jar could be filled and drained, and an additional feedthrough lead to an expansion tank which allows for fluctuations in the volume of the scintillator as a result of thermal expansion.

A diagram of the electronics and data acquisition system used for these measurements is shown in Figure 6.10. The PMT was powered by an Ortec 556 high voltage power supply operated at positive high voltage. The signal from the PMT went through a Canberra 2005 preamplifier to an Ortec 672 amplifier operated at $0.5\mu\text{s}$ shaping time. SPE peaks were measured with an amplifier gain of 1.5×1000 ; γ -ray spectra were measured at a gain 100 times smaller. A sample SPE distribution taken with this setup can be seen in Figure 2.10. The unipolar triangular output of the amplifier was passed to an Ortec ASPEC-927 multichannel analyzer, which then transmitted the data to a computer where it was recorded using Ortec's Maestro multichannel analyzer software.




Figure 6.11: A picture of the distillation column used to purify PC and TMB for prototype tests.

6.2.3 Small Cell Tests

Preparation

Before filling the cells with the scintillator cocktails, it was first necessary to purify the PC and TMB. Figure 6.11 shows the distillation column used for this purification. The pre-distilled fluid (either PC or TMB) was first filled into a three-neck flask, which sat on a heating mantle and magnetic stirrer. A stirring rod and boiling stones were added to the pre-distilled fluid to create nucleation points to help the fluid boil. A nitrogen tank was connected to one of the necks of the pre-distilled fluid flask to run a constant stream of nitrogen through the whole system.

A variac (not shown in Figure 6.11) was used to control the heat of the heating mantle, and a temperature sensor (also not shown) was placed between the flask and the mantle to monitor the temperature. The boiling point of PC is $\sim 170^{\circ}\text{C}$ and the boiling point of TMB is $\sim 69^{\circ}\text{C}$. The heating mantle was set to $\sim 190^{\circ}$ for PC and $\sim 90^{\circ}$ for TMB.

The vapors from the fluid then travel up the Vigreux column, which was insulated with additional foam tubing not shown in Figure 6.11. At the top of the column, a thermometer measured the temperature of the vapors, and vapors that made it up the column could then enter the condenser, which was cooled with a constant flow of water.

Vapors would then condense in the condenser, while those that do not (which may contain impurities with a lower boiling temperature than the fluid being distilled) will




Figure 6.12: (*Left*) A photograph of the setup used to sparge and seal the scintillator cells. (*Right*) A diagram of this setup, labeling each of the parts.

enter the vapor-release port. These vapors flow into the vapor-trap which is made into a cold trap with the addition of an ice bath (not shown).

Fluid that has re-condensed at the bottom of the condenser is then trained into a second three-neck flask which holds the distillate.

A bubbler is connected to the vapor-trap and distillate-holding flask that allows impurities and nitrogen gas to exit the system by bubbling through mineral oil without letting air from the outside enter the system.

When a sufficient volume of distillate has been collected, the heating mantle and stirrer are turned off, the water flowing through the condenser and the nitrogen tank are stopped, and valves are closed to isolate the distillation system and distillate-holding flask from each other and from the environment. The distillate-holding flask is then disconnected and transferred to a glove box, which is under slight constant positive pressure from nitrogen gas flowing through it. A humidity monitor showed that the relative humidity of the glove box was consistently between 3–6%.

Meanwhile, the fused silica cells were cleaned with ethanol and baked in an oven to remove any remaining ethanol and any moisture that may have adsorbed onto the inner surfaces of the cell. The cells were then transferred to the glove box as well.

We then mixed the PC, TMB, and primary and secondary wavelength shifters together inside of the glove box and used a syringe to fill each cell with the appropriate cocktail.

A rubber stopper was then placed on the ends of the necks of the filled cells and they were quickly transferred to the setup shown in Figure 6.12. This setup consists of a manifold connecting to a vacuum pump on one end and a mineral oil bubbler on the other end, each of which were separated from the bulk of the manifold by a valve. A needle connecting to a nitrogen gas tank entered the top of the manifold through a septa which kept the inside of the manifold separated from the environment while the needle was inserted. The bottom of the manifold had a threaded port where the stem of the cell could be connected via a coupling.

The valve connecting the vacuum pump to the manifold was initially shut, while the valve leading to the bubbler was opened. The needle connected to the nitrogen tank could then extend down into the cell and bubble nitrogen through the scintillator;



Figure 6.13: A picture of three scintillator cells after filling and sealing them.

the escaping nitrogen could then exit through the bubbler, while the entire manifold held a nitrogen environment.

Once the scintillator had been sparged this way, we removed the needle from the cell and submerged the cell into a liquid nitrogen bath until it froze. Freezing the scintillator allowed us to pull vacuum on the cell without pumping out scintillator vapor. We then closed the nitrogen gas tank, sealed off the valve leading to the bubbler, and opened the valve to the vacuum pump. After evacuating the entire manifold, we used a torch to seal off the cell by melting the stem.

This work was done with much help from M. Souza, who blew the fused silica used to make the cells and sealed them off after they had been sparged.

A picture of three cells after they had been filled and sealed is shown in Figure 6.13.

Measurements

The bell jar was turned on its side and allowed to rotate on Teflon pads so that different cell positions relative to the PMT could be tested, as shown in Figure 6.8. The cell was placed at the flat base of the bell jar, and a ^{54}Mn source was placed outside of the bell jar. The spectrum from the 835 keV ^{54}Mn γ -ray was measured, followed by a measurement of the SPE peak, for each cell. The location of the Compton edge of the spectrum and the SPE peak for each measurement were used to determine the light yield in PE/keV for each cell.

Light yield measurements were taken with the cell across from the PMT for different scintillator cocktail and reflector combinations. These measurements are summarized in Table 6.3. The light yield was found to vary by less than 5% with the position of the cell relative to the PMT. Errors in the light yield measurements come

Table 6.3: Measured light yields of the bell jar in the scintillator cell configuration with a ^{54}Mn γ -ray source and Tyvek or Lumirror as the reflector, in PE/keV. For all measurements the bell jar was filled with air and the scintillator cell was placed opposite the PMT.

Light Yields (PE/keV)		
Cocktail	Tyvek	Lumirror
PC+PPO (1.5 g/L)	0.115 \pm 0.008	0.472 \pm 0.019
PC+TMB+PMP (5 g/L)	0.403 \pm 0.015	—
PC+TMB+PPO (1.5 g/L) + POPOP (25 mg/L)	0.286 \pm 0.014	0.301 \pm 0.011
PC+TMB+PPO (1.5 g/L) + bis-MSB (25 mg/L)	0.267 \pm 0.015	0.336 \pm 0.013
PC+TMB+PPO (3 g/L) + POPOP (25 mg/L)	—	0.408 \pm 0.011
PC+TMB+PPO (3 g/L) + bis-MSB (15 mg/L)	—	0.389 \pm 0.016

primarily from the spread of the Compton edge due to the finite resolution of the system.

These studies show that Lumirror is much more effective than Tyvek 4077D both wet and dry, and that 3 g/L of PPO results in higher light yields than 1.5 g/L of PPO does. However, no significant difference was seen between the use of POPOP and bis-MSB as a secondary wavelength shifter. Since Lumirror reflects well at the emission wavelengths of PPO, as seen in Figure 6.9, we decided to use a mixture consisting of PC (50%) + TMB (50%) + PPO (3 g/L) for future studies.

6.2.4 Filled Vessel Tests

Once we had chosen a scintillator cocktail and reflector combination, the next step was to scale it up to a full prototype detector filled with the scintillator.

Preparation

Before we could fill the bell jar with the scintillator, we first had to thoroughly clean every component to ensure that we would introduce moisture into the system. To do so, we cut strips of Lumirror to line the inner surfaces of the bell jar and the mesh grid, as shown in Figure 6.8. These strips were then cleaned with ethanol and placed in a vacuum to dry. After about one day, the vacuum was back-filled with nitrogen gas, and the Lumirror was left in the nitrogen environment for several days.

The bell jar was scrubbed clean using a cleaning solution consisting of a 1:50 ratio of Detergent 8 to deionized water, and then rinsed down with deionized water to remove remaining detergent. The bell jar was then closed and sealed and a spray ball was inserted into the bell jar through the feedthrough on the lid, and ethanol was repeatedly forced through it at high nitrogen pressure. This process forcibly washed

the inside of the bell jar with ethanol and moisture from the inner surfaces. Nitrogen was then left blowing through the bell jar for a couple of days to further dry the system out.

We then filled the bell jar with argon gas while we quickly removed the lid, installed the Lumirror, and resealed the lid. Nitrogen was then left flowing through the bell jar for several more days. After flowing nitrogen through the entire system, we pumped it down to a vacuum of 2.2×10^{-5} mbar and tested for leaks with a helium leak detector.

To fill the bell jar with PC and TMB, we set up a distillation column similar to the one shown in Figure 6.11. However, instead of distilling in a nitrogen environment, we connected a vacuum pump to the system and distilled under a vacuum. Distilling under a vacuum removes oxygen impurities that distill with the PC and TMB, achieving the same effect as sparging. Additionally, instead of distilling into another three-neck flask, we drained the distillate into a large Erlenmeyer flask.

Before beginning the distillation, the Erlenmeyer flask was filled with the amount of PPO needed to achieve a final concentration of 3 g/L of PPO. The entire system was left under nitrogen flow for a day to dry it out before beginning distillation.

PC and TMB were then distilled into the flask until it was approximately halfway full. The flask was then shaken until the contents were well-mixed, and we then poured the contents directly into the bell jar, which was connected to the flask. We repeatedly distilled more PC and TMB this way until we had distilled a total of 10 L of each.

Measurements

The entire bell jar was stored inside a house of lead bricks to reduce background. The light yield in the bell jar was measured over the span of 42 days using a $1 \mu\text{Ci}$ ^{54}Mn γ -ray source placed on the top of the bell jar near the center. After each γ -ray spectrum measurement, the source was removed and background and SPE spectra were recorded.

The light yield was determined by considering a GEANT4 simulation of the energy deposited in the bell jar from this γ -ray source. The background-subtracted γ -ray spectrum was then scaled by the SPE peak. The simulated γ -ray spectrum was convolved with a Gaussian response function which accounted for the resolution and light yield of the detector as fit parameters. This function was then fit around the measured Compton edge, as shown in Figure 6.14.

Shortly after filling the bell jar, we measured a light yield of 0.466 ± 0.001 PE/keV. However, the light yield was found to steadily decrease over time at an average rate of $0.52 \pm 0.02\%$ /week. This steady decrease can be seen in Figure 6.15. This decrease may be due to PC reacting with metal oxides in the stainless steel walls of the bell jar, as discussed in Section 6.1 and shown in Figure 6.4. While a direct comparison between the effects observed in Figure 6.4 is hard to do because those studies used pure PC, an order of magnitude estimate shows that the rate of degradation observed in the bell jar is roughly consistent with the degradation seen in this figure.




Figure 6.14: A sample ^{54}Mn spectrum measured in the bell jar, background subtracted and scaled to the number of photoelectrons by the SPE peak. (*Red*) A smoothed simulated spectrum has been fit to the data around the Compton edge; *blue* the simulated spectrum before smoothing. The large narrow spike around 390 PE is the 835 keV full energy peak, while the broader spike just below 300 PE is the Compton edge. The part of the spectrum beyond the Compton edge is due to multiple scattering. The fit was performed specifically around the Compton edge in order to avoid effects from ionization quenching.

Neutron Signals

Having measured the light yield of the bell jar, it is interesting to see the bell jar's response to neutron signals.

To measure the neutron spectrum, we put an $^{241}\text{AmBe}$ (α, n) neutron source outside of the lead housing surrounding the bell jar, so that the lead would block many of the γ -rays coming from the source.

For these measurements, the PMT signal was fed to a CAEN V1720 250 MHz digitizer and was recorded and analyzed using the daqman software developed by B. Loer [168].

Measurements of the $^{241}\text{AmBe}$ spectrum were followed by measurements of the ^{54}Mn Compton spectrum and the SPE peak to calibrate the energy scale of the detector.

For the $^{241}\text{AmBe}$ runs, we expect two signals to be produced by the neutrons, as previously discussed. The prompt signal is a combination of the neutron thermalization signal and the 3–4 MeV γ -rays in prompt coincidence with the $^{241}\text{AmBe}$ neutrons. The capture signal is expected to follow the prompt signal with a time




Figure 6.15: The light yield (in photoelectrons per keV) measured in the bell jar over 42 days, neglecting the exponential tail in the single photoelectron distribution, decreasing at a rate of $0.52 \pm 0.02\%$ per week.




Figure 6.16: An example ADC waveform of a neutron capturing on ^{10}B in the bell jar. The thermalization signal can be seen around $0 \mu\text{s}$, and the neutron capture signal is seen around $1.5 \mu\text{s}$. The blue curve on the waveform shows the total integral of the curve from the start of the acquisition window to a given time.




Figure 6.17: The background-subtracted prompt energy spectrum from the AmBe source. The peak at high energy is the Compton edge of the 4.4 MeV γ -ray produced by the ${}^9\text{Be}(\alpha, n){}^{12}\text{C}$ reaction with a branching ratio of $\sim 64\%$.

constant of $2.2\,\mu\text{s}$. An example waveform of a neutron thermalizing and capturing can be seen in Figure 6.16. In this figure, the prompt signal can be seen around $0\,\mu\text{s}$, and the capture signal can be seen around $\sim 1.5\,\mu\text{s}$.

In order to see both the prompt and capture signals of the neutron, we record $13\,\mu\text{s}$ worth of data with each trigger.

When selecting events for analysis, we made two primary cuts. First, we required that the digitizer not be saturated by any pulse in the event. Second, we required that at least two pulses be observed.

For selected events, we then selected pulses from within the ADC waveform by identifying signals where the ADC went below a threshold of a few PE in amplitude. We then grouped together pulses that were within $10\,\text{ns}$ of each other in order to group together multiple scatters of the same neutron. We also required that the delayed pulse follow the prompt pulse by at least $0.65\,\mu\text{s}$ to remove contamination from afterpulses, which peak around $0.55\,\mu\text{s}$ after the first pulse. Because of the high rate of low energy noise, we accepted only the highest energy pulse in this time window for the delayed spectrum.

After analyzing data in this way and subtracting a measured background spectrum, we obtained the prompt and delayed spectra shown in Figures 6.17 and 6.18, respectively. The prompt spectrum in Figure 6.17 is mostly dominated by the 4.4 MeV γ -ray produced in 64% of all ${}^9\text{Be}(\alpha, n){}^{12}\text{C}$ reactions. The peak at higher energies is the Compton edge, at an energy of 4.2 MeV.




Figure 6.18: The background-subtracted delayed energy spectrum from the AmBe source. The peak at ~ 30 keV_{ee} represents the energy deposit of the ${}^7\text{Li}$ and α alone, and the broad peak around 400 keV_{ee} represents the Compton spectrum of the γ -ray plus the ${}^7\text{Li}$ and the α .

Since the bell jar is relatively small compared to the mean free path of the 478 keV γ -ray produced by the neutron capture on ${}^{10}\text{B}$, the signal from this γ -ray will be highly suppressed, and the Compton spectrum from this γ -ray will dominate its full energy signal. The delayed spectrum in Figure 6.18 shows a broad peak at ~ 400 keV_{ee}. This peak corresponds to the neutron capture on ${}^{10}\text{B}$ that goes to ${}^7\text{Li}^*$, and the light collected in this peak is from the Compton spectrum of this 478 keV γ -ray plus the light produced by the α and ${}^7\text{Li}$. The lower energy peak comes from the α and ${}^7\text{Li}$ alone, with results from the ${}^{10}\text{B}$ capture that goes directly to the ground state of ${}^7\text{Li}$, or cases where the γ -ray from ${}^7\text{Li}^*$ escapes without leaving any signal. This peak occurs around 30 keV_{ee}.

The α and ${}^7\text{Li}$ from the neutron capture branch that goes to ${}^7\text{Li}^*$ are at a lower energy than those from the capture that goes directly to the ground state of ${}^7\text{Li}$. This lower energy means that these ions will also be more heavily quenched than in the higher energy case. Since 94% of neutron captures on ${}^{10}\text{B}$ go through the ${}^7\text{Li}^*$ state, we expect these lower energy products to dominate the signal seen in the bell jar, where most of the 478 keV γ -rays will escape. It is therefore possible that the peak position of 30 keV_{ee} measured here is an underestimation of the actual quenched energy. Furthermore, since these measurements were taken several months after the bell jar is filled, it is possible that the impurities that cause the scintillator to degrade also cause ions to be more heavily quenched. In this case, a larger detector with a purer scintillator can expect to see a stronger signal from this low energy peak.

Nevertheless, the 30 keV_{ee} signal is high enough that a detector with a comparable light yield to Borexino should be able to efficiently detect neutron capture signals.

6.3 Neutron Beam Measurements

Having established that the LSV design is likely able to see the neutron capture signal, the next step is to determine if the design is likely able to detect the thermalization signal.

Since the ²⁴¹AmBe source used on the bell jar produces a very broad neutron energy spectrum and is overwhelmed by γ -ray backgrounds, it is very difficult to characterize the neutron thermalization signal from that source alone. Instead, we will fall back upon the Birks model of neutron quenching, given in Equation 5.2. Our goal is to generally characterize the quenching of protons in the liquid scintillator by determining Birks constant kB .

To accomplish this goal, we used the accelerator complex at University of Notre Dame, in a setup similar to those described in [92, 169]. Much of the work done in this section was done in close collaboration with J. Xu, E. Shields, and F. Froberg. This data was also collected with significant help from W. Tan, H. Beck, and the crew of University of Notre Dame beam operators.

6.3.1 Liquid Scintillator Detector

For these studies, a scintillator cell containing equal parts PC and TMB and 3 g/L of PPO was produced as described in Section 6.2.3. A diagram of the stainless steel canister holding the detector is shown in Figure 6.19. This cell was placed in a Lumirror cup (not shown) that covered the lateral sides and the base of the cell with the stem; the flat base was optically coupled to a Hamamatsu R11065 PMT. The cell sat on top of a Teflon cup with a hole cut in it through which the stem of the cell could sit. High voltage and signal cables went from the base of the PMT to two feedthroughs at the top of the canister, where they were attached to the data acquisition system and power supply. The entire structure was held in place by a stand filling the vertical length of the canister and two support disks attached by a spring holding the PMT, Teflon cup, and cell together.

The outer region of the canister near where the cell rested was shaved down to a thinner thickness. This helped mark where the cell was from the outside and also reduced the amount of mass that a neutron would pass through before and after entering the scintillator cell.

6.3.2 Neutron Beam Setup

The neutron beam was made by putting a LiF target at the end of one branch of the University of Notre Dame 11 MeV FN tandem proton accelerator. Protons accelerated to 2.44 MeV would collide with the target and undergo the ⁷Li(p,n)⁷Be reaction on the




Figure 6.19: The canister holding the scintillator cell for measurements at the University of Notre Dame neutron beam. (*Cyan*) The scintillator cell, (*yellow*) the PMT, (*gray*) the canister holding everything in place, (*red*) high voltage and signal cables, (*white*) Teflon support cup, (*purple*) support structure and springs holding everything in place.

target with a Q-value of -1.644 MeV. The energy of the outgoing neutron is correlated with its momentum relative to the proton beam.

The rate was controlled by a three-part pulsing system with a time resolution of 2 ns and a period of 101.5 ns. By selecting the fraction of proton pulses that went down the branch with our detector, we could further control the rate of neutrons; a branching ratio of n means that 1 out of every n pulses sends a bunch of protons down the relevant branch. Data reported here was collected using a branching ratio of 6 or 8.

The proton beam has a cross sectional diameter of 3 mm and is stable in energy to within 1 keV.

Since the energy of the outgoing neutron is correlated with its direction, a 22 cm diameter by 22 cm long polyethylene collimator with a 2.5 cm diameter hole was put between the end of the beam and the detector. This collimator selects for forward-scattering neutrons, which will have the highest energy coming out of the LiF target. A picture and a diagram of this setup are shown in Figure 6.20.

After passing through the collimator, neutrons may scatter in the liquid scintillator detector. Then, the neutron may scatter a second time in one of the coincidence detectors. Two types of coincidence detectors were used for this experiment: $2''$ diameter by $2''$ long Eljen 510-20 \times 20-9/301 and $5''$ diameter by $5''$ long Eljen 510-50 \times 50-1/301. Both of these detectors use the commercial reflector EJ-510 and scintillator EJ-301,



2" Coincidence Detectors

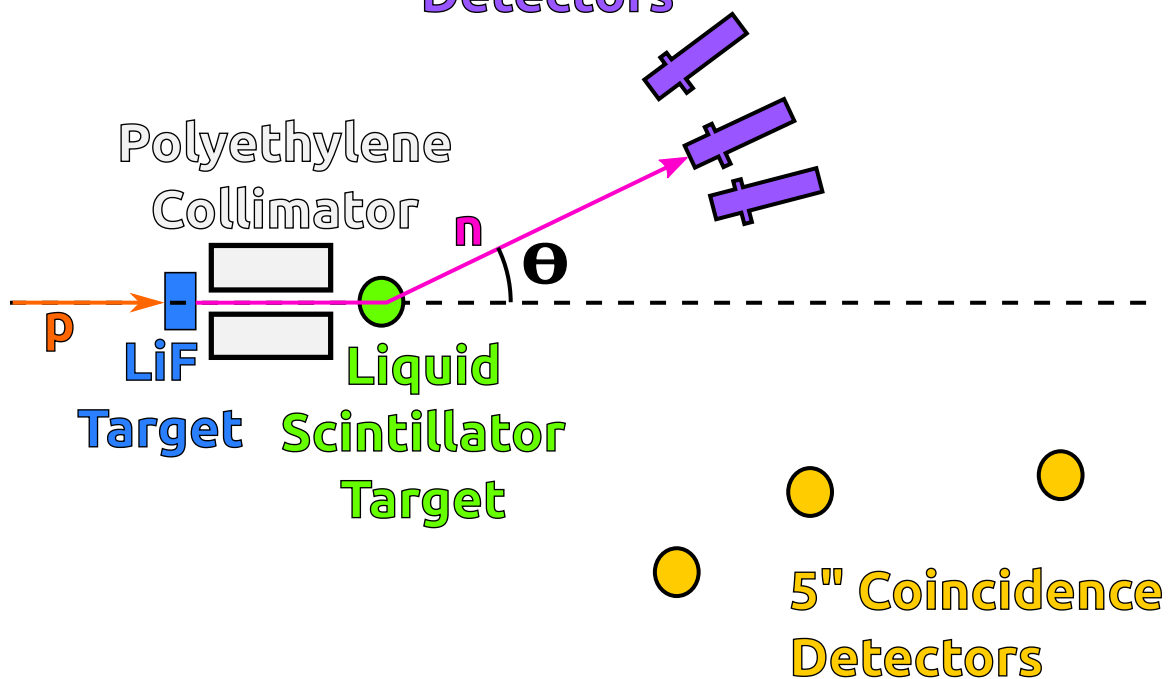


Figure 6.20: (Top) A picture of the detectors arranged around the neutron beam. (Bottom) A diagram of the experimental setup.

Table 6.4: Beamline angles and distance between the center of the liquid scintillator detector and the coincidence detectors, and the corresponding $p(n, n')p$ scattering expected recoil energy for a 690 keV neutron.

Detector	Diameter (inch)	Angle (deg)	Distance (cm)	Recoil Energy (keV)
1	5	59.02	150.3	507.22
2	5	41.26	148.8	300.10
3	5	24.88	200.0	122.16
4	2	48.04	69.9	381.51
5	2	32.37	69.6	197.75
6	2	18.13	69.7	66.83

which is a xylene-based liquid scintillator with organic fluors. These detectors have a typical light yield $\sim 1 \text{ PE/keV}_{ee}$.

Since the positions of the coincidence detectors and the liquid scintillator detector have all been precisely measured, the neutron scattering angle can be inferred from the coincidence detector that was hit. Since the neutron’s kinetic energy falls within a narrow band, the neutron’s scattering angle θ off of a proton in the liquid scintillator is kinematically related to the energy it deposited in the scintillator by

$$\Delta T_n = T_n \sin^2(\theta) \quad (6.2)$$

where T_n and ΔT_n are the kinetic energy of the neutron before scattering and the recoil energy of the proton, respectively.

The liquid scintillator detector was placed 50 cm from the end of the LiF target, with a measured offset perpendicular to the beamline of 0.2 cm. The positions of the liquid scintillator detector and the 2” coincidence detectors were measured by lowering a plum-bob from the center of the detector to the ground, and marking the ground. For the 5” coincidence detectors, it was not possible to measure the position of the center of the detector directly; instead, a plum-bob was used to trace out points around the outer diameter of the detector, and the center was inferred from these points. A yard stick was then used to measure distances and angles between the liquid scintillator detector and each of the coincidence detectors.

A list of measured recoil angles relative to the beamline, measured distances from the liquid scintillator detector, and estimated proton recoil energies is given in Table 6.4.

In this table, the expected recoil energy is inferred from Equation 6.2. The uncertainty in the positions of the liquid scintillator detector and the 2” coincidence detectors is estimated to be $\sim 1 \text{ cm}$. The uncertainty in the position of the 5” coincidence detectors is estimated to be $\sim 2 \text{ cm}$. This greater uncertainty is due to the extra step needed to infer the position of the center of the detector. Because the proton recoil angle is correlated with the recoil energy, uncertainty in the position of the coincidence detectors (and hence the recoil angle) translates to uncertainty in the expected recoil energy. By calculating the spread around each of these points for the given uncertainties in the positions, we estimate the uncertainties in the recoil




Figure 6.21: Simulation of 2.44 MeV protons propagating in the LiF target. (*Left*) Proton energy versus depth. (*Right*) Proton angle relative to the beamline versus depth.

energies due to this effect to range from 1.5–12%, depending on how quickly the recoil energy varies with the recoil angle at each angle.

Selecting the distance between the liquid scintillator detector and the coincidence detector is a matter of balancing two competing factors: the larger the solid angle subtended by the coincidence detector, the higher the event rate; however, a larger solid angle also means more uncertainty in the recoil angle (and hence recoil energy) of the target nucleus. Distances for each detector were chosen to roughly balance these factors, additionally accounting for the probability of a neutron scattering to a give angle. 5" detectors at more oblique angles where the neutron is less likely to scatter to are therefore placed closer to the detector.

6.3.3 LiF Target

In order to predict the proton recoil energy at each angle, it is necessary to first understand the neutron spectrum emitted from the LiF target.

The LiF target has a thickness of 0.52 mg/cm^2 ($\sim 0.0002 \text{ cm}$). While the spread in proton energy is very narrow, protons incident on the target will travel some distance into the target before they undergo the (p, n) . Energy loss before they capture may cause protons to capture at a broader energy range and at nonzero angles relative to the beamline. In order to understand this effect, we simulated protons propagating in the target using GEANT4 [170]. Figure 6.21 shows the results of these simulations; as the protons propagate further into the target, their mean energy decreases, while the spread in energy increases. However, most protons remain with 1.5° of the beamline.

In order to determine the cross section for protons undergoing these interactions at each depth, we used ${}^7\text{Li}(p, n){}^7\text{Be}$ cross sections measured by Burke et al. [38].




Figure 6.22: Energy-angle relationship for outgoing neutrons in the ${}^7\text{Li}(p,n){}^7\text{Be}$ reaction with a 2.44 MeV proton.

These measurements give the cross section for protons of a given energy to produce a neutron at a certain angle relative to the incoming proton's momentum. The angle of the outgoing neutron θ is kinematically related to its kinetic energy T_n , and can be calculated from relativistic two-body reaction kinematics using the equation,

$$\begin{aligned}
 T_n^2 &= \left(\frac{\sqrt{m_n^2 + p_{\text{cm}}'^2} \cos \theta \sinh \chi + \cosh \chi \sqrt{p_{\text{cm}}'^2 - (m_n \sin \theta \sinh \chi)^2}}{1 + \sin^2 \theta \sinh^2 \chi} \right)^2 - 2T_n m_n \\
 s &= (m_p + m_{\text{Li}})^2 + 2m_{\text{Li}} T_p \\
 \chi &= \ln \left| \frac{p_{\text{cm}} + \sqrt{m_{\text{Li}}^2 + p_{\text{cm}}^2}}{m_{\text{Li}}} \right| \\
 p_{\text{cm}} &= \sqrt{\frac{(s - m_p^2 - m_{\text{Li}}^2)^2 - 4(m_p m_{\text{Li}})^2}{4s}} \\
 p_{\text{cm}}' &= \sqrt{\frac{(s - m_n^2 - m_{\text{Be}}^2)^2 - 4(m_n m_{\text{Be}})^2}{4s}}
 \end{aligned} \tag{6.3}$$

where T_p is the mass of the incident proton, m_n , m_p , m_{Be} , and m_{Li} are the masses of the neutron, proton, ${}^7\text{Li}$, and ${}^7\text{Be}$, respectively, s is a Lorentz invariant quantity, χ is the rapidity, and p_{cm} and p_{cm}' are the center of mass momenta of the reactants and products, respectively. A clear and concise derivation of this calculation is given by S. Sjue in [171].

The relationship between neutron energy and angle for a 2.44 MeV incident proton is shown in Figure 6.22.

Figure 6.23 shows the cross section for this reaction as a function of the proton kinetic energy and either the neutron beamline angle or kinetic energy. Burke et al. give the cross sections as a function of the outgoing neutron angle; to convert




Figure 6.23: ${}^7\text{Li}(p,n){}^7\text{Be}$ cross section data from Burke et al. [38]. In both of these plots, the color represents the cross section in mb/sr. (*Left*) Neutron angle with respect to the proton versus the proton energy versus reaction cross section. (*Right*) Neutron kinetic energy versus proton energy versus cross section, made by transforming the left plot as described in the text.

this into a map of the neutron’s kinetic energy, we transform this data according to Equation 6.4.

After combining the data in Figures 6.21 and 6.23 and adding the proton beamline angles to the angle of the neutron with respect to the proton momentum, we can determine the neutron kinetic energy versus beamline angle distribution.

Figure 6.24 shows the energy and angular distribution of neutrons coming from the LiF target in this experimental setup. Forward-scattering neutrons, which we expect to be most likely to make it past the polyethylene collimator, have a mean kinetic energy of ~ 690 keV.

6.3.4 Electrons and Data Acquisition

The data acquisition system used for these measurements is illustrated in Figure 6.25.

The data acquisition system is designed to trigger whenever there is a coincidence between the liquid scintillator detector and any one of the coincidence detectors. Additionally, a periodic signal produced by the proton beam pulser gives timing information that differs by an approximately constant offset from the time at which a neutron is produced in the LiF target. This information allows us to make time of flight cuts that are useful for particle identification: by selecting a window around the expected time of flight for neutrons, we can remove random coincidence backgrounds as well as backgrounds from γ -rays produced in the beamline or LiF target in coincidence with the neutron.




Figure 6.24: Neutron energy and angular distribution from the ${}^7\text{Li}(p, n){}^7\text{Be}$ reaction of a 2.44 MeV proton on a 0.52 mg/cm^2 LiF target.

Signals produced by the pulser are read by a LeCroy 621AL quad discriminator set to a 423 mV threshold. Logic pulses produced by this discriminator are read out by a CAEN V1720 digitizer (12 bit, 250 MS/s) and, when the digitizer receives a trigger signal, recorded to a computer, starting $2 \mu\text{s}$ before the trigger was received.

The PMT viewing the liquid scintillator detector is operated at 1.4 kV and 0.4 mA, supplied by an Ortec 556 High Voltage Power Supply. Signals from this PMT are passed through a Gran Sasso low noise $\times 10$ amplifier module provided by collaborators at LNGS. Amplified signals are then fed to a LeCroy 428F linear Fan-In/Fan-Out module, which supplies signal to the CAEN V1720 digitizer as well as to a Phillips Scientific 711 6 channel discriminator set to a 10 mA threshold.

The coincidence detectors were operated between 1.15–1.6 kV, supplied by a LeCroy HV4032A Multi-Channel High Voltage Power Supply. Signals from these detectors were passed to a Phillips Scientific 779 32 Channel $\times 10$ Amplifier module and then to a LeCroy 428F Linear Fan-In/Fan-Out module. One copy of each of the signals was then read out by the CAEN V1720 digitizer while the other went to one of two discriminators. Coincidence detectors 2–6 went to the same Phillips Scientific 711 discriminator as the liquid scintillator detector, while detector 1 went to the LeCroy 621AL detector that also read out the pulser signal. Discriminator thresholds were set to a voltage equivalent to ~ 2 PE on each channel.

Logic pulses read out from both discriminators from the coincidence detectors were all read out by a Logical Fan-In/Fan-Out module which functioned as an OR gate, which would produce a logic pulse whenever any of the coincidence channels produced a signal in the discriminators. This pulse was then fed to a LeCroy 375L Coincidence module, which also read out the logic pulses produced by the discriminator reading the liquid scintillator detector channel. This coincidence module functioned as an AND gate with a 400 ns coincidence window. Whenever the liquid scintillator detector and

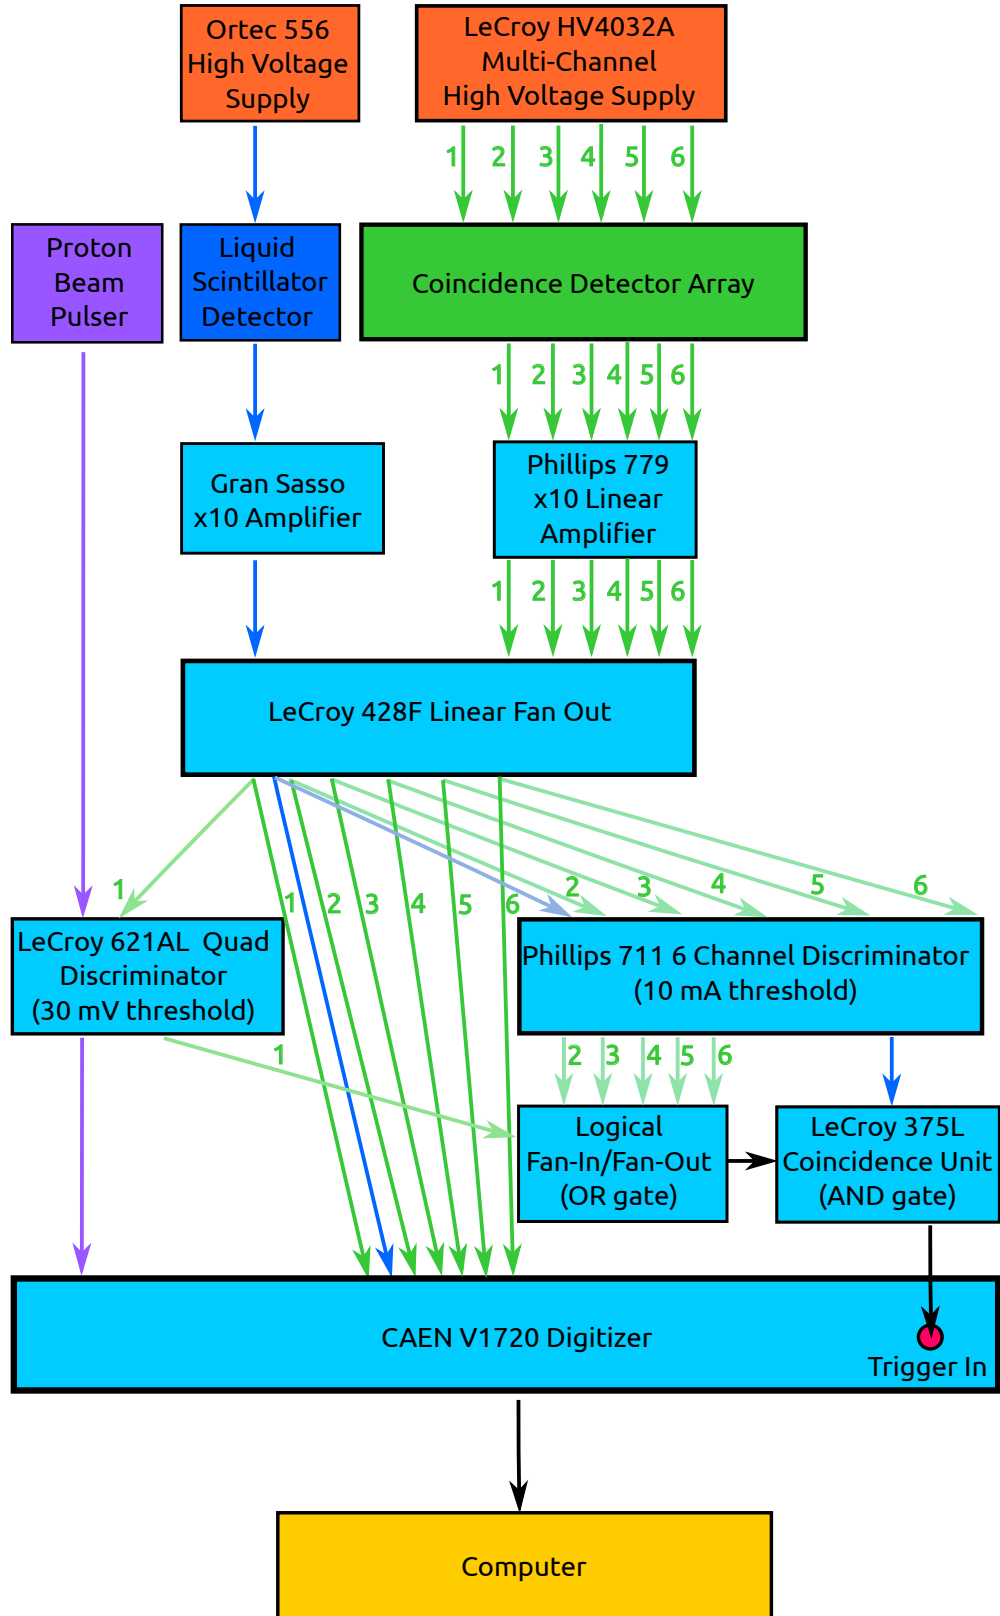


Figure 6.25: Schematic diagram of the data acquisition system used for the neutron beam measurements.

any of the coincidence detectors produced a signal in the AND gate within 400 ns of each other, this AND gate would produce a logic pulse.

The logic pulse produced by the AND gate was then supplied to the digitizer as a trigger signal, telling the digitizer to record the waveforms from each detector. Data was recorded and processed using the daqman software written by B. Loer [168].

6.3.5 Stability

In order to get enough statistics, we had to collect data at the neutron beam for most of a day. Since we later combine data collected through the entire day, it is necessary to study the stability of the detector to avoid biasing the data from drifts in detector parameters.

While the proton beamline energy stability was monitored by the operator, it was important to monitor to ensure the stability of the liquid scintillator detector light yield and the SPE mean of the PMT.

To measure the SPE mean of the PMT, we looked at the tail of scintillation pulses, where scintillation photons separated enough in time that the resulting photoelectrons can be readily distinguished. By measuring the charge of many such pulses, we can generate a SPE charge distribution, similar to the one shown in Figure 2.10. Measuring these peaks over the course of the experiment showed that the SPE mean remained stable to within $\sim 3\%$, with no trend of how the SPE mean changed over time. Furthermore, the relative SPE variance was measured to be $\sim 10\%$ for the PMT used.

The light yield was measured with a ^{241}Am source, which emits a 59.54 keV γ -ray in 36% of all decays. To calculate the light yield of the detector, we taped the ^{241}Am source to the outside of the canister and recorded the γ -ray spectrum in the detector. Doing so produced a Gaussian distribution corresponding to the full energy peak of the γ -ray with a mean of 81.5 PE. This corresponds to a light yield of 1.368 ± 0.002 PE/keV. To test the stability of the light yield over the duration of the experiment, we measured the light yield once before collecting neutron data and once after. Both calibration runs showed consistent light yields.

6.3.6 Thermalization Signal

The time difference between the buncher signal and the liquid scintillator detector typically shows two peaks: one corresponding to γ -rays produced in the $^7\text{Li}(p, n)^7\text{Be}$ reaction scattering in the detector, and another corresponding to neutrons doing so. Similarly, the difference in trigger times between the liquid scintillator detector and each coincidence detector also show peaks corresponding to coincident γ -rays and neutrons, as well as random γ -ray backgrounds.

By making time of flight cuts, we can select events in which neutrons scattered in the liquid scintillator detector and exactly one of the coincidence detectors. Doing so, we obtain the energy spectra shown in Figure 6.26.

While, as discussed in Equation 6.2 and Table 6.4, the proton recoil energy can be calculated if the scattering angle is known, there is some background from events

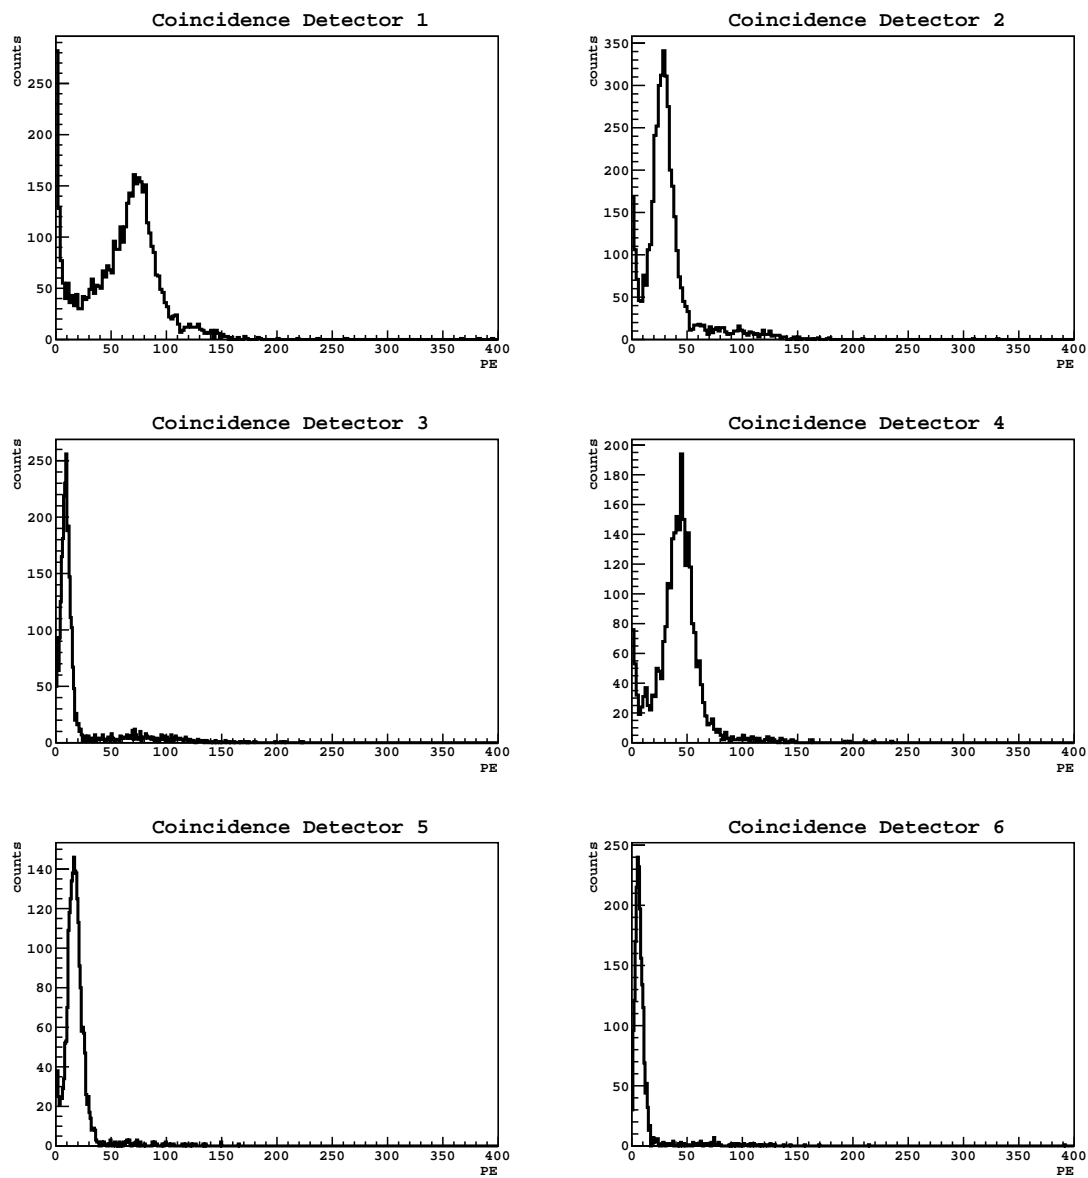


Figure 6.26: Spectra in the liquid scintillator detector for events with coincidences in coincidence detectors 1–6.

Table 6.5: Summary of expected, simulated, and observed recoil energies in the liquid scintillator detector when there is a coincidence with each coincidence detector.

Detector	E_r^{exp} [keV]	E_r^{sim} [keV]	E_r^{obs} [PE]
1	507.2	485.5	73.4
2	300.1	288.9	28.3
3	122.2	117.7	8.9
4	381.5	362.1	43.7
5	197.8	182.8	17.2
6	66.8	66.9	6.1

in which a neutron scatters multiple times in the liquid scintillator detector before scattering in a coincidence detector. Furthermore, energy lost in the detector materials and spread in scattering angle due to the angle subtended by each detector and uncertainty in the detector positions may cause inaccuracies in these predicted recoil energies.

In order to understand these spectra, we ran a detailed simulation using the GEANT4 toolkit of this experiment. The energy and angle distribution of neutrons produced in this simulation were drawn from the distribution shown in Figure 6.24. To save disk space, we only recored events in which energy was deposited in the liquid scintillator detector and at least one of the coincidence detectors. We then analyze the simulation data replicating the data analysis cuts as closely as possible in the simulation. After doing so, we obtained a series of energy deposition spectra in the liquid scintillator detector corresponding to events in which there was a coincidence in exactly one of each of the six coincidence detectors.

The expected recoil energy E_R^{exp} , the peak recoil energy in the simulation E_R^{sim} , and the peak number of PE detected for in the liquid scintillator detector for a coincidence in each of the coincidence detectors E_R^{obs} is given in Table 6.5.

In order to determine the proton quenching factor at each recoil energy, we convolved the simulated spectra with a Gaussian response function. For each energy bin E in the simulated spectra, the response function had an overall multiplicative factor accounting for the event rate, a constant vertical offset assuming a constant background rate, and a mean and variance given by

$$\begin{aligned}\mu &= E \times LY_{\text{rel}} \\ \sigma^2 &= (1 + \sigma_{\text{spe}}^2) \times E \times LY_{\text{rel}}\end{aligned}\tag{6.4}$$

where σ_{spe}^2 is the measured relative SPE variance and LY_{rel} is the relative light yield for proton recoils corresponding to the given recoil angle. The overall rate, constant offset, and relative light yield were left as free parameters and the simulated spectra was fit to the data for each coincidence spectrum.

Since quenching causes the detector response to become non-linear, each spectrum was fit around the proton recoil peak only. Furthermore, some background exists from proton recoils on carbon. Since carbon is quenched much more heavily than hydrogen due to its larger mass, the carbon peak is farther separated from the proton recoil




Figure 6.27: An example proton recoil spectrum with a simulated spectrum fit to it. The example shown here is for coincidence detector 2. (*Black*) the measured spectrum. (*Red*) The fit simulated spectrum. The departure of the fit spectrum from the measured spectrum is due to carbon recoils, which are quenched more heavily than was accounted for in this simulation; the simulated spectrum was fit to the measured spectrum above where we see the carbon recoil contribution to the spectrum.

peak in the data than we see in the simulation, which does not treat the two nuclei differently. However, by limiting the fits to the proton recoil peaks, we were able to avoid contributions from these backgrounds.

An example of such a fit can be seen in Figure 6.27, showing the simulated spectrum fit to the measured spectrum for coincidence detector 2. In this example, the simulated spectrum was fit to the measured spectrum between 10–50 PE, though the full spectrum was drawn for illustrative purposes. As can be seen here, there is a low energy peak in the simulated spectrum that deviates from the data. This is due to neutron recoils off of carbon, which are more heavily suppressed than proton recoils and therefore not accounted for in this fit.

Assuming that the ^{241}Am γ -ray was quenched very little, we can determine the quenching factor corresponding to each of these fit relative light yields by dividing the relative light yield from each fit by the light yield measured from the ^{241}Am calibration runs.

Similarly, we can determine the recoil energy corresponding to each fit quenching factor by dividing the peak measured in the data by the fit relative light yield. Doing so, we obtain the distribution quenching factors at each sampled recoil energy as shown in Figure 6.28. For comparison, we also show quenching factor measurements reported by Hong et al. [39] for a different organic liquid scintillator.




Figure 6.28: (*Black dots*) Quenching factor measurements in the liquid scintillator detector. (*Red triangles*) Measurements reported by Hong et al. [39] for a different organic liquid scintillator for comparison. (*Blue and green curves*) Two different quenching models (see text) fit to the measurements presented here.

Table 6.6: Summary of fits of scintillation quenching Models I and II explored here.

Model	kB [cm/MeV]	C [cm ² /MeV ²]	χ^2/NDF
I	0.0153 ± 0.0002	—	19.4
II	0	$(2.19 \pm 0.04) \times 10^{-5}$	1.6

6.3.7 Modeling

Based on the measured quenching factors presented here, we discuss three different scintillation quenching models.

$$\text{I) } \frac{dF}{dx} = \frac{S \frac{dE}{dx}}{1 + kB \frac{dE}{dx}}$$

$$\text{II) } \frac{dF}{dx} = \frac{S \frac{dE}{dx}}{1 + kB \frac{dE}{dx} + C \left(\frac{dE}{dx} \right)^2}$$

$$\text{III) } \frac{dF}{dx} = \frac{S_e \frac{dE}{dx} \Big|_e + S_n \frac{dE}{dx} \Big|_n}{1 + kB_e \frac{dE}{dx} \Big|_e + kB_n \frac{dE}{dx} \Big|_n}$$

where $\frac{dF}{dx}$ is the amount of light produced per unit track length of the scintillating particle, $\frac{dE}{dx}$ is the energy lost by the particle per unit length, S is the scintillation yield, kB is Birks' constant, and C is another constant characterizing the quenching.

Subscripts e and n represent specific electronic and nuclear contributions to their respective terms.

Model I is the standard form of Birks' law, as described in Chapter 5.2.1. Model II is a semi-empirical modified version of Birks' law presented by Craun and Smith [172], where C is a fitted parameter. Lastly, Model III is a model explored by Hong et al. [39] that treats scintillation and quenching from nuclear and electronic energy loss separately.

In Figure 6.28, we fit Models I and II to our measured quenching factors, where stopping powers were calculated using SRIM [96]. Model III was not included in this analysis; protons predominantly lose their energy to electrons, and so the nuclear energy loss terms are too small to be measured in this data. Quenching measurements of heavier ions are needed to adequately explore this model.

In Table 6.6, we can see the parameters from the fit of each model to the data, as well as the reduced χ^2 statistic for each fit. As we can see, Model I does not describe the data very well. Furthermore, when we fit Model II, we find that kB goes to zero and the quadratic quenching term dominates. We find that this model describes the data very well, though because the quadratic term in the denominator of Model II dominates the linear term over the range of proton recoil energies studied here, this measurement is not sensitive to kB .

Given this level of quenching in the scintillator, we may expect that the signal produced by the thermalization of a fast neutron—typically in the 1–10 MeV energy range—to be detectable by this LSV design. As will be discussed in Chapter 9.3, we see that Model II describes the proton scintillation of the DarkSide-50 scintillator cocktail well for proton recoils using roughly similar parameters; however, we find that a single set of parameters cannot simultaneously describe the quenching of all particles in the scintillator, indicating that this model is still an incomplete description of the quenching mechanism.

Chapter 7

Optical Simulations

In Chapter 6, we discussed a series of prototype studies that were done to choose a design for the DarkSide-50 LSV. While these prototypes all showed promising results, it is important to consider how the signals might change when we scale the detector from the small prototype to the much larger full LSV.

In order to determine if the larger LSV will perform well enough to be an effective neutron veto, we performed a series of optical simulations.

In this chapter, we will briefly discuss the general framework for these optical simulations, and then we will go on to discuss simulations done of the prototype bell jar detector which were used to validate the framework, and the simulations done to model the optics of a larger detector. Details of these simulations as well as an independent set of custom-built optical simulations performed by E. Shields are presented in [13].

7.1 Framework

The optical simulations discussed here were built upon the GEANT4 Monte Carlo simulation toolkit [170], primarily using GEANT4-9.5.0.

GEANT4 simulates the propagation and interactions of particles through materials, using both data (either supplied directly by the user or compiled from external databases) and theoretical models to predict interaction cross sections. The user can specify a geometry, including material compositions and optical properties, and a set of initial conditions (*i.e.* which particles will exist at what positions and momenta at the start of each simulated event). GEANT4 will then track each particle as travels through the specified geometry until it falls below some user-set threshold energy or leaves the geometry.

At each step in each particle’s track, GEANT4 determines the set of possible physical processes. Then, based on the given and calculated cross sections and user-specified data, GEANT4 determines a length scale for each process. For most processes, the distance that a particle will travel before it undergoes that process follows an exponential probability distribution. For each of these processes, GEANT4 draws an “interaction length” from the process’s associated exponential distribution. How-

ever, some processes, such as boundary processes, will happen at a fixed distance from the particle’s position at the start of the track. For these processes, GEANT4 determines what the fixed interaction length is. GEANT4 then chooses the process with the shortest interaction length, applies the process to the particle, changing its energy and momenta and incrementing the clock as appropriate, and then chooses the next process.

Some processes will create secondary particles. In these cases, if the energy of the secondary particle is above some user-specified threshold, the secondary particles are added to the stack and are tracked as well.

Optical processes typically begin with a process that converts energy lost by one particle into light, such as the Cherenkov and scintillation processes. GEANT4 treats γ -rays and low energy photons—“optical photons”—as completely different particles. “Optical photons” are those typically on the eV scale where PMTs may be sensitive. These photons usually undergo very different processes from the typically much higher energy γ -rays, such as absorption and emission by wavelength shifters, refraction upon entering a new material, and reflections off of boundaries.

Once a particle has undergone a process that creates optical photons, GEANT4 will determine the number of optical photons to produce based on the user-specified material properties and the physics of the process creating them. These optical photons will be added to the stack and propagated until they are either absorbed or escape the geometry.

Since light-producing processes frequently produce thousands of optical photons, and these photons may undergo many steps before their tracks end, optical simulations are frequently very slow to perform and may be very memory-intensive.

In general, the micro-physics underlying these optical processes are too complicated to model from first principles. Instead, GEANT4 takes a higher-level approach: the user specifies optical properties of all relevant materials, including refractive indices, absorption lengths, and reflectivities. Because these parameters are often not known for many materials, assumptions often have to be made to run the simulation, adding a lot of intrinsic uncertainty to the results. Additionally, because these parameters often affect observable quantities in complicated, overlapping ways, tuning too many parameters becomes difficult.

For the simulations presented here, we use measured values for as many parameters as possible to avoid tuning quantities.

7.2 Prototype Validation

Based on all of the measurements discussed in Chapter 6 and measurements the Borexino collaboration has made of their similar scintillator [164], we constructed an optical model of the prototype bell jar in GEANT4.

7.2.1 Optical Model

These simulations included the stainless steel bell jar vessel, the Lumirror reflector with the measured reflectance shown in Figure 6.6 (modeled as a boundary between the bell jar and its contents), the PMT, and the quartz window around it. The quartz window was assigned a refractive index of 1.45, and the liquid scintillator was given a refractive index of 1.43, the average refractive index of PC and TMB.

The R11065 PMT was given the quantum efficiency reported by Hamamatsu [14], and the photocathode was assigned a diameter of 6.51 cm and a collection efficiency of 85%. When a photon hit the photocathode, it was detected with a probability equal to the product of the quantum efficiency and collection efficiency; if it was not detected, it was given a 20% chance of reflecting off the photocathode (this number is an approximate value given to us in personal communications with Hamamatsu). If it is neither detected nor reflected, the optical photon is absorbed and the track is terminated.

Since GEANT4 evaluates whether or not a photon reflects before it evaluates whether or not the photon is detected, it is necessary to redefine the reflection probability and PMT quantum efficiency to account for these relative probabilities. This can be done through a straightforward application of Bayes' theorem, and gives,

$$\begin{aligned} QE'(\lambda) &= \frac{QE(\lambda)}{1 - R(1 - QE(\lambda))} \\ R'(\lambda) &= (1 - QE(\lambda))R \end{aligned} \quad (7.1)$$

where QE and QE' represent the quantum efficiency given by Hamamatsu and the quantum efficiency used in these simulations, respectively, R and R' are the photocathode reflectivity reported and used in the simulation, and λ is the photon's wavelength.

Reflections off the walls of the bell jar were treated as purely diffuse, following an angular distribution with respect to the normal described by Lambert's law, given by,

$$I(\theta) = I_0 \cos \theta \quad (7.2)$$

where $I(\theta)$ is the probability of a photon scattering to an angle θ with respect to the normal, and I_0 is the intensity at $\theta = 0$.

Reflections off of the photocathode were treated as being entirely specular, where the outgoing photon's reflection angle is equal to the angle of incidence with respect to the normal. The quartz cup containing the PMT is given a decadic extinction coefficient of 0.005 cm^{-1} , as is typical for quartz, to allow light to attenuate within it.

The scintillator's scintillation yield is assumed to be 12.7 photons/keV. This number is based on the yield measured by Borexino [164], adjusted for the 50% TMB dilution and the increased yield from the extra PPO concentration, as described in Table 6.3. Scintillation photons were produced using the PPO scintillation spectrum, as shown in Figure 5.5.

Light propagation is handled using measured values of PC and TMB absorption spectra. At any given wavelength, the absorption length is calculated by assuming equal concentrations of PC and TMB and by assuming that the molar extinction coefficients of a mixture add linearly proportional to the components' concentration [164]. Using Beer-Lambert's law, as given in Equation 5.3, we find that

$$\frac{1}{\lambda_{\text{PC}} + \lambda_{\text{TMB}}} = \frac{1}{2\lambda_{\text{PC}}} + \frac{1}{2\lambda_{\text{TMB}}} \quad (7.3)$$

Additional effects due to the absorption and re-emission of light by PPO are accounted for by using the measured absorption length of PPO to determine if a photon gets absorbed by the PPO. If it does, then we produce another photon, drawn from the emission spectrum of PPO and re-emitted isotropically, with a probability of 82%, corresponding to the quantum efficiency of PPO [34].

7.2.2 Simulation Results

For these simulations, a ^{54}Mn source was placed on the top of the bell jar and allowed to decay. The 835 MeV γ -ray was then tracked as it deposited some amount of energy into the scintillator. We kept track of the total energy deposited in the scintillator as well as the number of photoelectrons detected in each event. This ratio was used to determine the light yield of the simulation.

The primary uncertainty in this measurement comes from the systematic uncertainty in the optical properties of the scintillator, especially the scintillation yield, which has a $\sim 13\%$ uncertainty. This uncertainty comes from the uncertainty reported by Borexino [164] compounded with effects of varying the PPO concentration, discussed in Chapter 6.2.3. Additional uncertainties from the other optical properties such as the attenuation length of the scintillator and the quartz reflectivity give a total systematic uncertainty of $\sim 15\%$.

The total light yield from the simulation of the ^{54}Mn source on top of the bell jar is 0.46 PE/keV_{ee}. This number is in very close agreement with the initial light yield of 0.466 ± 0.001 PE/keV_{ee} experimentally measured, as discussed in Chapter 6.2.4.

This close agreement shows that the optical model we have built for the bell jar, which is very similar to the one we will use for simulating the LSV, can be used to perform an accurate optical simulation of the system.

7.3 DarkSide-50 Performance Prediction

Having validated the optical model on the bell jar, the next step was to apply this model to the full DarkSide-50 LSV design. A labeled drawing of the LSV as generated in the simulation is shown in Figure 7.1.

The PMTs used for the DarkSide-50 LSV are the Hamamatsu R5912 LRI built with low-radioactivity glass bulbs. A schematic diagram of the PMT is shown in Figure 7.2. These PMTs were chosen for their low radioactivity glass, low afterpulse rate, and a 1σ transit time spread of ~ 1 ns, much smaller than the scintillation pulse




Figure 7.1: A labeled drawing of the DarkSide-50 LSV as used for the GEANT4 simulations.

width. The PMT has a hemispherical photocathode with a radius of curvature of 11 cm and a diameter of 20.2 cm, resulting in a surface area of 380 cm². The bi-alkali photocathode is made of CsKSb.

The PMT was constructed in the Monte Carlo by joining together a front-facing and a back-facing paraboloid that approximate the dimensions of the PMT bulb. The photocathode was implemented as a smaller sub-volume of the front-facing paraboloid with the same cross-section as the photocathode in the physical PMTs. A cylinder was added to the back of the PMTs to represent the stem and mount of the PMTs, leading into the stainless steel wall.

The quantum efficiency of these PMTs is shown in Figure 7.3. As can be seen in this figure, the PMTs have a maximum quantum efficiency of 38.7% at 384 nm, and a full width at half maximum of 231 nm.

The PMTs in the simulation were given a quantum efficiency as described by Figure 7.3 and a 20% chance of reflecting photons off the photocathode, adjusted by Equation 7.1.

We assumed the cryostat, which is made of electropolished stainless steel, has a constant, purely specular reflectance of 58%. This number was measured by L. Cadonati on a sample of stainless steel and used as an approximation of the actual TPC cryostat stainless steel's reflectance.

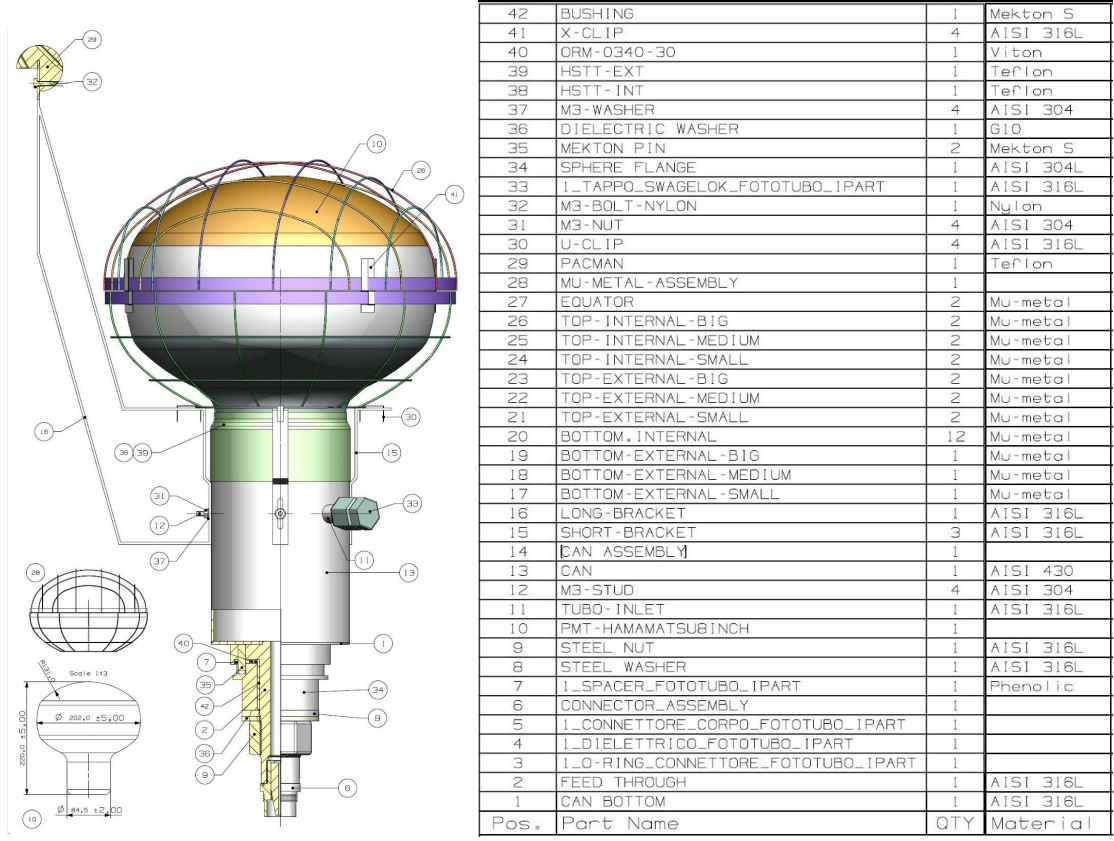


Figure 7.2: A schematic drawing of the Hamamatsu R5912 LRI PMTs used in the DarkSide-50 LSV, from [40].




Figure 7.3: Hamamatsu R5912 LRI PMT quantum efficiency curve as reported by Hamamatsu in [41].




Figure 7.4: The light yield versus position map generated by H. Qian. The x-axis is the square of the radius, the y-axis is the z-coordinate (parallel to the rotational axis of symmetry of the cryostat), and the color represents the light yield, normalized to give ^{60}Co scintillation events a light yield of 1.

Once we implemented the LSV in the Monte Carlo, we simulated 478 keV γ -rays (such as those we might see in the neutron capture reaction on ^{10}B) near the cryostat. Measuring the total amount of photoelectrons collected in the simulation per unit energy deposited, we predicted a light yield of 0.48 ± 0.09 PE/keV in the LSV.

As we will discuss in Chapter 8.4, this light yield is very close to the yield measured in the DarkSide-50 LSV.

7.3.1 Position Dependence

While we predicted a light yield for “typical” scintillation events in the LSV with the optical Monte Carlo, it is worth noting that the light yield of the actual detector is not completely homogeneous. Scintillation events that happen near the PMTs may be more likely to have their resulting photons hit a photocathode, while those happening near the cryostat may have much of their light blocked by it. A full description of the optics of the LSV therefore requires an understanding of the position dependence of the light yield.

As will be discussed in Chapter 9.2.2, S. Davini used calibration measurements at different positions to estimate the reflectance of the cryostat more precisely than we used for the above simulations. Doing so, he estimated a reflectance of $\sim 48\%$.

Using this reflectance value, H. Qian simulated 30 keV electrons scintillating uniformly distributed in the LSV, and estimated the light yield of LSV as a function of the position of the scintillation event. This position model is illustrated in Figure 7.4. As can be seen in this figure, event very close to the cryostat have smaller light yields, while those very close to the PMTs have higher light yields. However, the light yield is fairly uniform throughout the bulk of the scintillator, with little variation away from these areas.




Figure 7.5: A drawing of the LSV with a 3.8 tonne LArTPC, used for exploring how the optics of the LSV scale with a larger detector.

7.3.2 Predictions for a Larger Detector

In order to explore how this design scales with larger detectors, we implemented a design with a larger TPC. To do so, we constructed a 3.8 tonne LArTPC inside of a 2.459 m tall by 0.91 m radius stainless steel cryostat, as shown in Figure 7.5. This design resembles the design that was, at one time, being considered for the DarkSide G-2 detector.

Since the cryostat is much larger, photons from a given scintillation event have a substantially higher chance of hitting it than they did in the case with the DarkSide-50 cryostat. We therefore expect the reflectivity of the cryostat to effect the light yield more severely than it did in the earlier case. We therefore considered two cases: one where the outer walls of the cryostat were made of bare electropolished steel, and another case where the cryostat was covered in Lumirror.

In the first case, we predicted that the LSV with the larger cryostat would have a light yield of 0.47 ± 0.09 PE/keV, comparable to the value obtained in the DarkSide-50 case. However, we found that the predicted light yield went up to 0.63 ± 0.12 PE/keV when the cryostat was covered with Lumirror.

These results show that a high light yield can be obtained with this design for experiments scaled up to much larger sizes than DarkSide-50, and that covering the cryostat with Lumirror can be used to further improve the light collection efficiency.

Chapter 8

The DarkSide-50 Outer Detectors

We have now characterized the neutron backgrounds present in a dark matter experiment and, in particular, DarkSide-50. We have also discussed the general plan and research and development efforts that went into designing a veto system consisting of two nested outer detectors that can be used to efficiently eliminate and characterize these backgrounds.

The construction of these outer detectors, along with the commissioning of the TPC, was completed in 2013, thanks to much of the work done by the DarkSide-50 engineering and technical staff, as well as the rest of the collaboration.

In this chapter, we will discuss the analysis structure used for processing and analyzing data taken by the outer detectors. We will also discuss specific details of the design of the WCV and LSV in DarkSide-50 and characterize their performance. Much of the work in this chapter is also discussed in detail in [40] and [173]. Much of the work done here was done in conjunction with S. Davini, as well as the rest of the DarkSide-50 outer detector analysis team.

So far, the DarkSide-50 collaboration has published two dark matter limits. The first one, described in [50] was taken during a ~ 50 day campaign in which the TPC was filled with atmospheric argon. The second limit, described in [42, 12] showed the additional data collected during a ~ 70 day campaign taken with the TPC filled with underground argon. As discussed in Chapter 3.1.1, underground argon has a substantially lower contamination of ^{39}Ar , making it more suitable for high-exposure dark matter searches.

8.1 Electronics and Data Acquisition System

The electronics and data acquisition system for the DarkSide-50 outer detectors—the LSV and WCV—were designed by S. Davini and are discussed in [174].

8.1.1 Hardware

The PMTs are operated between 1600–1800 V and are powered by CAEN A1536 HV boards [175] housed inside of CAEN Mainframe SY4527 [176]. A $50\,\Omega$ RG213

single-cable transmission line is used to conduct both the signal and high voltage; they connect to $50\,\Omega$ terminals in the front-end where the high voltage and signal are decoupled. These cables have an electrical length of 246.8 ± 0.25 ns and an impedance of $50 \pm 2\,\Omega$.

The signal from the PMTs is read by an array of custom-built front-end analog modules, where they are decoupled from the high voltage, amplified by a factor of 10, and given a specifiable offset voltage. A copy of the $10\times$ signals is then discriminated and passed to a set of front-end digital board modules for generating a trigger signal when the detectors are operated in self-trigger mode and to monitor the PMTs' trigger rates.

The $10\times$ analog signals are then read out by National Instruments (NI) PXIe-5162 flash ADC modules. The flash ADCs are synchronized with the TPC DAQ using a NI PXIe-6674T timing and trigger module, which synchronize all devices with the TPC DAQ system. A NI PXIe-7961 FPGA module synchronizes the system's clocks with the LNGS GPS via a 1 pulse per second signal and provides a time-stamp indicator for each trigger. Storing the waveform data on the flash ADC memory buffers begins when the timing and trigger module receives a trigger signal either from the discriminators or the TPC. The waveform data are successively readout via a PXIe bus by a NI PXIe-8133 controller module, which also performs on-the-fly zero-suppression on the waveform data, and transmits the event data fragment to event builder software. The zero-suppression is discussed in more detail in Section 8.1.2. The event builder receives the event data fragments from the four NI PXIe-8133 modules being used, and finally writes the complete event structure to disk. The event structure is composed by an header which includes the trigger number and the time-stamp, followed by the zero suppressed waveform data on each channel.

The electronics are organized into 256 channels, of which 110 readout signals from the LSV, 80 readout signals from the WCV, and the remainder are unused. The digitizers had a peak-to-peak dynamic range of -0.9 to 0.1 V, and the PMTs produced signals with negative voltages.

Typical data taken during a dark matter campaign triggers on the TPC trigger signal.

8.1.2 Zero-Suppression

Between the 110 LSV PMTs and 80 WCV PMTs, the amount of storage space required to continuously record hundreds of microseconds worth of data at a 1.25 GHz sample rate would quickly exceed our record capabilities. Instead, we apply a zero-suppression algorithm to the waveforms as they are read in. This zero-suppression is meant to pick out pulses from the waveform that correspond to photoelectrons in the PMTs without recording an excess of the noise.

The zero-suppression algorithm is illustrated in Figure 8.1, and begins execution when the digitizer receives a trigger signal. The NI PXIe-8133 controller module continually reads digitized waveforms into a circular buffer. When the waveforms go below a programmable threshold, the module begins recording the waveform, starting a preset number of samples before the trigger. The module continues to record until




Figure 8.1: Illustration of the zero-suppression algorithm applied to the outer detector data. Made by L. Pagani.

the waveform goes back above the threshold. Once this happens, it records the waveform for some specified amount of time. If the waveform goes back below this threshold before this time ends, the algorithm waits for the waveform to go back above threshold and starts counting down the post-threshold samples again. Once the algorithm has reached the end of a pulse, it compares the amount of time that the pulse spent below the threshold. If this width is longer than the specified minimum width, controller writes the pulse to disk; otherwise, it discards the snippet of waveform.

From here on out, we will refer to these units of zero-suppressed waveform as *pulses*. A summary of the zero-suppression parameters used for DarkSide-50 data acquisition is given in Table 8.1. The parameters were set to allow photoelectrons to create a pulse, without letting electronic random noise do so. The typical pulse created by a single photoelectron has a width of approximately 15–20 ns.

8.1.3 Offset Calibration

Due to small variations in the electronic properties of each channel, each channel may have slightly different baseline voltages. In order for the zero-suppression threshold to have the same meaning across all channels, it is necessary to calibrate the offset voltage of each channel.

Table 8.1: Summary of the zero-suppression parameters used for DarkSide-50 data acquisition.

Parameter	Value
threshold	-50 mV (~ 0.25 PE)
pre-samples	20 samples (16 ns)
post-samples	20 samples (16 ns)
minimum width	4 samples (3.2 ns)

To measure the offset voltage of each channel, we triggered the digitizers with a pulser and measured the baseline noise in each channel. We then varied the offset voltages for each channel and determined the offset voltage needed to set the mean baseline to zero.

8.2 Outer Detector DAQ and Low-Level Analysis

Before we can discuss the analysis of the data taken with the outer detectors, it is worthwhile to discuss the low-level data reconstruction—that is, the process of reading the raw output from the digitizers and translating them into physically meaningful quantities.

The software we developed to perform this reconstruction is called DarkArtOD, and is based upon the *art* framework developed at the Fermi National Accelerator Laboratory [177].

A diagram of the reconstruction software analysis flow is shown in Figure 8.2. In accordance with the *art* framework, there are two primary entities in the software: “data products” and “modules”. In short, data products are objects that can be computed, passed between modules, and written to disk. Modules, on the other hand, take some input and can create data products.

In DarkArtOD we add the additional structure of the “submodule”. A module can call some number of submodules which carry out various tasks that the module may need to perform in order to produce a final data product. As is illustrated in Figure 8.2, we can conceptually divide the reconstruction into four levels: raw data-level, pulse-level, sum waveform-level, charge estimator-level, and output-level.

While we will discuss many of these modules below, it is worthwhile to first offer an overview of the entire reconstruction process. Reconstruction begins with the raw data—the binary file written to disk by the digitizers, consisting of encoded zero-suppressed waveforms, run-level data such as DAQ settings and the run number, and event-level data such as the output of the timing modules and the event numbers. At the raw data level, we decode this binary information and record the waveform, run, and event information for future use.

At the pulse-level, we apply a number of corrections to the pulses to account for various effects that might affect the shape and integral of these pulses.

At the sum waveform-level we stitch the pulses seen by each channel together into an event-level picture of photoelectron detection times. The sum waveform generator

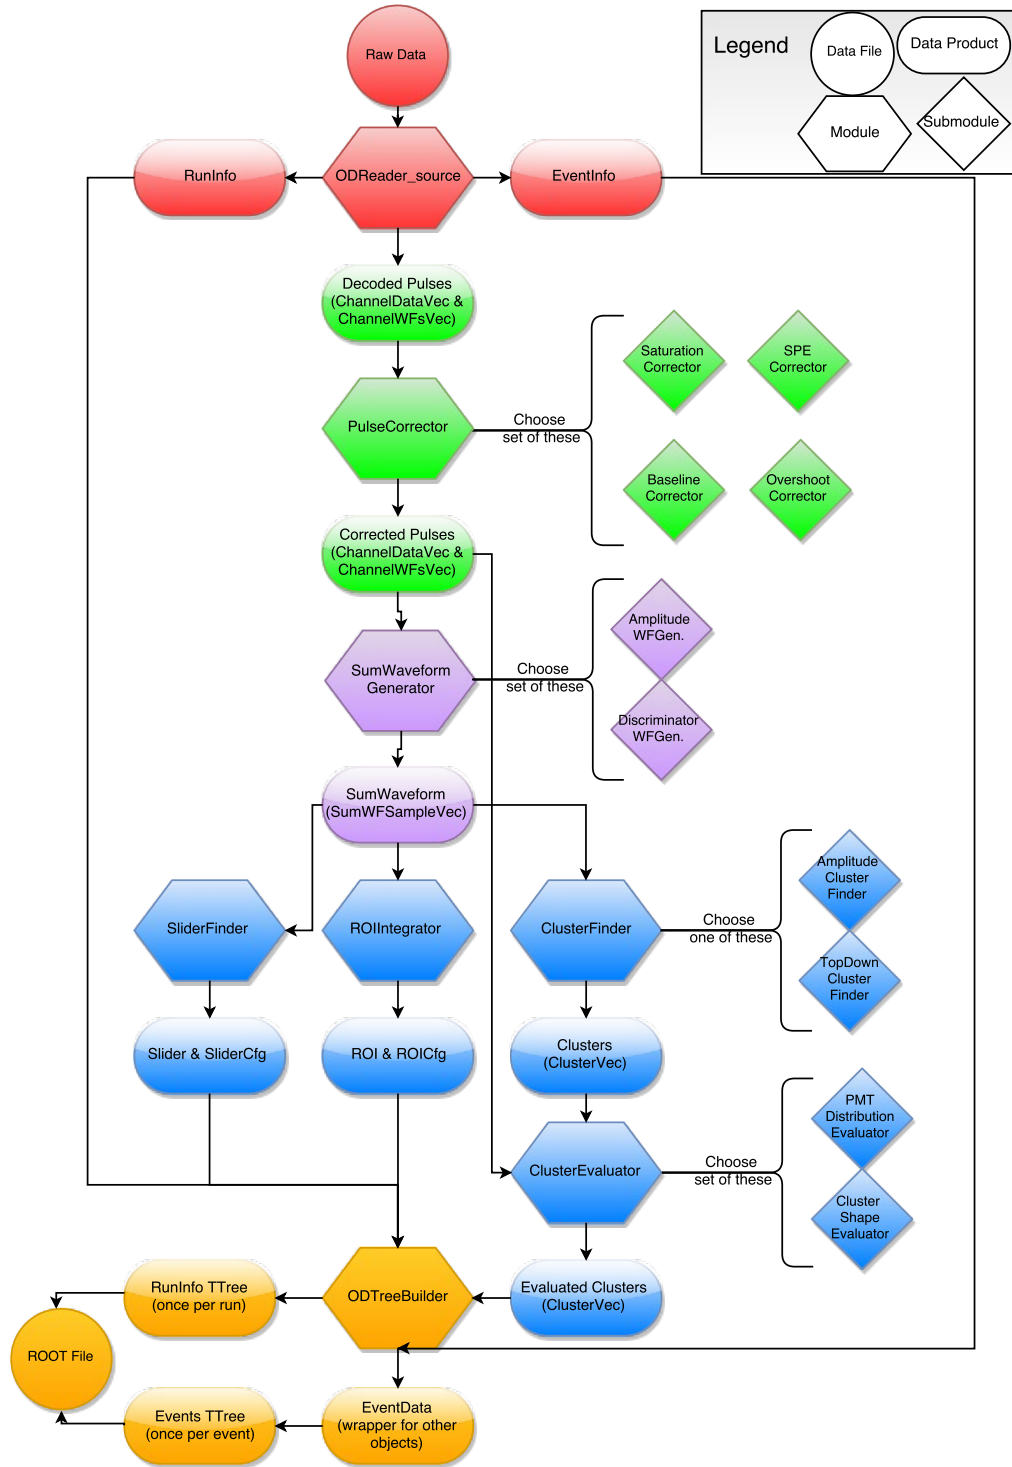


Figure 8.2: A diagram showing the analysis flow of the DarkArtOD data reconstruction software. Red nodes are the raw data-level, green nodes are pulse-level, purple nodes are sum waveform-level, blue nodes are charge estimator-level, and yellow nodes are output-level.

module creates an “amplitude waveform” which is the sum of all pulses, offset by their start times, as well as a “discriminator waveform” that emulates a discriminator acting on the data. The amplitude of the amplitude waveform at any time is the total current read out by all of the PMTs (after zero-suppression) at that time; the amplitude of the discriminator waveform is approximately the number channels that have triggered within some specifiable amount of time prior to that time.

These sum waveforms are then passed down to the charge estimator-level. At this level in the event reconstruction, we look at various integrals of the sum waveforms and try to interpret physically relevant quantities from the waveform shape. This consists of a set of sliding windows with a fixed width that scan a sub-region of the sum waveform, looking for the window of greatest size, integrals of fixed temporal regions of interest, and clusters. Clusters are temporally similar photoelectrons that may correspond to a scintillation event of some energy.

Lastly, these data structures are passed to the output-level, where they are written to disk in a ROOT file. These ROOT files are used for higher level analysis, and are further trimmed down to smaller sizes related to processed TPC data files through the SLAD software written by H. Qian. These SLAD files contain a tree of outer detector events where each event number matches the appropriate TPC event number, and all time coordinates are converted to the time relative to the S1 start time seen in the TPC. These SLAD files are what we ultimately analyze when we use the outer detectors to veto events in the TPC.

8.2.1 Pulse Reconstruction

As mentioned above, after the raw data has been decoded, the first step in reconstructing an event is to reconstruct the individual pulses. The raw pulse data is read into DarkArtOD as a vector of ADC counts with a sample number corresponding to where, relative to the start of the acquisition window, the pulse begins. The decoder converts these ADC counts into a voltage measurement based on the digitizer settings read in at the start of the run.

Since the amplitude of the waveforms is converted to volt and the measure time, integrals over the waveform have units proportional to charge (accounting for resistance in the proportionality constant). It is therefore common to refer to waveform integrals as “charge”.

Saturation Correction

Because the digitizers have a finite range, any pulse whose amplitude exceeds this range will be saturated. The typical photoelectron has an amplitude of approximately -200 mV. However, if multiple photoelectrons are detected on a single channel at approximately the same time, their amplitudes may pileup to exceed the -0.9 V dynamic range of the digitizers and the pulse will become saturated.

Typically, a non-saturated pulse looks like the one shown in Figure 8.1. When a pulse is saturated, the peak has exceeded the maximum value that can be recorded by the digitizer, and so we see the bottom of the pulse flat, instead of pointed. This is

identified by three or more consecutive samples within the lower 90% of the amplitude range.

If a pulse is saturated, the saturation correction submodule attempts to correct this saturation. The algorithm does so by calculating a line connecting the two samples immediately before the saturated region, and another line connecting the two samples immediately after the saturated region. These two lines form a triangle above the saturated samples; the area of this line is added to the area of the pulse as an approximate correction to the pulse charge.

The accuracy of this correction was estimated by L. Pagani and S. Davini by artificially saturating pulse waveforms between 4–7 PE and comparing corrected waveforms to their actual waveforms. Doing so, they found that the corrected pulse integrals were correlated with the true integrals by a linear trend passing through the origin with a correlation coefficient $> 99\%$.

Overshoot Correction

Since photoelectron signals in the PMTs create pulses of negative voltage, any signals decoded to have positive voltage must be non-physical. This can happen as a result of noise; however, positive voltage noise tends to be relatively small in amplitude compared to a photoelectron. However, large signals in the PMTs create a high current; after this current dies away and the voltage returns to the baseline, the electronics may overshoot and create a significant positive voltage before returning to the baseline.

To avoid factoring the overshoot voltage into the pulse integral, we artificially set all samples with a voltage greater than 15 mV to 15 mV.

The effects of this correction are most evident in muon signals, which saturated all of the channels for several microseconds, and are followed by very large overshoot noise. By removing this noise, we increase the integral of muon signals and make the total charge collected in these events more consistent.

Baseline Correction

Although the offset calibration brings the mean baseline on each channel close to 0, low and medium frequency oscillations in the baseline mean that the mean will not necessarily be close to 0 for each pulse. This may create a pedestal background that individual pulses can sit on and may degrade the resolution of the detector by adding extra noise.

We therefore apply a baseline correction to each pulse to remove this source of noise. To do so, we integrate over the first 15 pre-threshold samples (12 ns) of each pulse to calculate the mean baseline. We then uniformly subtract this baseline from the entire pulse. Since the typical pulse length is around 15–20 ns wide, this baseline subtraction remove baseline fluctuations with frequencies small compared to the pulse width.

SPE Correction

Each PMT and channel will naturally have slightly different gains, and these gains may change over time. Since we sum over signals from many different channels, it is important to normalize each channel by their respective gains.

To normalize each channel by the gain, we measure the charge mean charge of a single photoelectron. By normalizing pulses by the SPE mean, we can measure charge in the number of photoelectrons detected by each PMT. This normalization is very similar to that discussed in Chapters 2.4 and 6.2.

In order to measure the SPE charge distribution in the LSV PMTs, dedicated laser runs are performed. As shown in Figure 7.2, each PMT is equipped with an arm extending from the base of the PMT to a position in front and to the side of the photocathode. An optical fiber is mounted on this arm, pointing into the photocathode. The optical fibers run through a system of optical feedthroughs to an external 405 nm laser with picosecond timing. This setup allows us to deliver a low intensity laser pulse directly to each PMT.

The laser is then pulsed at a low intensity, sending photons through the optical fibers to many of the PMTs. The low intensity of the laser ensures that the PMTs will primarily see 0 or 1 PE with each pulse. The average number of photoelectrons detected per trigger is referred to as the occupancy; the laser intensity is set so that the typical PMT will have an occupancy between 1–2%.

This laser is pulsed at 500 Hz, and trigger signal is simultaneously sent to the outer detector digitizer boards. During these dedicated laser runs, the outer detector digitizers open with a $2\ \mu\text{s}$ window when they are triggered. Pulses are then selected from a 150 ns wide window around the trigger time to reduce background from random coincidence scintillation events.

We then subtract the baseline from these pulses and integrate them to determine the SPE charge distribution. We then fit a Gaussian distribution plus an exponential around the peak of the SPE distribution to determine the SPE mean and relative variance for each channel. SPE means are typically $\sim -10^{-9}\text{V}\cdot\text{s}$, and the relative variance is typically $\sim 10\%$.

Since SPE means are negative, normalizing the waveforms by the SPE means inverts the waveforms. Therefore, while physical pulses typically produce negative voltages that dip below the baseline, SPE-normalized waveforms show pulses as positive excursions above the baseline.

8.2.2 Charge Integral Estimators

In order to be able to interpret physical events from these pulses, we must understand how the photoelectrons are arranged in time. Then, by identifying groups of photoelectrons, we can identify the time and energy of scintillation pulses that may have eventually produced these photoelectrons.

For the DarkSide-50 analysis, we use three different types estimators to study these temporally grouped photoelectrons.

Clusters

The first estimator we use is known as *clusters*. These are photoelectrons that are clustered together in time, and can be used to identify the time and energy of a scintillation event.

An acquisition window opened with the outer detector DAQ receives a trigger signal may contain any number of clusters, which are identified by scanning the total waveform.

Two clustering algorithms were explored. The first algorithm is referred to as the “Amplitude Cluster Finder”. This algorithm scans the amplitude sum waveform, starting at the first sample, looking for samples where the spe-normalized current exceeds a specified threshold (typically set to a value comparable to twice the amplitude of a single photoelectron). This threshold crossing triggers the start of a cluster. Clusters found by this algorithm were given a specified width (typically 300 ns); after triggering, the algorithm steps back a specified amount of time (typically set to 20 ns) and integrates for the specified duration. Once one cluster ends, the algorithm then immediately begins looking for the next cluster by searching for the next threshold crossing.

The second algorithm is referred to as the “Top-Down Cluster Finder”. This algorithm is illustrated in Figure 8.3 and works by identifying the maximum value of the SPE-normalized amplitude waveform. The algorithm checks if this maximum is above some specified threshold (typically equivalent to two photoelectrons); if it is not, the algorithm terminates. If the amplitude is above the specified threshold, it begins looking for the start of the cluster. Regions of the sum waveform where a pulse is not present are exactly 0 in amplitude due to the zero-suppression. This algorithm therefore identifies the beginning of a cluster by looking for gaps between pulses. Once it finds a gap larger than a specified value (typically set to 20 ns), it begins looking for the end of the cluster by finding another gap of the same length. A minimum cluster width (typically set to 50 ns) can also be specified; if so, the algorithm will not identify the end of a cluster until at least the specified amount of time has elapsed, or the end of the sum waveform has been reached.

Once the start and end of a cluster have been identified, the algorithm records the integral of the cluster, removes the samples from the sum waveform, and reiterates until there are no more samples in the waveform above the threshold.

Most of the analysis involving clusters presented in this document was done using the Top-Down Cluster Finder algorithm. Since this algorithm has a variable width, we believe that it can more accurately determine the energy of scintillation events; however, it is also slower than Amplitude Search Cluster Finder algorithm.

These clusters are useful for studying the scintillation spectrum in the LSV, as it can identify the time and energy of scintillation events in and out of prompt coincidence with TPC signals. It can also be used to identify individual Cherenkov signals in the WCV.




Figure 8.3: An example of a waveform in the LSV with its corresponding cluster, illustrating the Top-Down Cluster Finder algorithm. The variable in the horizontal axis is the time in ns with respect to the trigger.

Temporal Regions of Interest

In order to avoid the risk of missing a neutron event in coincidence with the TPC due to subtleties in the clustering algorithm (for example, a cluster from a high energy random background shortly after the coincidence signal may pick up the tail of the coincidence signal, lowering the reconstructed charge of the coincident cluster), a series of *temporal regions of interest* are defined relative to the TPC S1 pulse start time. In each region of interest, the LSV sum waveform is integrated between a specified start and end time each time the LSV is triggered. By doing so, the total amount of scintillation light in a specific time window can be calculated reliably, without risking being lost due to subtleties of the algorithm.

These regions of interest are primarily useful during the dark matter search for identifying the prompt coincidence signal. For most of the WIMP-search data taken, a window extending from 50 ns before the TPC saw the event to 250 ns afterwards is integrated over. Since we expect neutrons to thermalize in well under 100 ns, we expect this window to be able prompt coincidences very effectively.

Additional regions of interest are defined several microseconds before and after the prompt region in order to provide estimates of random backgrounds that might be present in the prompt window as well.

Sliding Windows

Lastly, we define a set of *sliding windows*. The sliding window algorithm takes a start time, an end time, and a window width. The algorithm then scans the amplitude sum waveform between the start and end times, searching for the sub-window of the specified width that has the greatest charge. The analysis code allows multiple sliding windows to be defined.

This charge estimator provides a reliable tool for identifying delayed coincidences between the TPC and LSV. To do so, the algorithm searches for a 500 ns wide window between the time corresponding to the S1 start time and the end of the LSV acquisition window, searching for the region with the greatest charge. If a neutron capture happens in the LSV, this algorithm is therefore likely to find its capture signal some time after the event in the TPC.

Another sliding window of the same width is defined to scan the region before the TPC signal to search for any signals that may precede a signal in the TPC, such as might be seen from a high energy external neutron entering the LSV from the WCV.

8.3 The Water Cherenkov Veto

A photograph taken from inside the DarkSide-50 WCV can be seen in Figure 8.4. This vessel was eventually filled with water, providing passive shielding and a muon veto signal to the inner detectors, as discussed in Chapter 5.1.




Figure 8.4: The inside of the DarkSide-50 WCV. The inner surfaces are all covered in Tyvek, as is the LSV, seen in the middle, standing on four stilts. Taken by Y. Suvarov.

8.3.1 Detector Design

The stainless steel water tank that now constitutes the WCV was originally built for the Borexino Counting Test Facility (CTF). As can be seen in Figure 8.4, the WCV is a stainless steel cylinder with a diameter of 11 m and a height of 10 m. When filled, it contains ~ 1000 tonnes of water. It is built from carbon steel and is protected from corrosion by a coating of Permatex resin.

The 80 8" PMTs along the sides and floor are each equipped with light-collecting collars that improve their light collection efficiency.

A port on the top of the vessel has several tubes feeding down, into the LSV. Some of these tubes carry fluid or cabling to and from the LSV and its PMTs, while others continue on into the TPC, where they are used for powering the TPC electronics and filling and draining the LAr. The top of this port is viewed by a large flange from the DarkSide-50 radon-free clean room.

The water that fills the WCV is distilled using the Borexino water purification plant, described in [8]. This process purifies the water to the levels shown in Table 8.2.

The process of filling the WCV began on October 2, 2013 and ended November 15, 2013.

8.3.2 Muon Signal

Since the WCV only produces a physical signal when high energy charged particles produce Cherenkov light in it, the background rate in the WCV is substantially lower

Table 8.2: Water contamination levels before and after purification with the Borexino water purification plant, as reported by [8]

Isotope	LNGS Water [Bq/kg]	Purified Water [Bq/kg]
^{238}U	1×10^{-3}	3×10^{-7}
^{226}Ra	0.3	$< 1 \times 10^{-6}$
^{222}Rn	10	$< 3.4 \times 10^{-6}$
^{232}Th	1×10^{-3}	3×10^{-7}
^{40}K	1×10^{-3}	$< 1.6 \times 10^{-6}$




Figure 8.5: Total charge collected in the WCV for each event in the ~ 70 day underground argon campaign [42]. (Blue) All events, including those with cosmogenic signals and those without. (Red) Events with coincidental high energy signals in the LSV, predominantly due to cosmogenic signals in both detectors.

than that in the LSV. The primary source of background in the WCV is the dark noise of the PMTs.

During the standard DarkSide-50 dark matter search analysis, the primary quantity that we measure in the WCV is the total number of photoelectrons detected during the acquisition window of that event.

Figure 8.5 shows the distribution of photoelectrons collected per event over a ~ 70 day period of running. The outer detectors triggered on the LSV signal during these runs, which means a muon signal was only recorded when there was an accompanying event—either directly from the muon itself, or from a cosmogenic neutron—in the TPC. At low energies, we can see the contribution of dark noise in the detector, while at signals $\gtrsim 1000$ PE, we see another hump corresponding to cosmogenic signals in the WCV. Since these signals are likely to produce a similarly large signal in the




Figure 8.6: An example waveform of a cosmogenic signal in the WCV. The inset plot is zoomed in around the time that the muon passed through the detector.

LSV, we can verify that these high energy events are likely cosmogenic in origin by finding that events in this hump are very likely to be in coincidence with large signals in the LSV.

By collecting data in a run where the outer detectors triggered on the LSV, we were able to measure the rate of muons passing through both the LSV and WCV. Doing so, we saw $380 \pm 5 \mu/\text{day}$, consistent with the rate of $370 \mu/\text{day}$ reported by Borexino [178] after scaling for the detector size.

An example WCV waveform during a cosmogenic event is shown in Figure 8.6. As can be seen in this figure, passing muons produce a very large signal in the WCV that is easily detected. Around $6.6 \mu\text{s}$ before the trigger signal was received, we see a very large spike in the WCV PMTs as a large burst of Cherenkov radiation produces ~ 1000 PE. A long tail of afterpulses and electronics noise caused by this large signal then follows for well over $100 \mu\text{s}$ following the muon.

For the DarkSide-50 underground argon campaign, we vetoed all events in which a signal over 200 PE was seen in the WCV, to be sure that we can remove all cosmogenically-induced events. Data taken in which we triggered the outer detectors with a pulser to collect data during random triggers indicate that $\sim 0.36 \pm 0.03\%$ of all non-cosmogenic events will be removed by this cut, primarily due to the dark noise in the PMTs.




Figure 8.7: Two photographs of the inside of the DarkSide-50 LSV, taken by Y. Suvarov. (*Left*) The TPC cryostat suspended inside the LSV, held in place by connective pipes. (*Right*) The Lumirror reflector and PMTs mounted on the walls of the detector.

8.4 The Liquid Scintillator Veto

The LSV can be seen inside the WCV in Figure 8.4. Figure 8.7 shows two photographs of the inside of the LSV with the TPC cryostat lowered inside of it. This vessel was filled with boron-loaded liquid scintillator, as discussed in Chapter 5.2.

8.4.1 Detector Design

The stainless steel sphere comprising the LSV, as shown in Figure 8.7, has an inner diameter of 4 m. When it filled, it contains ~ 30 tonnes of liquid scintillator. In this figure, we can see the reflective Lumirror behind the 110 R5912 LRI PMTs.

As seen in the right-hand photograph in Figure 8.7 as well as the diagram in Figure 7.2, a 1.0 mm diameter μ -metal cage surrounds the PMTs. The large size of the PMTs means that they are sensitive to Earth’s magnetic field. This highly conductive μ -metal cage prevents the Earth’s magnetic field from interfering with the performance of the phototubes. Since μ -metal may accelerate oxidation of PC, we coat this cage in a $\sim 20\ \mu\text{m}$ thick lining of phenolic paint from Morton [179].

PMT Potting

The PMTs are mounted in the stainless steel sphere of the LSV. They each connect to feedthroughs in the walls, which connect them to cables carrying both the high voltage power and the signals. These cables run through the water of the WCV and eventually into the control room where they interface with the data acquisition system. The back-end sealing of the PMTs and feedthroughs had to be compatible with PC, TMB, and water. To develop such a seal, we encapsulate the base and neck

of the PMTs in cylindrical stainless steel housing with an outer diameter of 90 mm. The housing is affixed to the glass in the PMT neck using EP45HT epoxy resin from Master Bond, which provides structural support as well as a seal to keep the inside of the housing dry. The end-cap of the housing was welded and helium leak tested to ensure the quality of the seal.

These housing units are inserted into holes in the sphere which serve as feedthroughs, secured with a rear nut. These interfaces are secured with Viton o-rings and a layer of phenolic resin, which additionally electrically decouples the PMTs from the sphere to avoid ground loops. The feedthroughs were sealed from the back end with underwater jam-nut connectors produced by Framatome [180]. These connectors are potted using the same EP45HT epoxy resin to prevent water infiltration, and the inside of the cylindrical housing is filled with inert mineral oil which prevents water from condensing on the divider without stressing the delicate joints of the metal pins and the glass. Finally, a heat-shrink Teflon tube is glued to the PMT neck and steel can by the epoxy resin. To achieve this connection, we used the Tetra-Etch technology by the Gore company [181], which employs a preliminary surface etching on the Teflon film to allow a stronger adhesion with other surfaces.

Through these mechanisms, we were able to mount the PMTs inside of the LSV while protecting the electronic components from the scintillator and water on both ends while securely keeping the scintillator and water separate, with multiple layers of redundancy offering additional security.

Fluid Handling System

The liquid scintillator fluid handling system was largely based off of the design used for Borexino [182, 183], including using many of the fluid storage facilities from Borexino.

A block diagram of this system is shown in Figure 8.8. Pre-distillation fluids were transferred from their storage vessels into the distillation column, where they were distilled under a vacuum to an intermediary vessel. The re-condensed fluids were then run through a packed nitrogen-stripping column to further remove impurities before being transferred to the DarkSide-50 filling station.

The scintillator cocktail was prepared in two stages. The first stage involved preparing a *master solution*. The master solution was prepared by distilling the entire mass of PPO and dissolving it in distilled PC.

The second stage involved separately distilling PC and TMB. The distilled PC, TMB, and master solution were then mixed together in-line to ensure a homogeneous mixture when the fluids finally filled the LSV.

8.4.2 Run Phases

We can divide the data taken with the LSV into two experimental phases. The two phases correspond to the TPC being filled with atmospheric and underground argon, respectively, and each have different scintillator cocktails. The differences and run times of both phases are summarized in Table 8.3. Phase-I corresponds to the data published [50], while Phase-II was used for the data published in [12].

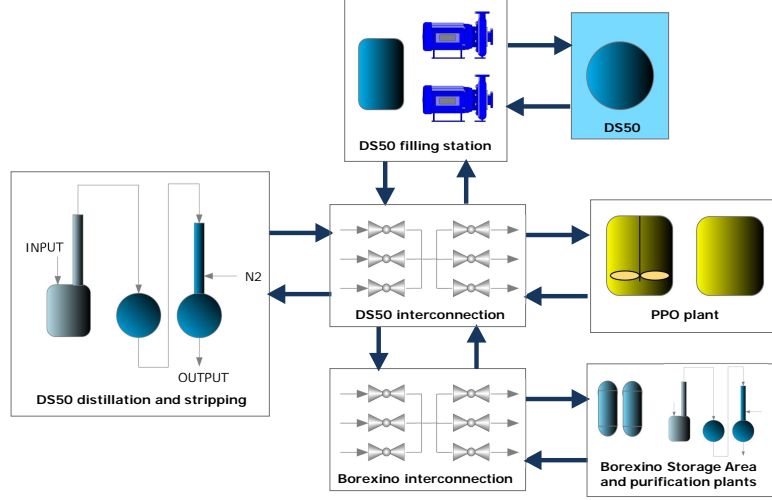


Figure 8.8: Block diagram of the fluid handling systems for the DarkSide-50 LSV, from [40].

Table 8.3: Summary of Phase-I and Phase-II of LSV operation, including start and end dates of both phases, the type of argon used in the TPC for each phase, the scintillator composition, the observed ^{14}C decay rate, the thermal neutron capture time, the calculated ratio of thermal neutrons capturing on ^1H to those capturing on ^{10}B , and the acquisition window used relative to when the outer detector DAQ received the trigger signal.

	Phase-I	Phase-II
Start	Nov. 2013	Feb. 2015
End	Jun. 2014	<i>present</i>
TPC Fill	Atmospheric Ar	Underground Ar
LSV Cocktail: PC	50%	95%
LSV Cocktail: TMB	50%	5%
LSV Cocktail: PPO	2.5 g/L	1.4 g/L
^{14}C Rate	$\sim 150 \text{ kBq}$	$0.25 \pm 0.03 \text{ kBq}$
Neutron Capture Time	$2.2 \mu\text{s}$	$22 \mu\text{s}$
$\text{BR}(\text{n} + ^1\text{H} \rightarrow ^2\text{H} + \gamma)$	0.8%	7.7%
O.D. DAQ Window	$[-4.5, 65.5] \mu\text{s}$	$[-7, 133] \mu\text{s}$ $[-10.5, 190] \mu\text{s}$

As we will discuss later, the primary reason for changing the scintillator composition was an overwhelming high rate of ^{14}C contamination observed in the TMB. After finding replacement TMB with much lower levels of contamination, we reconstituted the scintillator using the cocktail in Phase-II, using lower concentrations of TMB and PPO due to price and time constraints.

Furthermore, as can be seen in Table 8.3, we recorded data for different lengths of time in the outer detector during both phases. During Phase-I, the $2.2\,\mu\text{s}$ thermal neutron capture time means that nearly all neutrons will capture within $\sim 20\,\mu\text{s}$ after they thermalize. However, neutrons that thermalize in the stainless steel of the cryostat or the Teflon of the TPC reflector will take a longer time to capture. It has been shown in [157] that waiting $60\,\mu\text{s}$ after the thermalization signal allows us to veto on the γ -rays produced by neutron captures on ^{19}F , ^{56}Fe , and ^{52}Cr .

During Phase-II, the neutron capture signal in the scintillator was ten times longer, meaning that $65\,\mu\text{s}$ was no longer enough time to reliably detect neutron captures. As a result, we initially extended the window to $140\,\mu\text{s}$ after the trigger signal was received, allowing ~ 6.4 times the capture signal to elapse. After about half of the data published in 70 day underground argon campaign [12] was taken, we decided that the background rate in the LSV was low enough to further extend the acquisition window to $190\,\mu\text{s}$ after the trigger time. This longer time window now covers ~ 8.6 times the thermal neutron capture time, increasing the likelihood that we a neutron captures during the acquisition window.

8.4.3 Carbon-14 Background

As described in Chapter 3.1.6, ^{14}C is a naturally present background in any organic scintillator. The low energy spectra seen in the LSV during Phase-I and Phase-II can be seen in Figure 8.9, where the contribution of ^{14}C , with a $156\,\text{keV}$ endpoint β -spectrum, is clearly visible in Phase-I. Much of the work studying the ^{14}C spectrum was done in collaboration with L. Crippa.

While the PC used for the scintillator is the same measured by Borexino to have a very low ^{14}C contamination, the TMB came from a Dow Chemicals. Dow Chemicals has a large chemical plant in the United States, and another in the Netherlands. The prototype tests were all performed using TMB produced by the company’s US plant, and a high rate of ^{14}C was not observed. However, shortly after filling the LSV, the ^{14}C spectrum shown in Figure 8.9 was observed.

Upon further investigation, we learned that Dow Chemical’s plant in the US produces TMB from petroleum, while the plant in the Netherlands derives the TMB from biogenic methanol. This difference means that the TMB we used for our prototype tests had a very low ^{14}C contamination, much like the PC which is similarly derived from petroleum, while the TMB we used to fill the LSV had nearly atmospheric levels of ^{14}C .

During Phase-I, we observed a ^{14}C background rate of $\sim 150\,\text{kBq}$. This rate is about a tenth of what we would expect from a sample of purely atmospheric carbon, indicating that the TMB produced at Dow Chemicals’ plant in the Netherlands might actually use some combination of petroleum-derived and biogenic substrates. An




Figure 8.9: (*Black*) The low end of the photoelectron spectrum seen in the LSV in Phase-I, with ^{14}C -rich TMB, at 50% concentration, compared to (*red*) the Phase-II low- ^{14}C TMB at $\sim 5\%$ concentration. The dramatic decrease in ^{14}C background is evident.

alternative explanation is that the plant-matter used to derive the TMB may have been collected before the 1960's, when nuclear weapons testing caused a large spike in the atmospheric ^{14}C concentration.

Since the endpoint of the β -spectrum is 156 keV, this high rate covers much of the energy region where we expect to see neutron thermalization signals, as well as the neutron capture on ^{10}B channel that goes directly to the ground state of ^7Li . Trying to veto on these signals would therefore result in an intractably high dead time for the experiment, as nearly every event would have an accidental veto signal from a randomly coincident ^{14}C decay. Therefore, for the dark matter search performed during this phase, we had to use laxer cuts than we would be able to sustain for a longer campaign, while suffering an acceptance loss $\sim 11\%$.

TMB Replacement

In order for DarkSide-50 to remain background free for the duration of the experiment, the LSV needed to be able to make more stringent cuts than those used during Phase-I. In order for these cuts to be feasible, it was necessary to decrease the ^{14}C background in the experiment. It was therefore necessary to find a source of TMB that would have as low a ^{14}C rate as possible.


Much of the work done in locating lower- ^{14}C TMB was done by B. Buchholz at Lawrence Livermore National Laboratory. We identified five TMB samples of interest to us. In order to determine the ^{14}C contamination in these samples, we sent them to the Accelerator Mass Spectrometry facilities at Lawrence Livermore National Laboratory.

The five samples tested included two produced by the Dow Chemical plant in the USA, two by the Dow Chemical plant in the Netherlands, and one produced by Kemira in Finland. While the TMB produced by Dow Chemical in the USA and Kemira are both petroleum-based, the TMB from Kemira would likely be cheaper to purchase than the TMB from Dow Chemical due to its proximity to Italy.

In addition to these five samples, measurements were taken of diesel and heptamethylnonate, two purely petroleum-derived samples, for comparison, as well as the ox-1 and ox-2 standard NIST samples of oxalic acid measured in 1950 and 1977, respectively, the Australian National University standard sample of sucrose, and the TIRI-Wood AMS standard sample.

Sample Preparation Samples were prepared by extracting ~ 5 mg into a vial with ~ 100 mg of CuO stones. These vials were then frozen with liquid nitrogen, pumped down to a vacuum, and sealed shut with a torch before being left to combust overnight. This process converts the sample into carbon dioxide. The vial is then transferred into a sleeve and broken; the escaping CO_2 then travels through a tube through several cold traps. Each cold trap consists of dry ice and ethanol, and condenses out most of the impurities that may escape the vial with the CO_2 .

The gaseous CO_2 is then condensed in a liquid nitrogen cold trap, which separates out impurities with lower boiling points than CO_2 such as helium. When the CO_2 condenses, it enters a T-junction consisting of a pressure transducer that measures the ingress of CO_2 , a terminal containing iron powder, and a terminal containing $\text{Mg}(\text{ClO}_4)_2$ (magnesium perchlorate). Hydrogen is then injected into the T-junction and reacts with the CO_2 through the reaction



The water formed in this reaction is absorbed by the magnesium perchlorate, while the elemental carbon forms a graphite layer on the iron powder. The graphite samples were then hammered into a small depression in an aluminum cylinder and loaded onto a wheel where they could be inserted into the accelerator.

This process produced 150–350 μg of graphite for each of the TMB samples and 970 μg of graphite for the diesel samples.

Measurements To measure each sample, the wheel rotated the chosen sample into a Cs^- beamline. Data taken from the first two minutes of Cs^- bombardment were thrown out to exclude backgrounds from surface impurities. Some carbon atoms hit by these ions became negative ions and were accelerated across a series of plates of increasing voltage. The velocity to which the ions were accelerated was chosen so that the ions would turn 90° in a magnetic field at the end of the track. This bend allowed ions of specified masses to be selected from the beamline; different parameters were used to choose ^{14}C and ^{13}C separately.

Selected ions were then accelerated down a 10 MV tandem accelerator, at the end of which they passed through a carbon foil, which stripped them of their electrons, turning them into positive ions. The positive carbon ions were then accelerated down

Table 8.4: Accelerator mass spectrometry measurements comparing the $^{14}\text{C}/^{13}\text{C}$ fraction for each measured sample to the modern fraction.

Sample	^{14}C fraction/modern ^{14}C fraction
Dow Chemical USA 1	0.0019 ± 0.0001
Dow Chemical USA 2	0.0017 ± 0.0001
Dow Chemical Netherlands 1	0.1458 ± 0.0009
Dow Chemical Netherlands 2	0.1966 ± 0.0007
Kemira Finland	0.0047 ± 0.0001
Heptamethylnonane	0.0020 ± 0.0001
Diesel	0.0009 ± 0.0001

an additional 10 MV potential through a series of plates decreasing in voltage back down to ground.

A second magnet was then more finely tuned to send ^{14}C and ^{13}C ions down separate tracks. The current of the ^{13}C track was measured to determine the number of ^{13}C nuclei in the sample, while the ^{14}C track passed through another 90° bend with a magnetic field to further reduce impurities. The ^{14}C ions then passed through a final chamber with perpendicular electric and magnetic field, designed to select ions by their velocities. Finally, the ^{14}C ions passed through a mylar filter and entered an argon gas proportional counter, which differentiated between ^{14}C ions and $^7\text{Li}_2$ contaminant ions by their track length.

These measurements gave the $^{14}\text{C}/^{13}\text{C}$ fraction for each sample, which were then normalized to the ^{14}C fraction for modern atmospheric carbon samples.

Results The results from the measurements are summarized in Table 8.4. The ox-1, ox-2, ANU, and TIRI-Wood standard samples were found to agree with their expected values, validating these measurements. Each sample was measured five times and the results were averaged together. The diesel background measurements were subtracted from the Dow Chemical Netherlands measurements, but not from any of the other ones reported here.

As can be seen in Table 8.4, the Kemira and Dow Chemical USA samples all showed very low ^{14}C fractions. Comparing these measurements to the measurements of heptamethylnonane, which should be approximately a measurement of the background, show that the Dow Chemical USA samples were ^{14}C -pure below the background of the experiment, while the Kemira Finland samples was only a little bit above the background.

The dominant sources of background in these measurements comes from $^7\text{Li}_2$ and ^{13}CH , which may contaminate the graphite sample placed in the Cs^- beam. Other surface contaminants not removed in the first two minutes of the measurement may also provide a background.

As a result of these measurements, we decided to use the TMB produced by the Dow Chemical plant in the USA, despite the additional shipping cost.

Scintillator Replacement

The first step in replacing the TMB in the scintillator was to remove all of the original TMB. This removal was accomplished using the Borexino distillation plant because the DarkSide plant only has a single stage and therefore cannot be used to separate PC and TMB. Five repeated distillations were used to reduce the TMB concentration by a factor of ~ 1000 , treating the TMB as the light product of the distillation and sending it to one of the storage vessels for removal.

Meanwhile, the PC and PPO were removed from the bottoms of the distillation product vessels, and sent back to be re-purified. After refilling the LSV with the PC and PPO solution, we found that the light yield had dropped by a factor of ~ 2 . Further investigation revealed that heavy contaminants in the bottom of the storage vessel had contaminated the scintillator, making an addition re-purification cycle necessary.

This second re-purification cycle was performed using the Borexino distillation plant followed by nitrogen stripping. While this cycle removed the contaminants and restored the light yield of the LSV, it also removed much of the PPO in the scintillator, which we eventually had to add back in, which we did via a master solution. However, due to price constraints, we only added enough PPO to restore the concentration to 1.4 g/L.

Similarly, once we had received the new batch of low- ^{14}C TMB, we distilled it and mixed it in-line with PC and PPO, though only restored the concentration to 5% rather than 50% due to price and time constraints.

8.4.4 Measured Energy Spectra

The energy spectrum of clusters observed in the LSV during Phase-II is shown in Figure 8.10. The most straightforward way to look at the spectra of energy deposited in the LSV is to look at clusters identified in different time windows. The versatility in the cluster variables in being able to identify several features per event at different, non-pre-specified times, with variable energies makes them powerful tools for studying the activity seen by the LSV.

The first distinction we can make in different types of clusters that we see in the detector is based on the time at which they occur. Figure 8.11 shows the time profile of cluster start times, broken down into that of high and low energy clusters. As we can see in this figure, there is a large spike around $-6\ \mu\text{s}$. This spike corresponds to prompt coincidences between the LSV and TPC. Since most of these coincidences are due to γ -rays scattering in both detectors, they tend to produce relatively high energy signals. The fact that the rate of these events is so low outside of this prompt spike can be taken as a testament to the cleanliness of the detectors.

However, we see three broader peaks following the prompt spike in the low energy distribution. These peaks correspond to afterpulses in the LSV PMTs. Before installing them, P. Lombardi tested many of the PMTs and found that they typically have afterpulse rates $\sim 10\%/PE$ and can happen within $20\ \mu\text{s}$ following the original signal. High energy scintillation events (such as those we typically see when we have a




Figure 8.10: Photoelectron spectrum of the LSV clusters during Phase-II: (*black*) for all clusters and (*red*) for clusters starting within a ± 25 ns prompt coincidence with the TPC. Region A is the low energy signal that consists largely of afterpulses, Region B is the ^{60}Co peak, and Region C is the ^{208}Tl peak.

prompt-coincident γ -ray) produce large signals in each PMT, meaning that these low energy afterpulses are quite common following a prompt signal. Indeed, the clusters that constitute these peaks tend to primarily be a few PE concentrated in a very small number of PMTs, supporting the hypothesis that these are afterpulses.

The rest of the low energy cluster time distribution is evenly distributed over time. These are the random backgrounds due to contaminants that produce a signal only in the LSV.

Random Background Spectrum

We can further divide random backgrounds into three categories: there are randomly coincident γ -rays coming from the LSV materials (such as the stainless steel sphere and PMTs) and γ -rays from the TPC cryostat that don't first scatter in the TPC; there are backgrounds due to PMT nose, such as dark noise and afterpulses; and there are backgrounds intrinsic to the scintillator. These backgrounds all limit the performance of the LSV by providing false coincidences that can decrease the acceptance of the LSV cuts and limit the strength of the cuts we can impose. For example, this was the case with the ^{14}C backgrounds during Phase-I: the high rate of ^{14}C forced us to make cuts that missed much of the neutron capture signal.




Figure 8.11: LSV cluster start time during the ~ 70 day underground argon campaign, showing the time distribution relative to the LSV trigger time of scintillation events in the LSV. The large spike around $-6 \mu\text{s}$ corresponds to prompt coincidences between the LSV and the TPC. The blue and green curves highlight high and low energy clusters, respectively, showing that the different origins of the features in the cluster start time profile.

The first category of random background is, by design, small. While the γ -rays from the detector material will have energy distributions much like those we expect for true coincidences, we have chosen materials to have very little radioactivity in order to maintain an activity in the LSV below 100 Hz. This low rate makes the probability of such a background appearing in the LSV relatively small.

Similarly, the dark rate in the PMTs is ~ 1 kHz. However, since dark noise tends to be ~ 1 PE, these events may contribute to very low energy signals, but the probability of the dark rate piling up between all the PMTs to produce a fake signal of an appreciable size is very small.

The primary contributors to the random backgrounds therefore come to the ^{14}C and ^3H intrinsic to the scintillator, as discussed in Chapter 3.1, and afterpulses following a prompt or random coincidence signal. The highest energy of these signals is ^{14}C , which β -decays with an endpoint of 156 keV, explaining why most of the random backgrounds have relatively low energies.

Low Energy Features: Tritium As can be seen in Region A of Figure 8.10, afterpulses dominate most of the low energy spectrum. To remove these afterpulses, we removed all clusters that followed another cluster by less than $30 \mu\text{s}$. Since this cut removes many clusters that are not afterpulses, we determine the probability of a non-afterpulse cluster being removed by this cut by integrating the spectrum




Figure 8.12: (*Top*) Photoelectron spectrum of the ^{14}C in the LSV in Phase-II, evaluated with the clustering method. The black curve is the measured low energy spectrum after applying cuts to remove afterpulses, and the red curve is the theoretical ^{14}C spectrum, convolved with a Gaussian response function, fit over the range of 35–130 PE. (*Bottom*) The ratio of the fit spectrum to the measured spectrum over the fit energy range.

produced by this cut between 160–200 PE. We compare this integral to the integral of a random coincidence spectrum made without this cut over the same PE range. To correct for the acceptance loss of this cut, we scale the spectrum with afterpulses removed by the ratio of these two integrals; this is the spectrum shown in Figure 8.12.

The lowest energy peak extending up to ~ 10 PE, likely comes from ^3H decays in the bulk of the scintillator. As discussed in Chapter 3.1.5, tritium is naturally found in both atmospheric and subsurface samples of hydrogen. Underground, it is produced by various nuclear reactions that maintain an equilibrium concentration of ^3H . An aggregate study of radioisotopes concentration in various subsurface samples has been performed by Lehmann et al. [184]. This study has found that various subsurface source of hydrogen have tritium ratios of $^3\text{H}/\text{H} \approx (0.12\text{--}2.7) \times 10^{-18}$.

The total LSV scintillator volume contains $\sim 1.8 \times 10^{30}$ hydrogen atoms. Since the TMB and PC were derived from subsurface petroleum, we can expect them to have tritium ratios roughly in this range. This assumption leads us to predict a tritium activity ranging from $\sim 390\text{--}8700$ Bq in the LSV.

As will be discussed later, estimates of the light yield of the LSV indicate a value of ~ 0.5 PE/keV. Since ^3H β -decays with an endpoint energy of 18 keV, we expect the endpoint in the observed spectrum to be ~ 10 PE, consistent with what we see examining the spectrum by eye.

As mentioned in Section 8.2, the top-down clustering algorithm used to generate this spectrum requires clusters to have an amplitude greater than 2 PE. Due to variance in the SPE height, scintillation events that produce fewer than 3 PE will rarely pass the clustering threshold. This is consistent with the data, where we see zero events below 2 PE and a dip in the spectrum between 2–3 PE. To estimate the rate of ^3H decays in the LSV, we integrate the observed spectrum from 3 PE to 10 PE. This 3 PE lower bound corresponds to an energy of 5.4 keV.

As can be seen in Figure 8.12, the ^3H spectrum seems to be sitting on top of another background. While the origin of this background is still unknown, a Gaussian fit to the falling edge from 10–20 PE reveals a mean of 7.8 PE and a total rate of ~ 850 PE. One possible explanation for this feature would be contamination of ^{210}Pb in the bulk of the scintillator; ^{210}Pb β -decays with a 16% branching ratio to a β with an endpoint of 63.5 keV and mean energy of 16 keV. The high value of Z for the ^{210}Pb nucleus means that the Fermi function in the β -spectrum significantly lowers the energy of the outgoing β . This rate of ^{210}Pb β -decays would be consistent with a total of ~ 1.9 ng of ^{210}Pb dissolved in the scintillator, which it may have picked up from the surfaces of the fluid handling system or other detector components that have been washed with water. While the 22 year half-life of ^{210}Pb means that it can be found out of secular equilibrium with the rest of the ^{238}U decay chain above it, it does not, however, explain why we do not see the 1.162 MeV endpoint β -spectrum from ^{210}Bi , which has a half-life of 5 days and is produced by the β -decay of ^{210}Pb .

Approximating this Gaussian background as the true shape of the background underneath the ^3H spectrum, we can subtract this spectrum from the ^3H spectrum to get the rate of ^3H decays producing a signal above 5.4 PE. The β -spectrum of ^3H is described in detail in Chapter 3.1.5. By integrating the expected spectrum, we find that 52% of all ^3H β -decays fall within this window. Scaling the background-subtracted rate, we estimate the ^3H contamination in the LSV to be ~ 300 Bq, consistent with reported measurements of the subsurface ^3H ratio.

Low Energy Features: Carbon-14 At higher energies in Figure 8.12, we see the ^{14}C spectrum, with an endpoint of 156 keV. This spectrum was discussed at length in Chapter 3.1.6.

In order to determine the ^{14}C contamination of the scintillator, we convolved the ^{14}C β -spectrum with with a Gaussian response function, which accounted for the Poisson and SPE charge variance in the variance. Since we did not generate the fit spectrum by Monte Carlo for this analysis, we could not use the model for the position dependence of the light yield, discussed in Chapter 7.3.1. Instead, we treated the relative variance of the light yield with the position as a free fit parameter, and added this to the Gaussian variance of the response function. The total ^{14}C decay rate, a constant background term, and the light yield relative variance were left free as fit parameters.

We assumed the β -spectrum was described by Equation 4.3, where we defined the shape factor as,

$$C(E) = 1 + aE \quad (8.1)$$

where a is an empirically determined parameter. This equation comes from the shape factor described in Equation 3.2, with $b = c = 0$ and a left as a free parameter in the fit.

Electron stopping powers were obtained using the NIST ESTAR tool [122]. However, this tool does not give stopping powers for electrons lower than 10 keV. Al-Ahmad et al. [185] have measured the the stopping power of 1–10 keV electrons in polyethylene and found that the stopping power decreases approximately linearly with electron kinetic energy in this range. We therefore linearly extrapolated the stopping power down to 0 keV for these calculations. Stopping powers calculated with Models I and II were generally between 0.97–1, with lower values at lower energies. Since we performed the fit between 35–130 PE, scintillation light produced by electrons below 10 keV constitutes a very small fraction of the total amount of light detected, and so errors due to this extrapolation are likely small.

The fit was performed over the range 35–130 PE. The lower bound of this range was chosen to be over the ^{210}Pb endpoint, to avoid interference from the lower energy parts of the spectrum, and the upper bound was chosen to extent beyond the ^{14}C endpoint to fit the constant background rate.

For this fit, we modeled scintillation quenching with Model II, using the scintillation parameter values we found in the $^{241}\text{Am}^{13}\text{C}$ calibration discussed in Chapter 9.3: $kB = 0.022 \text{ cm/MeV}$, $C = 5.0 \times 10^{-6} (\text{cm/MeV})^2$, and the light yield equal to 0.60 PE/keV. We obtained a total ^{14}C rate of $294 \pm 3 \text{ Bq}$ sitting on top of a constant background of $0.584 \pm 0.009 \text{ Hz}$, with a shape factor parameter $a = -0.65 \pm 0.69 \text{ MeV}^{-1}$. The low sensitivity of this measurement to a in this case means that this shape factor parameter is consistent with the measurements reported by Wietfeldt et al. [20], though it also consistent with $a = 0$.

It should be noted that the ^{14}C scintillation rate observed in the LSV during Phase-II is consistent with the factor of ~ 1000 reduction in the concentration of the original TMB batch, indicating that most of the current ^{14}C contamination may be explained by the residual high- ^{14}C TMB leftover in the scintillator.

In addition to uncertainty in the fit parameters intrinsic to the fit and the different models tested, there is also uncertainty due to potential contributions from other random backgrounds producing backgrounds in this energy range that vary in rate with energy.

Prompt Coincidence Spectrum

The full LSV energy spectrum can be seen in Figure 8.10. This figure shows the energy spectrum of all clusters as well as those whose start times fell within $\pm 25 \text{ ns}$ of the S1 signal in the TPC. As can be seen in this figure, most of the high energy signals in the LSV are contained within this 50 ns window, while much of the lower energy structure goes away when we constrain the signal to this prompt region.

The clusters found in this 50 ns time window are primarily those resulting from γ -rays scattering in the TPC and then in the LSV. By studying these coincidence, we can assess the contamination of various radioisotopes in the TPC components. H. Qian is currently leading an effort to validate the background model by fitting

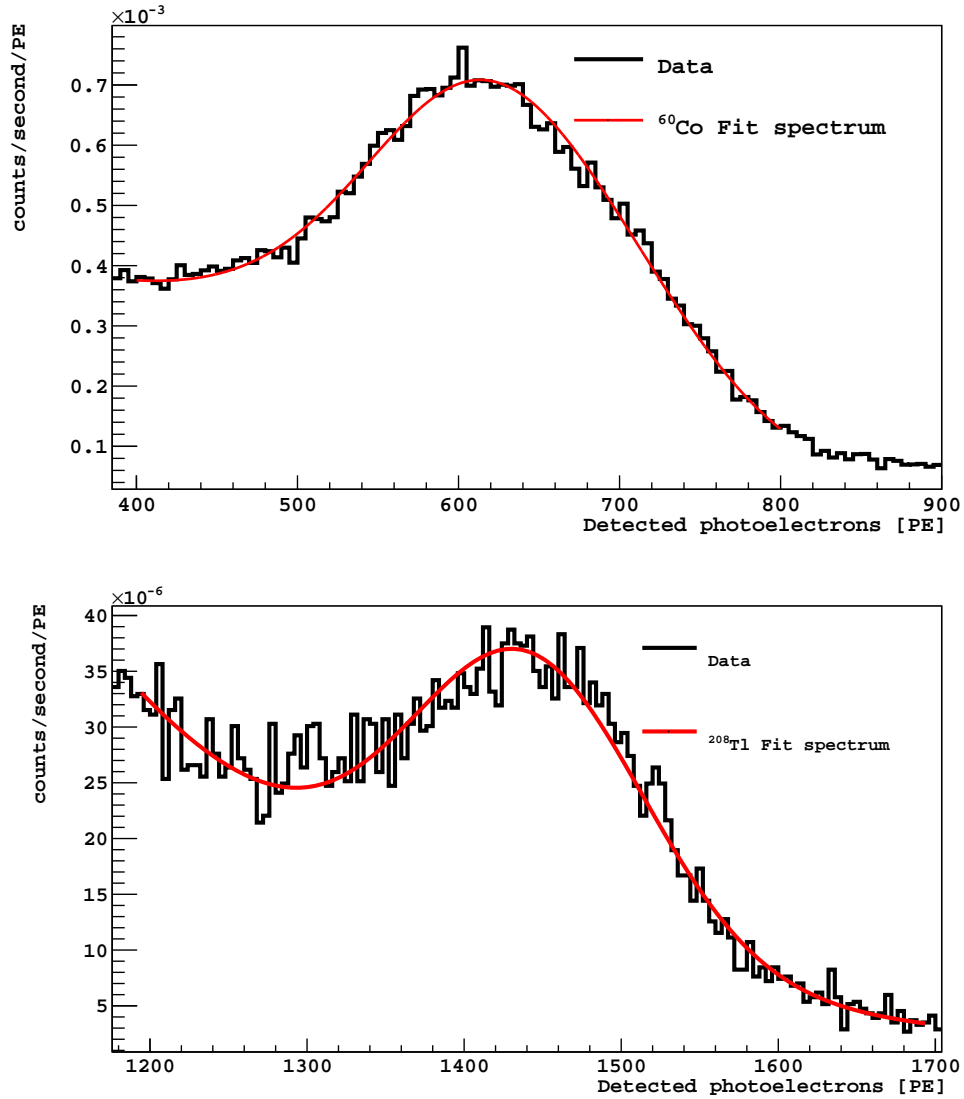


Figure 8.13: (*Top*) The (*black*) photoelectron spectrum of the ^{60}Co in the LSV in prompt coincidence with a signal in the TPC compared to (*red*) a fit of a simulated energy distribution convolved with a Gaussian response function. The fit, performed around the two peaks from the 1.17 and 1.33 MeV γ -rays, gives a light yield of 0.592 ± 0.011 PE/keV. (*Bottom*) The (*black*) photoelectron spectrum of the ^{208}Tl in the LSV in prompt coincidence with a signal in the TPC, compared to (*red*) a fit to a Gaussian plus a falling exponential fit to it. The 2.6 MeV γ -ray peak is seen at 1440 ± 5 PE.

simulated background spectra to this prompt spectrum to determine the isotopes' concentration. While his work is still in progress, so far the results appear consistent with the measurements described in Chapter 3.2.

While these results are not within the scope of the present document, there are two particular features in the prompt spectrum worth drawing attention to: the ^{60}Co spectrum and the ^{208}Tl spectrum. These spectra can be seen in more detail in Figure 8.13

Cobalt-60 The ^{60}Co spectrum can be seen in Region B of Figure 8.10. Since, as discussed in Chapter 3.1.3, ^{60}Co produces coincident 1.173 keV and 1.332 keV γ -rays in 99.85% of its decays, it produces very prominent prompt coincidence signals in the TPC and LSV when one γ -ray scatters in one detector while the other scatters in the other detector. Since both γ -rays are produced within picoseconds of each other, the only difference between when they produce signals in the TPC and LSV is the time of flight from their material of origin into the respective detectors.

GEANT4 simulations indicate that this time of flight difference is on the order of a couple of seconds, making the ^{60}Co signal a powerful tool for aligning the time variables in the TPC and outer detectors. Most of the spike around $-6\mu\text{s}$ in Figure 8.11 is due to these coincidences.

Simulating these decays in the cryostat, convolving the energy spectrum with a Gaussian response function, and fitting it to the data, as shown in Figure 8.13, gives a light yield of 0.592 ± 0.011 PE/keV. Due to the finite resolution of the LSV, the individual peaks from both γ -rays are not distinguishable.

Thallium-208 The ^{208}Tl spectrum can be seen in Region C of Figure 8.10, as well as in Figure 8.13. As discussed in Chapter 3.1.7, ^{208}Tl is part of the ^{232}Th decay chain and produces a 2.6 MeV γ -ray. This γ -ray is the highest energy γ -ray found in any of the common radioisotopes, making it easy to identify despite the relatively low rate of these events compared to ^{60}Co .

Furthermore, $\sim 85\%$ of these γ -rays are also accompanied by a 583 keV γ -ray. Much like we saw for ^{60}Co , if the 583 keV γ -ray scatters in the TPC while the 2.6 MeV γ -ray scatters in the LSV, we will reliably get a prompt coincidence signal, making this isotope very prominent in the LSV prompt coincidence spectrum.

Modeling the background as an exponential, we can fit a Gaussian distribution the full capture peak of the ^{208}Tl spectrum to estimate the light yield of the LSV. Doing so, we find the peak located at 1440 ± 5 PE, consistent with a light yield of 0.551 ± 0.002 PE/keV.

8.4.5 Muon Signals

In addition to creating a large signal in the WCV, a muon passing through the LSV makes an even larger signal, since the LSV can produce Cherenkov light in addition to scintillation light. Much of the work done studying these cosmogenic signals was done in collaboration with L. Marini and X. Xiang.




Figure 8.14: An example waveform of a cosmogenic signal in the LSV. The inset plot is zoomed in around the time that the muon passed through the detector.

An example waveform from such an event can be seen in Figure 8.14. The cosmogenic signal typically produces enough light to saturate the digitizers of each channel for $\gtrsim 100$ ns. This is followed by a very large overshoot in each of the digitizers, which is removed from the waveform by the overshoot corrections in the pulse reconstruction code. The waveforms then bounce back to a very high amplitude, followed by several microseconds of large signals produced by cosmogenically-induced electronic noise and afterpulses.

Figure 8.15 shows the total charge collected in the acquisition window for events during the ~ 70 day underground argon campaign. There are three noteworthy features in this distribution.

Above $\sim 100 \times 10^3$ PE, we see a population of events in coincidence with very large signals in the WCV. As we saw in Figure 8.5, this corresponds to events in which muons produced cosmogenic large signals in both the WCV and the LSV. Around $\sim 50\text{--}80 \times 10^3$ PE we see a second population of events with a large signal in the LSV and very little in the WCV. Inspecting the waveforms for several of these events has shown that these are due to electronics noise in which there are large oscillations in the LSV front-end boards. Similar noise is seen in the WCV electronics, where we see very little signal in the LSV and a large signal in the WCV; these signal have some overlap with events in which there is a cosmogenic signal in the WCV but not in the LSV.

During the DarkSide-50 ~ 70 day underground argon campaign, we vetoed signals in which the total charge collected in the LSV exceeded 2000 PE to remove cosmogenic

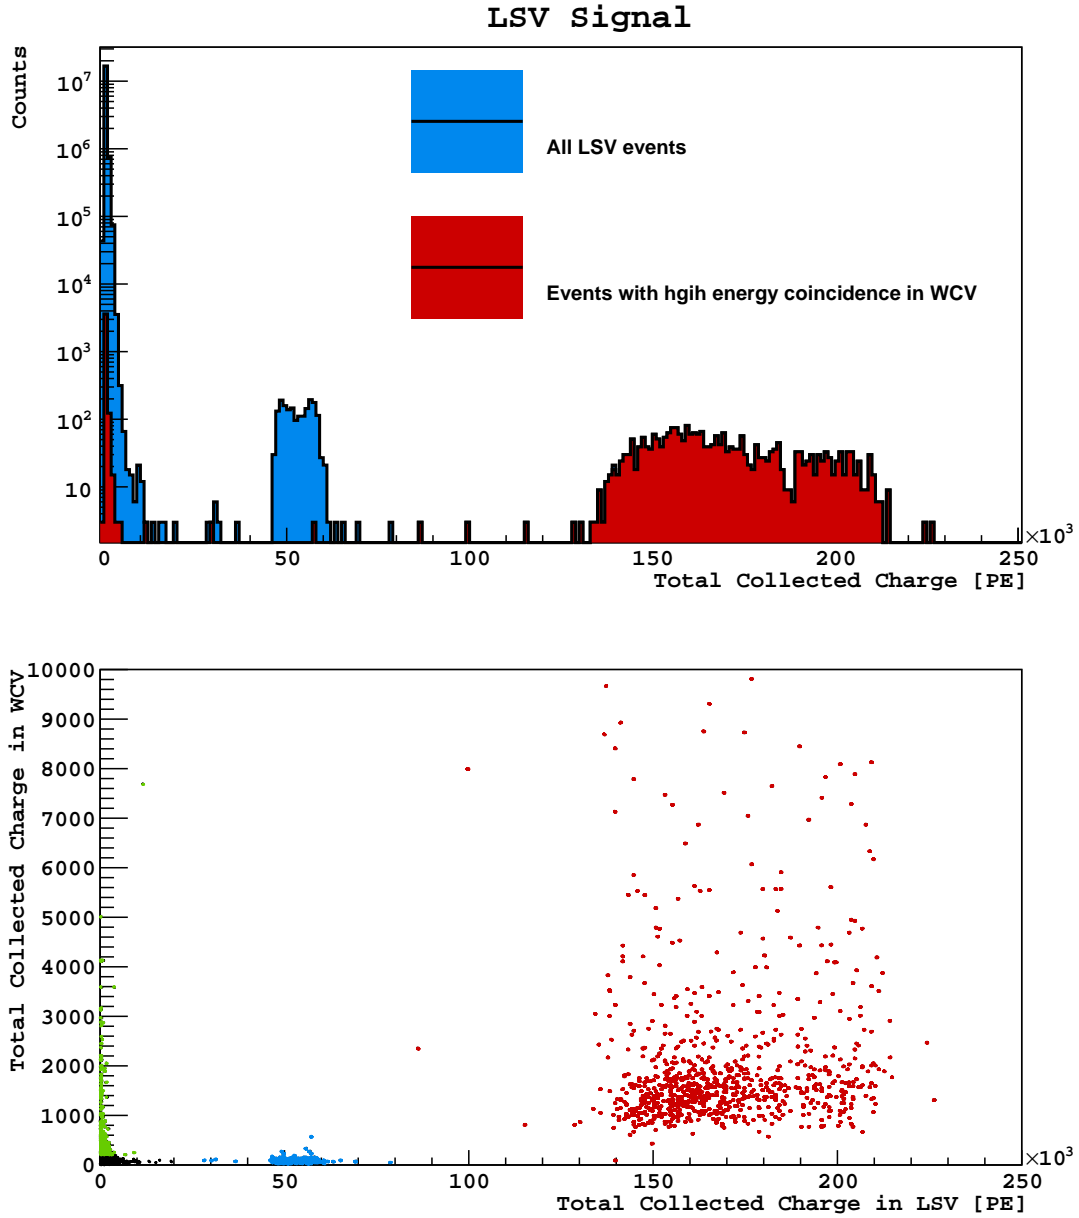


Figure 8.15: (*Top*) Total charge collected in the LSV for each event in the ~ 70 day underground argon campaign [42]. (*Blue*) All events, including those with cosmogenic signals and those without. (*Red*) Events with coincidental high energy signals in the WCV, predominantly due to cosmogenic signals in both detectors. (*Bottom*) Total charge collected in the LSV and in the WCV per event. (*Red*) Cosmogenic signals in the LSV and WCV. (*Blue*) Front-end oscillations in LSV electronics. (*Green*) Some front-end oscillations in WCV electronics and some cosmogenic signals. (*Black*) Standard LSV and WCV signals.

backgrounds. Data taken during a pulser run, as described in 8.3.2, indicate that this cut decreases the acceptance loss of the dark matter search by $0.14 \pm 0.10\%$.

8.4.6 Stability

The long-term stability of the LSV was assessed according to three metrics: the light yield stability, the LSV-TPC coincidence timing offset stability, and the event rate stability.

Light Yield Stability

In order to run for a long time, it is necessary that the light yield (which depends on the stability of the Lumirror reflectance, PMT quantum efficiency, and scintillator attenuation length and scintillation yield) remain high throughout the duration of the dark matter search campaign. The light yield stability is also important for placing cuts that will maintain their effectiveness over time.

The stability of the light yield was monitored by studying the position of the easily identifiable ^{60}Co peak over time. During Phase-I, which went from Nov. 6, 2013 to Jun. 3, 2014, the light yield increased with an average slope of 0.3 ± 0.1 PE/week. During Phase-II, which went from Apr. 8, 2015 to Jul. 31, 2015, the light yield increased with an average slope of 0.2 ± 0.1 PE/week.

While the causes for this slow increase in the position of the ^{60}Co peak are still not understood, this change in light yield is small enough that it can be neglected.

Coincidence Timing Offset Stability

In order to ensure that prompt cuts consistently identify prompt coincidences between the outer detectors and the TPC, it is important to monitor the stability of the timing offset between signals recorded by the two detectors. This offset relies on the time it takes the trigger signal to reach the outer detector DAQ, details of the trigger conditions used by the TPC DAQ, and other details relating to the TPC and outer detector electronics, DAQ, and trigger settings.

Possible timing offsets are monitored by looking at the distribution of the start times of clusters. True outer detector-TPC coincidences typically give a sharp timing peak in this distribution, as seen in Figure 8.11, and so cluster start times can be used to track the relative timing of the two systems with an accuracy < 10 ns.

The TPC coincidence timing remained stable throughout the duration of the dark matter search campaigns, with the exception of a couple small deviations where the trigger conditions or DAQ settings were changed. These deviations were accounted for in the outer detector reconstruction code.

Event Rate Stability

Lastly, we monitor other potential changes in the LSV's signal by tracking the event rate of the detector over time.

With the exception of a minor change in the data acquisition system made about a fifth of the way through Phase-II, which caused an increase in the event rate by $\sim 1.2\%$, the event rate was stable to within 0.03% /week over the course of the ~ 70 day campaign.

Following operations on the LSV scintillator, such as removing and replacing the original batch of TMB, we saw spikes in the region below ~ 20 PE. However, these spikes quickly diminished to the baseline rate in this energy range within a few hours of the operations ending. These low energy events tend to be very highly concentrated in a small number of phototubes given their total charge, compared to other clusters in this energy range that remained stable following the fluid handling operations. This concentration of the signal indicates that most of these clusters were caused by something producing a large signal in individual PMTs. Furthermore, these signals seemed to produce light with an unusually short time constant. One possible explanation is that the motion of the scintillator during the operations caused charge to build up, which then discharged on the μ -metal grids.

8.4.7 Dark Matter Search Cuts

As previously mentioned, DarkSide-50 has so far published two papers with WIMP exclusion curve. Each of these WIMP searches were done with different detector configures in both the TPC and the LSV, which are summarized in Table 8.3. As a result of these different run conditions, different cuts were made in the outer detectors.

During both phases, the strengths of the cuts were limited by their acceptance losses due to random backgrounds. During Phase-I, the high rate of ^{14}C events severely limited how strong the cuts could be, and special effort had to be made to design cuts with as high a chance of vetoing neutrons as possible without accidentally vetoing too many events. The much lower random background rate in the LSV during Phase-II meant that we were able to make much more aggressive cuts.

During Phase-I, we vetoed one event that we believe to be an (α, n) neutron that passed all TPC cuts based on its signal in the LSV, and one event that we believe to be a cosmogenic muon cut based on its signals in the WCV and LSV. During Phase-II, we saw one event that we believe to be a fission neutron that passed all TPC cuts except for a cut removing events with disproportionately large signals in a single PMT.

The cuts made during Phase-I and Phase-II along with their associated acceptances are summarized in Table 8.5. The acceptance is defined as the estimated probability of a WIMP event passing each cut, while acceptance loss is the estimated probability of a WIMP event failing the cut.

Muon Cuts

Example waveforms of a muon passing through the WCV and LSV can be seen in Figures 8.6 and 8.14, respectively. Both of these waveforms were recorded during the same event.

Table 8.5: Dark matter selection cuts made based on the WCV and LSV signals during Phase-I and Phase-II, along with their acceptances. These cuts are explained in more detail in the text.

Cut	Phase	Value	Acceptance Loss
Muon	I	total WCV charge ≥ 200 PE	$< 1\%$
	II	total WCV charge ≥ 400 PE	0.5%
		total LSV charge ≥ 2000 PE	
Prompt	I	prompt LSV ROI ≥ 10 PE	5%
	II	prompt LSV ROI ≥ 1 PE	0.9%
Delayed	I	delayed LSV slider ≥ 80 PE	6%
		late LSV slider ≥ 110 PE	
	II	delayed LSV slider ≥ 3 or 6 PE	16%
Pre-prompt	I	—	—
	II	pre-prompt LSV slider ≥ 3 PE	0.9%

For the Phase-I analysis, events were rejected if we saw more than 200 PE in the WCV during the course of the event.

As shown in Table 8.3, we lengthen the acquisition window part way through Phase-II in order to more reliably be able to detect neutron captures. The longer acquisition window meant that the total background per event increased in the WCV, and so we had to increase the threshold for this cut in order to maintain a comparably low acceptance loss. However, we added a cut on the total charge collected in the LSV to help compensate for events that might have lower signals in the WCV. For the Phase-II analysis, we vetoed the event if we saw more than 400 PE in the WCV or 2000 PE. For the most part, we expect the cut on the total charge in the LSV to be redundant with the other LSV cuts made, since a signal this large should fail at least one if not both of the other cuts. However, we included this cut as a safety precaution in case the other cuts fail to see the cosmogenic signal.

Prompt Cut

The prompt cut was designed to veto neutron thermalization signals, as well as γ -ray signals that may be in prompt coincidence between each detector and γ -rays that may be produced by neutrons inelastically scattering on TPC and cryostat materials.

This cut was defined using a temporal region of interest, as defined in Section 8.2.2. The time in the LSV acquisition window corresponding to prompt coincidences was tuned using prompt coincidence γ -rays—primarily ^{60}Co —as described in Section 8.4.4; this time was found to be around $-6.5\mu\text{s}$. A temporal region of interest was defined extending from $[-10, 200]$ ns and $[-50, 250]$ ns around the prompt time during Phase-I and Phase-II, respectively. The lower background rate in Phase-II allowed us to extend the size of this window without suffering additional acceptance loss while increasing our confidence that the prompt signal will be collected in this time window.




Figure 8.16: The waveform seen in the LSV for an event seen during the 70 day underground argon campaign published in [12]. This event produced a single nuclear recoil in the WIMP energy range in the TPC. This signal was vetoed by the LSV due to the large prompt signal seen around $-6.5 \mu\text{s}$ and in the left inset plot, as well as by the two capture signals: one at $7.8 \mu\text{s}$ and one (seen in the right inset plot) at $69.9 \mu\text{s}$. This signal was also vetoed because it produced light with a high concentration in one TPC PMT, indicating that it may have been a fission event in the TPC fused silica in which at least two neutrons were produced.

During Phase-I, we vetoed events with ≥ 10 PE in the prompt window, while we vetoed events with ≥ 1 PE in this window during Phase-II. The cuts are summarized in Table 8.5

Figure 8.16 shows an example waveform from an event seen during Phase-II, during the ~ 70 day campaign published in [12]. This event produced a single nuclear recoil in the TPC in the WIMP energy window. However, as we can see in this figure, it produced a very large prompt signal, with 1438.14 PE collected in the prompt window followed by a long tail of afterpulses. This signal is roughly equivalent to a $\sim 26,000$ keV nuclear recoil—more energy than we expect to see in an (α, n) reaction. However, as we will discuss shortly, we also see what appear to be two capture signals, indicating that two neutrons may have been produced in this event. While the (α, n) reaction does not produce more than one neutron at a time, fission reactions may do so. Since fission reactions may also produce high energy γ -rays, it is possible that this large prompt signal was caused by a γ -ray scattering in the LSV at the same time that two neutrons thermalized.

In addition to vetoing this event by its LSV signals, this event also produced a disproportionately large signal in one of the TPC PMTs. As discussed in Chap-

ter 4.1.2, we reject such events in order to avoid backgrounds due to Cherenkov light in the fused silica windows of the TPC. It is likely that this fission event, among the many typical fission products, also produced a high energy electron that produced Cherenkov light. Therefore, we can conclude that this event was likely due to a fission occurring in the TPC fused silica windows.

Since this event has clear prompt and delayed signals, it is an instructive example of the LSV cuts at work.

Delayed Cuts

In order to detect neutron capture signals, we defined a set of sliding window variables, as explained in Section 8.2.2.

As discussed in Chapter 5.2, thermal neutrons capturing on ^{10}B have a 6.4% chance of producing the ground state of ^7Li , along with a 1.775 MeV α . The signal from this recoiling ^7Li nucleus and α is very heavily suppressed by scintillation quenching. However, the reaction has a 93.6% chance of producing the excited state of ^7Li , which also produces a 478 keV γ -ray, which makes a large and easily detectable signal in the LSV.

During Phase-I, the background from ^{14}C below 156 keV was too high to veto on the signal from the neutron capture going to the ground state of ^7Li ; any attempt to do so would also result in an acceptance loss close to 100%.

Instead, we tuned the cuts to try to detect the higher energy neutron capture signals that we could see above the ^{14}C backgrounds. While doing so greatly reduces the maximum achievable neutron vetoing efficiency of the LSV, we found that it was enough for the 50 day atmospheric argon campaign that ran during Phase-I and was published in [50]. These less aggressive cuts also let us detect these higher energy signals with only taking an acceptance loss of 6%.

In order to increase the sensitivity of this cut, we defined two separate 300 ns wide sliding windows: the delayed sliding window scanned the region from the prompt time to 8.8 μs afterwards, and the late sliding window scanned the acquisition window from the end of the delayed window to the end of the acquisition window.

Since the thermal neutron capture time for Phase-I was 2.2 μs , the delayed sliding window covered the first four neutron capture times, and so had a lower threshold of 80 PE. The late window, which was less likely to detect a neutron capture, then had a higher threshold of 110 PE.

During Phase-II, the LSV had much lower background, making the full neutron capture signal detectable. For data collected during this phase, we defined a single 500 ns wide sliding window extending from the prompt time to the end of the acquisition window.

As we will discuss in Chapter 9, the ^{10}B neutron capture that goes to the ground state of ^7Li produces a signal with a peak between 20–30 PE, depending on how close the capture occurred to the cryostat. However, statistical variance in signal size—largely due to Poisson statistics and the SPE variance of the PMTs—means that some of these captures can produce signals below these values. In order to detect

these signals reliably, we therefore want to place the cuts as low as we can without accruing too much acceptance loss.

The long thermal neutron capture time of $22\,\mu\text{s}$ meant that time region scanned by this sliding window had to be much longer than during Phase-I. Scanning more time means that even with a low background rate, the probability of an accidental coincidence increases. Because we increased the size of the acquisition window about halfway through the ~ 70 day campaign, it was necessary to increase the value of the cut with the longer window, too, in order to keep the acceptance loss $\sim 16\%$.

Therefore, as described in Table 8.5, we vetoed events if the sliding window had a charge greater than 3 PE for the first part of this campaign, and if the window had a charge greater than 6 PE for the latter part. While this higher threshold increases the chances that a statistical fluctuation in the signal size of the neutron capture go below our cut threshold, we estimated that the efficiency gain from having a longer window was greater than the loss from the higher threshold. This trade-off motivated the change in DAQ settings.

In Figure 8.16, we can see two such neutron captures at $7.8\,\mu\text{s}$ and $69.9\,\mu\text{s}$, corresponding to two fission neutrons capturing in the LSV. These captures produce signal of 259.9 PE and 278.5 PE, respectively, corresponding to neutron captures on ^{10}B going to the excited state of ^7Li . Each capture signal is then followed by a train of afterpulses due to the large signals that they produce.

Pre-prompt Cuts

During Phase-II we defined a second 500 ns sliding window, scanning from the start of the acquisition window to the prompt time.

The purpose of this cut is to check if any signal has proceeded a signal in the TPC. This might happen if a delayed neutron is emitted near the TPC, or if an external neutron with enough energy to penetrate the LSV—such as we might see from a cosmogenic neutron—scatters in the LSV on its way to the TPC. Since we only trigger on the TPC, it is also possible that the outer detectors may only catch the end of a cosmogenic signal, leaving the TPC vulnerable to cosmogenic neutrons. Since this cut checks the beginning of the acquisition window, it may be able to catch these tails. Furthermore, if fluctuations in the timing move a prompt signal out of the prompt temporal region of interest, this sliding window may be able to catch it.

During the ~ 70 day underground argon WIMP search, we vetoed events with a pre-prompt signal ≥ 3 PE. Data taken during a run in which the outer detectors triggered on a pulser indicated that this cut had an acceptance loss of 0.9%.

Chapter 9

Neutron Source Calibration

We have now quantified the neutron backgrounds relevant for running a background free dark matter experiment, developed a neutron veto system to efficiently remove these backgrounds, and built such a veto for DarkSide-50. The next logical steps are to study this veto detector’s response to neutron interactions and to simulate the background neutrons to determine how efficiently these detectors can veto neutrons.

Although we studied much of the physics and optics related to these responses in the prototype tests described in Chapter 6, the current LSV uses a slightly different scintillator cocktail, and so all of the parameters describing the detector’s response cannot be assumed to be the same as in the prototype tests. However, the results of the tests are still useful to guide us in determining the current LSV cocktail properties, as they are likely not drastically different.

We accomplish our first goal of quantifying the detector’s response through two calibration campaigns, performed with two (α , n) calibration campaigns: one using an $^{241}\text{AmBe}$ neutron source, discussed in Section 9.2, and one with an $^{241}\text{Am}^{13}\text{C}$ neutron source, discussed in Section 9.3.

We will accomplish the second goal with a series of simulations based on what we’ve learned from these calibrations and the prototype studies discussed in Chapter 6, the details of the veto performance and cuts, and the expected neutron backgrounds. These simulations will tell us the efficiency of the veto and the expected neutron background of DarkSide-50.

9.1 Calibration Insertion System

The CALibration Insertion System (CALIS) was designed by the calibration group at University of Hawaii to deploy calibration sources inside of the LSV close to the cryostat. Figure 9.1 shows a diagram and a photograph of CALIS inserted into the LSV.

CALIS consists of a long arm that can be lowered into the LSV through one of the organ pipes at the top and a counterweighted joint that can bend at a 90° angle to hold the source up against the cryostat. The arm can be rotated around the axis

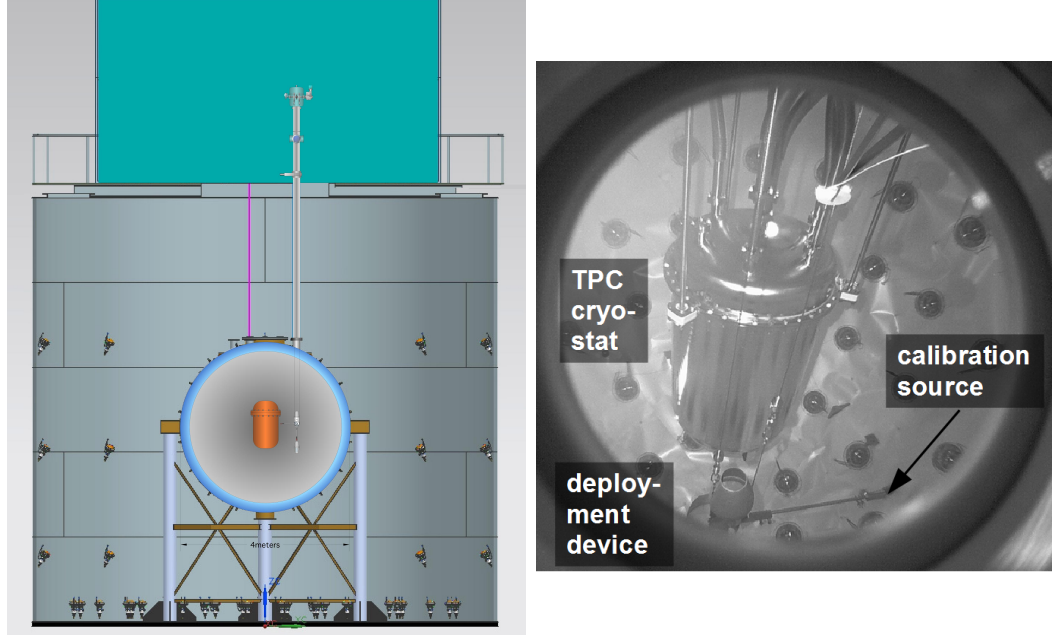


Figure 9.1: (*Left*) A diagram of CALIS lowered into the LSV from the port in the clean room, from [40]. (*Right*) Photograph of the inside of the LSV with CALIS deployed inside of it, from [40].

parallel to the organ pipe, so that the source can be moved to different positions relative to the cryostat.

While CALIS has been used for deploying various γ -ray and neutron sources, we will primarily discuss the calibrations done with the $^{241}\text{AmBe}$ source (discussed in Section 9.2) and with an $^{241}\text{Am}^{13}\text{C}$ source (discussed in Section 9.3).

9.2 Americium-Beryllium Source Calibration

The first calibration campaign took place January and February 2015. This calibration campaign used a $^{241}\text{Am}^9\text{Be}$ source to generate neutrons through the $^9\text{Be}(\alpha, n)^{12}\text{C}$ reaction. The work discussed in this section was done with extensive collaboration with the calibration group at University of Hawaii, including J. Maricic, B. Reinhold, E. Edkins, and B. Hackett, as well as H. Qian and M. Wada.

A photograph of the $^{241}\text{AmBe}$ source can be seen in Figure 9.2. The source is composed of a beryllium disk with ^{241}Am mixed in with it. ^{241}Am α -decays with a Q -value of 5.638 MeV. One of the three primary decay modes has a branching with a 5.486 MeV α and a branching ratio of 84%. This branch is accompanied by a 59.5 keV γ -ray 94% of the time, and by a 26.3 keV γ -ray followed by a 33.2 keV γ -ray 6% of the time. Another 13.1% of decays are accompanied by a 5.443 MeV α , followed by an array of γ -rays ranging from 27–103 keV, while 1.7% have a 5.388 MeV α with an array of γ -rays of similar energy. Since 100 keV γ -rays have an attenuation length of ~ 170 nm in lead [158], these γ -rays are very effectively blocked by putting the source in ~ 2 mm of lead shielding.



Figure 9.2: A photograph of the $^{241}\text{AmBe}$ source, opened up. Photograph taken by Y. Suvarov.

The source was deployed into the LSV using the CALIS insertion system described in Section 9.1.

9.2.1 Neutron Spectrum

Because the ^{241}Am is mixed in with the beryllium target, α s involved in the (α, n) reaction have an approximately isotropic angular distribution; as a result, neutrons emitted by this source are also isotropically emitted. As the α s travel through the beryllium and slow to a stop, they have some probability of capturing on a ^9Be nucleus and forming a compound $^{13}\text{C}^*$ nucleus with excitation energy ranging from 9.9–15.5 MeV, depending on the energy of the α when it captures.

As shown in Figure 9.3, the compound $^{13}\text{C}^*$ nucleus decays to ^{12}C by emitting a neutron through the reaction

$$^{13}\text{C}^* \rightarrow \begin{cases} ^{12}\text{C} + n & (\text{Br} = 36\%) \\ ^{12}\text{C}^* + n, ^{12}\text{C}^* \rightarrow ^{12}\text{C} + \gamma (4.439 \text{ MeV}) & (\text{Br} = 61\%) \\ ^{12}\text{C}^{**} + n, ^{12}\text{C}^{**} \rightarrow ^{12}\text{C}^* + \gamma (3.215 \text{ MeV}) & (\text{Br} = 3\%) \\ ^{12}\text{C}^* \rightarrow ^{12}\text{C} + \gamma (4.439 \text{ MeV}) \end{cases} \quad (9.1)$$

with a Q-value of -4.9 MeV. Neutrons emitted by this reaction therefore have energies ranging between 0–11.2 MeV.




Figure 9.3: Diagram illustrating the $^{241}\text{AmBe}$ (α , n) reaction. An α emitted by an ^{241}Am nucleus interacts with a ground state ^9Be nucleus to form a compound ^{13}C nucleus. This ^{13}C nucleus then decays to the ground state, first excited state, or second excited state of ^{12}C by emitting a neutron. Depending on which state the ^{12}C ends up in, it may release up to two γ -rays as it decays to the ground state. Black lines represent nuclear states. Numbers to the left of each black line represent the J_π value of that state; numbers to the right are the excitation energy of the state in keV. Branching ratios were tuned by Borexino, from I. Machulin.




Figure 9.4: Neutron energy spectrum from the $^{241}\text{AmBe}$ source, from [43]. Spectra are normalized so that the integral of each spectrum is the probability of the given decay mode. (Blue) The spectrum for the $^{241}\text{AmBe}$ reaction resulting directly in the ground state of ^{12}C . (Green) The spectrum for the $^{241}\text{AmBe}$ reaction going to the first excited state of ^{12}C , which emits a 4.439 MeV γ -ray. (Purple) The spectrum for the $^{241}\text{AmBe}$ reaction going to the second excited state of ^{12}C , which emits 3.215 and 4.439 MeV γ -rays.




Figure 9.5: The neutron capture signal measured during the $^{241}\text{AmBe}$ calibration campaign, showing the signal produced by neutrons capturing on ^{10}B . The signal produced by the $\alpha + ^7\text{Li}$ in the capture channel going directly to the ground state of ^7Li can be seen at $29.2 \pm 0.2 \text{ PE}$. The signal produced by the $\alpha + ^7\text{Li} + \gamma$ (478 keV) produced by the capture channel that goes to the excited state of ^7Li can be seen at $264.6 \pm 0.1 \text{ PE}$.

Figure 9.4 shows the neutron spectra for all three reaction channels. As can be seen in this figure, these spectra span a wide range of energies, comparable to the (α, n) spectra we expect to be present as potential background sources in DarkSide-50, as discussed in Chapter 4.4.3. However, the high energy γ -rays in frequent coincidence with these neutrons makes the neutron thermalization signal difficult to study, as much of the γ -ray signal will overlap with the signal created by the neutrons slowing down.

9.2.2 Neutron Capture Signal

While the neutron thermalization signal is difficult to study due to the prompt coincidence γ -rays coming from the source, the delayed capture signal should be similar to what we expect from background neutrons.

To measure the neutron capture spectrum, we took data triggering the outer detectors on the LSV with the source held close to the cryostat. In order to avoid afterpulses, we waited at least $20 \mu\text{s}$ after the prompt pulse to start looking for the neutron captures and only looked at the first cluster in this time window. A background spectrum was recorded by taking data with CALIS deployed empty; this spectrum was scaled for live time and subtracted from the neutron capture spectrum.




Figure 9.6: The time distribution for neutron captures measured during an $^{241}\text{AmBe}$ calibration campaign. (*Black*) The low energy and (*blue*) high energy peaks, corresponding to neutron captures on ^{10}B leading to the ^7Li ground state and excited state, respectively. (*Red*) An exponential plus a constant were fit to both of these curves (shown here only for the low energy peak curve), giving a decay time of $22.1\text{ }\mu\text{s}$, consistent with the expected capture time for thermal neutrons during Phase-II.

Figure 9.5 shows the neutron capture signal for the Phase-II scintillator. The ^7Li and α from the capture channel that goes directly to the ground state of ^7Li can be seen at $29.2\pm0.2\text{ PE}$; the signal going to the excited state of ^7Li with the coincident 478 keV γ -ray can be seen at $264.6\pm0.1\text{ PE}$. At lower energies, near 0 PE , we see some background, primarily due to afterpulses leaking past the cuts.

The higher energy signal appears with a branching ratio relative to the lower energy signal of $\sim 95.6\%$, comparable to the expected value of 94.6% . This agreement implies that few γ -rays created by these neutron captures enter the cryostat before depositing most of their energy in the scintillator, since neutron captures in which the 478 keV γ -ray is lost will look similar to those decays going directly to the ground state.

In order to measure the thermal neutron capture time in the LSV and verify the identity of the two capture peaks, we made a histogram of the time between the prompt signal and the capture signal during the $^{241}\text{AmBe}$ campaign. Figure 9.6 shows the time distribution for events falling in the energy range of the high and low energy peaks. In order to determine the neutron capture times, we fit an exponential plus a constant to each of these curves and extracted the decay time of the exponential. In this fit, the constant is a measure of the random backgrounds present in each energy window.

Performing these fits gives capture times of $22.1 \pm 0.1 \mu\text{s}$ and $22.1 \pm 0.2 \mu\text{s}$ for the low and high energy peaks, respectively. These measurements are consistent with the capture times expected 5% TMB.

PPO Concentrations Dependence

The $^{241}\text{AmBe}$ calibration campaign started between Phases I and II of DarkSide-50, and the scintillator was reconstituted during the course of this campaign. In particular, at the start of the campaign, the LSV had 0.7 g/L of PPO, and the PPO concentration was increased to 1.4 g/L part way through the calibration campaign.

While the spectrum shown in Figure 9.5 were taken with 1.4 g/L of PPO and are therefore applicable to analysis of the Phase-II data, these two different concentrations of PPO gave us the ability to study the effects of the PPO concentration on the scintillator response.

These studies were carried out by H. Qian, who found that increasing the PPO concentration from 0.7 g/L to 1.4 g/L increased the position of the low energy peak by 3 PE ($\sim 11\%$), but decreased the position of the high energy peak by 13 PE ($\sim 5\%$).

It is possible that the high energy peak decreased in energy either because of impurities accidentally added to the scintillator in process of adding more PPO, or that the higher concentration lead to more light being lost to PPO reabsorbing its own scintillation light and non-radiatively dissipating it. However, it is notable that despite the lower light yield of the high energy peak, the low energy peak increased in photoelectrons. This is likely evidence that higher PPO concentrations decrease the effects of ionization quenching within this concentration range.

This relationship is important because the efficiency with which the neutron veto can detect this low energy peak is a key component of the overall neutron veto efficiency. Since this low energy signal is heavily quenched, decreasing the effects of ionization quenching by increasing the PPO concentration could further improve the efficiency of a future neutron veto with a similar design. Further investigation into this relationship could reveal the optimal concentration of PPO needed to achieve the maximum strength from this capture signal.

Position Dependence

As can be seen in Figure 7.4, the light yield in the LSV is mostly uniform in the bulk of the scintillator, far from the PMTs and the cryostat. However, most neutrons that we need to veto with the LSV will typically thermalize and capture within a few centimeters of the cryostat. It is therefore important to understand how the light yield varies near the cryostat, so these optical effects can be accounted for in our full background analysis.

In particular, in order to construct Figure 7.4, we needed to determine the reflectance of the electropolished stainless steel comprising the outer surface of the cryostat, which was not measured *a priori*.

In order to do so, we took data with the source up against the cryostat, and with the source rotated 90° , placing it $\sim 72\text{ cm}$ away from the cryostat. Data was

collected with the LSV in self-trigger mode; if at least 20 PMTs registered a signal above ~ 0.5 PE within a 60 ns window, a trigger signal was sent to the outer detector digitizers.

Measuring the neutron capture spectrum with the source in both positions showed that the low energy peak decreased by $\sim 26\%$ when the source was close to the cryostat relative to when the source was rotated away, while the higher energy peak decreased by $\sim 10\%$. This effect is more dramatic for the signal from the $\alpha+{}^7\text{Li}$ than it is for the signal coming from the $\alpha+{}^7\text{Li}+\gamma$. This difference is likely because the γ -ray may travel ~ 20 cm before scattering, while the α and ${}^7\text{Li}$ will deposit all of their energy near where the neutron captured. Since the γ -rays travel farther before scintillating, they may travel far from the cryostat and into the bulk even while the source is rotated up against the cryostat, meaning that the γ -ray signal strength is less dependent on the neutron capture position than the lower energy signal is.

S. Davini studied this effect in an optical Monte Carlo of the source in both positions, and found that a comparable effect was seen in the simulation as was observed in the data if the reflectance of the cryostat were set to a constant value of 48%. We therefore used this reflectance in the optical Monte Carlo.

9.3 Americium-Carbon Source Calibration

As mentioned in Section 9.2, the ${}^{241}\text{AmBe}$ source produces too many high energy γ -rays and too broad a neutron spectrum to provide a clear signal for calibrating the response of the LSV to neutron thermalization signals. To circumvent this problem, we designed and built a ${}^{241}\text{Am}{}^{13}\text{C}$ neutron source that would produce ~ 5 MeV neutrons, with little coincident γ -ray background, similar to the source used by the Daya Bay Collaboration [186].

The overall design of this source was lead by J. Martoff, and the source itself was assembled by M. Wada and Y. Suvarov. The source was deployed into the LSV using the CALIS insertion system designed by the calibration group at University of Hawaii, including J. Maricic, B. Reinhold, E. Edkins, and B. Hackett. Furthermore, much of the analysis presented here was done in close collaboration with H. Qian and M. Wada using the G4DS implementation of GEANT4, whose design was lead by the Monte Carlo group, coordinated by D. Franco.

9.3.1 Source Design

A photograph and a schematic diagram of the ${}^{241}\text{Am}{}^{13}\text{C}$ source is shown in Figure 9.7. The ${}^{241}\text{Am}$ source, inside a stainless steel holder, is placed inside of a lead container with ~ 2 mm thick walls. A graphite pellet made from ${}^{13}\text{C}$ is wrapped in gold foil and placed on top of the ${}^{241}\text{Am}$ source. This lead container blocks the low energy γ -rays emitted by the ${}^{241}\text{Am}$ source.

The ${}^{241}\text{Am}$ source used here is Type AFR from The Eckert & Ziegler Group. It consists of a $1\text{ }\mu\text{m}$ *active matrix* of $100\text{ }\mu\text{Ci}$ of ${}^{241}\text{Am}$ embedded in a gold foil. This




Figure 9.7: (*Top*) A schematic diagram of the $^{241}\text{Am}^{13}\text{C}$ source holder, showing the lead holder and its lid, the ^{241}Am source in its container, the ^{13}C pellet wrapped in a $1\ \mu\text{m}$ thick sheet of gold foil, and the polyethylene spacer. (*Bottom*) A photograph of the $^{241}\text{Am}^{13}\text{C}$ source, covered in gold foil with the lid removed. Photograph taken by Y. Suvarov.

Table 9.1: Composition of the $^{241}\text{Am}^{13}\text{C}$ source graphite pellet.

^1H	0.21%
^2H	$4.22 \times 10^{-5}\%$
^{12}C	1.66%
^{13}C	97.02%
^{16}O	1.11%
^{17}O	$4.46 \times 10^{-4}\%$
^{18}O	$2.23 \times 10^{-3}\%$

matrix is plated on a 0.15–0.25 mm thick silver substrate, and a 2 μm protective gold layer.

When the ^{241}Am decays, it emits an α with kinetic energy typically between 5.4–5.5 MeV. 37% of these decays are accompanied by a 13.9 keV γ -ray, and 45.9% are accompanied by a 59.5 keV γ -ray. These low energy γ -rays, as well as an array of comparably low energy γ -rays with branching ratios in the range of 10^{-6} –1% are easily blocked by the lead shielding surrounding the source. There is also an array of higher energy γ -rays spanning all the way up to 921 keV [45], however, their branching ratios are typically on the scale of 10^{-7} – $10^{-6}\%$, rendering them unlikely to coincide with an (α , n) reaction.

As the α travels through the matrix, protective gold plating, and gold foil, it will slow down until it reaches the ^{13}C pellet. The pellet is composed of a powder of 99% pure ^{13}C mixed with a $\text{C}_4\text{H}_6\text{O}_2$ binding agent making up 3% of the total mass of the pellet. The powder and the binding agent are pressed together with a hydraulic press, and have a mass of 315 mg/cm². The approximate isotopic composition of the pellet is given in Table 9.1.

As the α thermalizes inside of this pellet, it may capture on a ^{13}C nucleus and undergo the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction through a ^{17}O compound nucleus. This reaction is exothermic with a Q-value of 2.2 MeV.

A diagram of this reaction is shown in Figure 9.8. Since ^{16}O , the daughter nucleus of this reaction, is doubly magic, its first excited state is at a relatively large excitation energy of 6.049 MeV. As we saw in Equation 4.21, if the excitation energy of the final state is greater than the Q-value of the decay, Q_m may be negative. If $Q_m < 0$, the radicands in Equation 4.21 may be less than zero, meaning that the outgoing neutron energy may not be real for some values of the α kinetic energy T_α . Requiring that the neutron energy be real, we find that the minimum α energy needed to excite the m^{th} energy level of the final state nucleus is given by,

$$T_\alpha > (E_{ex}^m - Q) \left(\frac{m_t + m_\alpha}{m_t} \right) \quad (9.2)$$

where E_{ex}^m is the excitation energy of the m^{th} excited state, Q is the Q-value of the reaction, and m_t and m_α are the mass of the target nucleus and α , respectively. Evaluating this expression for the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction, we find that the first excited




Figure 9.9: GEANT4 simulations of the energy deposited in the ^{13}C pellet by ^{241}Am α s propagating through the source, through (*blue*) $0\ \mu\text{m}$ of foil, (*green*) $1\ \mu\text{m}$ of foil, or (*purple*) $2\ \mu\text{m}$ of foil, and into the pellet.

layer, and the gold foil, allowing them scatter along the way. Finally, the total energy deposited in the ^{13}C pellet was recorded.

The distribution of energy deposited in the graphite pellet is shown in Figure 9.9. As can be seen in this figure, the energy deposition distributions all have a sharp high energy cutoff with a long tail going down to 0 MeV. The sharp cutoff comes from the case where the α travels normal to the surfaces between the ^{241}Am source and the pellet, taking the shortest possible path and therefore attenuating as little as possible. The long tail going down to 0 MeV is a result of the α s that travel at angles relative to surfaces and therefore travel a longer path and are more heavily attenuated before they reach the graphite pellet.

As can be seen in Figure 9.9, even with no additional layers of foil, the protective gold coating moderates the α s enough to prevent them from depositing more than ~ 4.5 MeV in the graphite pellet, already well below the 5.05 MeV cutoff for reaching the first excited state of ^{16}O . With $1\ \mu\text{m}$ of foil added between the ^{241}Am source and the pellet, the maximum energy deposited into the pellet decreases to ~ 4 MeV, and with two layers, making a total of $2\ \mu\text{m}$ of foil, the maximum energy is ~ 3.5 MeV.

Even though these simulations indicated that no foil was needed to sufficiently attenuate the ^{241}Am α s to avoid exciting the daughter ^{16}O nucleus, we decided to use one layer of foil to avoid leakage from the tail of the distribution.

9.3.2 Neutron Energy-Angle Distribution

Since the ^{241}Am and beryllium were mixed together in the $^{241}\text{AmBe}$ source, we were able to assume an isotropic angular distribution of the α s when they underwent the (α, n) reaction on a ^9Be nucleus. We were therefore able to ignore the relationship between the outgoing neutron's direction and energy and assume that the energy was uncorrelated with the direction. Since the ^{241}Am and ^{13}C are separated by the gold foil in the $^{241}\text{Am}^{13}\text{C}$ source, such an assumption is no longer valid. Instead, α s that enter the graphite pellet do so with a net momentum orthogonal to the gold-graphite interface. Understanding the neutron energy spectrum therefore requires understanding how the neutrons energy and angle relate to each other.

The neutron spectrum and energy-angle distribution was computed using both GEANT4 and TALYS. As described above, the $^{241}\text{Am}^{13}\text{C}$ source geometry was implemented in GEANT4, and ^{241}Am α s were generated uniformly in the active matrix and propagated into the graphite pellet.

At each step taken by the α as it thermalizes in the pellet, we determine that the α is traveling relative to the normal of the gold-graphite interface by calculating the vector connecting the current location of the α with its location at the previous step.

We assume that the probability an α capturing at any given step is small, and therefore that the probability of a neutron capturing at each step is independent of the previous steps.

Given the kinetic energy of the α at each step, we use TALYS to determine the cross section for the α undergoing the (α, n) reaction on each nucleus given in Table 9.1 and generate an (α, n) neutron spectrum for each isotope. The isotopes are then scaled by a factor of R_i , given by,

$$R_i = \frac{C_i \rho_p N_A s}{A_i}$$

where C_i is the mass fraction of isotope i , ρ_p is the total density of the pellet, N_A is Avogadro's number, s is the size of the simulated step, and A_i is the mass number of the isotope.

These scaled spectra are summed together to determine the total neutron spectrum produced at that step. The angle of the outgoing neutron relative to the momentum of the incident α can then be constrained by kinematics; this relationship is

described by,

$$\begin{aligned}
T_n^2 &= \left(\frac{\sqrt{m_n^2 + p_{\text{cm}}^2} \cos \theta \sinh \chi + \cosh \chi \sqrt{p_{\text{cm}}^2 - (m_n \sin \theta \sinh \chi)^2}}{1 + \sin^2 \theta \sinh^2 \chi} \right)^2 - 2T_n m_n \\
s &= (m_\alpha + m_t)^2 + 2m_t T_\alpha \\
\chi &= \ln \left| \frac{p_{\text{cm}} + \sqrt{m_t^2 + p_{\text{cm}}^2}}{m_t} \right| \\
p_{\text{cm}} &= \sqrt{\frac{(s - m_\alpha^2 - m_t^2)^2 - 4(m_\alpha m_t)^2}{4s}} \\
p'_{\text{cm}} &= \sqrt{\frac{(s - m_n^2 - m_d^2)^2 - 4(m_n m_d)^2}{4s}}
\end{aligned} \tag{9.3}$$

where m_n is the neutron mass, m_α is the α mass, m_t is the mass of the target nucleus (predominantly ^{13}C), m_d is the mass of the daughter nucleus (predominantly ^{16}O), and T_α and T_n are the kinetic energies of the incident α and outgoing neutron, respectively. This equation is similar to the relationship discussed in Equation 6.4, as explained in [171], with the appropriate mass and energy terms substituted in.

In order to determine the angle of the outgoing neutrons relative to the normal of the gold-graphite interface, we recognize that Equation 9.4 sets the polar angle between the neutron and α momenta, but not the azimuthal angle, which can vary uniformly between 0 and 2π .

To determine the angle between the outgoing neutron's velocity and the normal of the gold-graphite interface, θ_n , we fix the polar coordinate as constrained by Equation 9.4 for each energy bin in the outgoing neutron spectrum and vary the azimuthal angle between 0 and 2π in steps of 0.1 to define a vector of unit magnitude pointing in the direction of the neutron's velocity relative to the incident α , $\hat{v}_{\alpha,n}$. We then rotate this vector by the angle of the α relative to the normal and calculate θ_n using the relation,

$$\cos \theta_n = \hat{v}_{\alpha,n} \cdot \hat{n}$$

where \hat{n} is the unit magnitude vector pointing in the direction normal to the gold-graphite interface. This allows us to determine the energy-angle distribution of neutrons emitted by the $^{241}\text{Am}^{13}\text{C}$ source.

This simulation and analysis was performed for three different source configurations: one in which no additional gold foil was added between the ^{241}Am source and the graphite pellet, one in which one layer of $1\ \mu\text{m}$ gold foil had been inserted between them, and one in which two layers of gold foil were used. Figure 9.10 shows the neutron spectra predicted for all three configurations. These spectra are normalized by the number of simulated α -decays and therefore account for the geometric effects of whether or not the α will make it into the graphite pellet, as well as the probability of the α undergoing the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction once it enters the target.

As can be seen in Figure 9.10, increasing the amount of foil between the ^{241}Am source and the graphite pellet decreases the total rate of neutron production. Each




Figure 9.10: Neutron energy spectra coming from the $^{241}\text{Am}^{13}\text{C}$ source, as simulated for source configurations with (*blue*) no gold foil between the ^{241}Am source and the graphite pellet, (*green*) $1\text{ }\mu\text{m}$ of gold foil between them, and (*purple*) $2\text{ }\mu\text{m}$ of foil between them.

Table 9.2: The mean and RMS neutron energy and total neutron production rate of the $^{241}\text{Am}^{13}\text{C}$ source with a $100\text{ }\mu\text{Ci}$ ^{241}Am source and three different thicknesses of gold foil between the ^{241}Am source and graphite pellet, as simulated using GEANT4 and TALYS.

Foil Thickness	$\langle T_n \rangle$ [MeV]	RMS T_n [MeV]	Rate [n/s]
$0\text{ }\mu\text{m}$	4.678	0.714	7.0
$1\text{ }\mu\text{m}$	4.455	0.707	4.4
$2\text{ }\mu\text{m}$	4.210	0.717	2.0

additional layer of foil also decreases the mean energy of the outgoing neutrons by $\sim 230\text{ keV}$, although the width of these distributions varies little. This information is summarized in Table 9.2.

The energy-angle distribution calculated for neutrons emitted by the $^{241}\text{Am}^{13}\text{C}$ source with $1\text{ }\mu\text{m}$ of foil between the ^{241}Am source and the graphite pellet is shown in Figure 9.11. The periodic vertical stripes in this figure are an artifact of the discretization of the azimuthal angle relative to the α momentum at the time of the (α, n) reaction. This figure shows that the neutron energy is generally higher when it is produced with momentum in the forward direction, though the change in neutron energy with the angle tends to be fairly small.

For the most part, the neutron energy-angle distribution shown here qualitatively agrees in energy scale and width with the spectrum predicted by Daya Bay for their




Figure 9.11: The energy-angle distribution of neutrons emitted by the $^{241}\text{Am}^{13}\text{C}$ source with $1\text{ }\mu\text{m}$ of gold foil between the ^{241}Am source and the graphite pellet target. The angle of the neutron θ_n is given relative to the line normal to the gold-graphite interface. The colors of each bin correspond to the probability of a neutron being generated with the energy and angle given by that bin per α -decay in the ^{241}Am source.

$^{241}\text{Am}^{13}\text{C}$ source [186]. However, the distribution shown in [186] predicts a large backscatter peak of lower energy neutrons produced with momenta in the backwards direction which we do not see in our simulations.

9.3.3 Data Analysis and Results

The $^{241}\text{Am}^{13}\text{C}$ source was inserted inside of a stainless steel holder inside of CALIS, which lowered the source into the LSV just outside of the TPC cryostat, as described in Section 9.1. Data was collected using this source from December 16, 2015 to January 8, 2016. Collected outer detector data were transferred to Fermi National Laboratory where it was reconstructed with DarkArtOD, as described in Chapter 8.2, while TPC data were reconstructed at LNGS with DarkArt and then transferred to the Fermilab computers. Outer detector and TPC data were then merged, matching corresponding events together for further coincidence analysis.

The primary goal of this analysis was to tune an optical response model, including a model describing ionization quenching of the liquid scintillator, describing the prompt and delayed coincide neutron-induced signals in the LSV following a single neutron scatter in the TPC. Such a model can then be applied to neutron background simulations to determine the neutron vetoing efficiency of a set of cuts made using the LSV data.

Therefore, the general philosophy underlying this analysis was to analyze the data in a way that can be reproduced in Monte Carlo as directly as possible. This way, we can reduce the amount of uncertainty in comparison of the Monte Carlo and data.

In addition to the analysis of this data, we also performed a GEANT4 Monte Carlo simulation of this calibration campaign. The $^{241}\text{Am}^{13}\text{C}$ source was implemented by generating neutrons with the energy-angle distribution shown in Figure 9.11. The source and the source holder were added to the G4DS framework with the help from B. Reinhold and H. Qian.

Events with single-scatter events were identified in the TPC by requiring that the TPC reconstruction software identify two physical pulses: one S1 pulse, and one S2 pulse. Some TPC events have a third pulse, generally referred to as S3, which is caused by secondary ionized electrons produced by the S2 signal on the TPC cathode scintillating in the gaseous argon. We found that neglecting these events did not result in a significant loss of statistics, and so events with an S3 signal were not included in this analysis.

Events in which a particle scatters multiple times in the TPC will typically only produce one S1 signal, unless they are separated far enough in time to produce two distinguishable signals, though this is unlikely given the small size of the TPC. However, if two recoils occur at different heights in the TPC, they will ionize electrons at two different heights as well. G. Koh and A. Fan have found that ionized electrons drift through the LAr at a speed $\sim 1\text{ mm}/\mu\text{s}$ under a drift field of 200 V/cm . Therefore, two recoils that happen at different distances from the TPC cathode will produce electrons that take different amounts of time to reach the gas pocket and produce an S2 signal. As a result, multiple-scatter events at different heights in the TPC will produce separate S2 pulses, resulting in three or more total pulses being identified by the TPC reconstruction software.

By studying GEANT4 simulations of γ -rays scattering in the TPC, G. Koh has found that the TPC reconstruction software can typically identify multiple scatters if they are separated in height by more than 4 mm . In our analysis of a GEANT4 simulation of the $^{241}\text{Am}^{13}\text{C}$ source, we therefore tagged an event as a multiple-scatter event if G4DS recorded energy depositions in the TPC separated by more than 4 mm in height. Due to the $\sim 1\text{ mm}/\mu\text{s}$ drift speed of the electrons, this also means that we should reject events with TPC energy depositions separated by more than $4\mu\text{s}$ as well.

Since the Monte Carlo did not include optics, simulated energy depositions are measured in keV rather than in PE. To smooth the translation between these two factors, we combine S1 and S2 to determine the energy of an event in keV_{ee} . As described in Chapter 2.3, energy deposited into the LAr can be dissipated as heat, scintillation photons, or ionized electrons. By scaling S1 by the amount of energy carried away by a single scintillation photon, and S2 by the number of S2 photons produced per ionized electron and the amount of energy carried away by each ionized electron, we can determine the total amount of energy not dissipated by heat. This




Figure 9.12: The reconstructed quenched recoil energy vs. f_{90} for single scatter recoils observed in the TPC during the $^{241}\text{Am}^{13}\text{C}$ calibration campaign. Events with high $f_{90} \gtrsim 0.84$ are primarily events with Cherenkov light in the Teflon or fused silica. The upper band (approximately $0.5 \lesssim f_{90} \lesssim 0.84$) is due to nuclear recoils from $^{241}\text{Am}^{13}\text{C}$ neutrons scattering in the LAr, while the lower band with $0.1 \lesssim f_{90} \lesssim 0.5$ is due to electron recoils due to backgrounds in the TPC, and the lowest band with $f_{90} \lesssim 0.1$ is due to S2 triggers and events near the extraction grid.

calculation is done using the relation,

$$E_r = \frac{W}{1 + \alpha} \left(\frac{S1^{\text{corr}}}{\epsilon_1} + \frac{S2^{\text{corr}}}{\epsilon_2} \right) \quad (9.4)$$

where E_r is the energy of the recoiling ^{40}Ar nucleus in keV_{ee} , $S1^{\text{corr}}$ and $S2^{\text{corr}}$ are the S1 and S2 signals after applying position-dependent corrections to account for inhomogeneities in the S1 and S2 signal strength, $W \approx 23.6 \text{ eV}$ is the mean energy required to produce an electron-ion pair [92], $\alpha \approx 0.21$ for electron recoils and $\alpha \approx 1$ for nuclear recoils [92] is the ratio of excitons to ions produced in the deposition, and ϵ_1 and ϵ_2 are the number of S1 and S2 photoelectrons, respectively, detected per exciton or ion. These last two quantities have been measured by M. Wada and L. Pagani. At 200 V/cm , they have measured $\epsilon_1 \approx 0.1572 \text{ S1 PE/exciton}$ and $\epsilon_2 \approx 27.47 \text{ S2 PE/ion}$. It should be noted that the analysis presented here has been found to be insensitive to small variations in these quantities.

By integrating S1 pulses, we can determine the fraction of S1 scintillation light that came within the first 90 ns of the pulse. This fraction is the f_{90} parameter discussed in Chapter 2.3, used for pulse shape discrimination to distinguish between electron and nuclear recoils.

Figure 9.12 shows the distribution of LAr recoil energy in keV_{ee} versus f_{90} . There are four features worth discussing in this plot. At high values of f_{90} (conservatively, we use $f_{90} \gtrsim 0.84$) we see low energy events caused by Cherenkov light produced in the TPC and PMT fused silica or the Teflon PMT reflector. Below that, we see a band roughly in the range of $0.5 \lesssim f_{90} \lesssim 0.84$, corresponding to nuclear recoils in the TPC. At $\lesssim 0.1f_{90} \lesssim 0.5$ we see a band corresponding to electron recoils in the TPC, and finally at very low f_{90} , we see a band primarily consisting of surface backgrounds and events that triggered on S2. The surface backgrounds seen in this band typically result from pulses that happen close to the extraction grid, where S1 and S2 are so close in time that they cannot be separated. Additionally, events may trigger on S2 rather than S1 if no S1 signal is produced in an event, as we observed to be the case for many surface backgrounds as well as a few other classes of events that are currently being investigated by C. Stanford, G. Koh, and X. Xiang. The spread of events between the electron and nuclear recoil bands at energies greater than $\sim 20 \text{ keV}_{ee}$ is due to events with coincident nuclear and electron recoils, as may likely result from a neutron inelastic scatter.

This energy variable is the total energy deposited that is not dissipated as heat. The fraction of recoil energy dissipated as heat is described by the Lindhard-Birks model discussed in Chapter 2.3 and in [92]. In order to estimate the total energy deposited in a single nuclear recoil, we therefore divide the quenched energy E_r by the Lindhard-Birks factor for a nuclear recoil of quenched energy E_r .

As can be seen in Figure 9.12, the electron and nuclear recoil bands bleed together at low energies. While the dark matter search analysis published in [12] used a more complicated analysis to avoid electron recoils leaking into the nuclear recoil f_{90} region, we simplified our analysis of the $^{241}\text{Am}^{13}\text{C}$ source by only considering events with reconstructed energies greater than 20 keV . This lower bound on the energy also excludes nearly all Cherenkov events. Above 20 keV , we see that nuclear recoils can be cleanly distinguished from electron recoils by considering events with $0.5 \lesssim f_{90} \lesssim 0.84$.

Since our simulation of the $^{241}\text{Am}^{13}\text{C}$ source does not include optics, the f_{90} parameter is not readily available. Instead, we approximate this cut by considering the fraction of quenched energy deposited by nuclear recoils. We note that, as can be seen in Figure 9.12, the nuclear recoil band at higher energies has a mean $f_{90} \sim 0.78$, while the electron recoil band has a mean $f_{90} \approx 0.29$. We approximate a cut requiring that $f_{90} > 0.5$ by requiring that nuclear recoils contribute to greater than 60% of the total quenched energy in the simulation. We find that the analysis presented here is relatively insensitive to the Monte Carlo analysis: similar results are seen requiring that nuclear recoils constitute more than 40% of the quenched energy, instead.

As can be seen in Figure 9.10, most neutrons are produced with kinetic energy below 5 MeV . The maximum energy, T_r^{max} , that can be transferred to a nucleus with mass number A by a single scatter from a neutron with kinetic energy T_n is given by,

$$\frac{T_r^{\text{max}}}{T_n} = \frac{4A}{(1+A)^2} \quad (9.5)$$




Figure 9.13: Recoil energy spectrum in the TPC during the $^{241}\text{Am}^{13}\text{C}$ calibration campaign between 20–500 keV, after correcting for Lindhard-Birks quenching, in coincidence with a neutron thermalization and capture signal.

This is consistent with the recoil energy spectrum we see in Figure 9.13. To further clean up the signal, we require that the total energy deposited in the TPC be less than 500 keV. While, as we can see from the nuclear recoil spectrum shown in Figure 9.13, there are relatively few events above this threshold, it helps remove events that have mixed nuclear recoil and high energy γ -ray signals, which might be harder to reconstruct in the Monte Carlo.

Lastly, we require that the LSV observe a cluster in prompt coincidence and a cluster in delayed coincidence with the TPC. Prompt coincidences are identified by looking for clusters that start within a ± 200 ns window relative to the TPC S1 time, as calibrated using ^{60}Co coincidences.

Delayed coincidences are identified by looking for the largest cluster in the LSV that follows the TPC S1 signal by more than $20 \mu\text{s}$. This delay helps remove events in which an afterpulse might register as a capture signal. We also require that the prompt cluster charge be less than 50,000 PE to remove events with muon signals in the LSV.

Early and Late Prompt Signals

The prompt signal in the LSV during the $^{241}\text{Am}^{13}\text{C}$ campaign can be subdivided into two categories: *early prompt* and *late prompt*. We define early prompt to be neutrons that scatter in the LSV before entering the TPC, and late prompt to be those that scatter in the TPC before doing so in the LSV. Often early prompt signals will result from neutrons scattering in the LSV, scattering in the TPC, and then going back into the LSV to thermalize and capture. Since it can take a couple hundred nanoseconds




Figure 9.14: LSV cluster start time vs. charge for clusters in prompt coincidence with a single nuclear recoil in the TPC passing all cuts mentioned in the text. The *early prompt* region can be seen left of -12.8 ns, while the late prompt region can be seen to the right.

to collect of the light produced by a large scintillation event in the LSV, we typically only see one prompt cluster in these cases, and so as a result, the early and late prompt signals may look slightly different.

Figure 9.14 shows the cluster start time versus charge distribution for these prompt events. Since both γ -rays produced by a ^{60}Co decay produce signals in the LSV and TPC at approximately the same time, we were able to use the ^{60}Co signal to determine the timing offset dividing early and late prompt signals. Doing so, we call clusters with a start time earlier than -12.8 ns relative to the TPC S1 start time *early prompt*, and those with start times after this time *late prompt*.

We see much more spread in the neutron energy in the late prompt region than we do in the early prompt region. This difference is largely because early prompt neutron clusters all start shortly after the neutron leaves the $^{241}\text{Am}^{13}\text{C}$ source, and the initial neutron scatter before it enters the TPC starts with the neutron at its full energy. In the late prompt region, neutrons may have lost some of their energy to the TPC and cryostat materials before producing any signal at all. In this region, we see that the cluster start time is anti-correlated with the cluster charge; this anti-correlation results from the fact that lower energy neutrons have longer times of flight and will also produce less scintillation light in the LSV.

Figure 9.15 shows the spectrum in the LSV for events in prompt coincidence with nuclear recoils in the TPC. This spectrum is dominated by neutron recoils on protons in the scintillator. The kink around 700–800 PE is due to the 1461 keV γ -ray from the decay of the first excited state of ^{40}Ar from inelastic neutron scatters. Much of the




Figure 9.15: The prompt LSV spectrum showing LSV clusters in prompt coincidence with nuclear recoils in the TPC. (*Black*) Shows all prompt signals, (*blue*) shows the spectrum for early prompt events, and (*green*) shows the spectrum for late prompt events.

spectrum above 1000 PE is due to other inelastic recoils, many of which are described in Table 5.1.

In Figure 9.14, we see a long tail extending to beyond 100 ns at low energies around 100 PE and lower. This tail can be seen in the low energy peak in the late prompt spectrum of Figure 9.15. This long tail is too far delayed to be explained by low energy neutrons' time of flight. This feature is likely due to inelastic neutron recoils on fluorine in the Teflon, exciting the ^{19}F to the metastable second excited state with excitation energy of 197 keV and a half-life of 89.3 ns.

Overall, we see a large peak around ~ 200 PE in the early prompt signal; this is the mean signal produced by the $^{241}\text{Am}^{13}\text{C}$ neutrons as they thermalize. We do not see such a pronounced peak in the late prompt signal. As mentioned earlier, this is because neutrons may lose energy while traveling from the TPC to the LSV, and so the neutron may have lost significant amounts of energy before its first scatter in the LSV.

Model Fitting

As mentioned earlier, the goal of this analysis is to tune a model of the LSV response to nuclear recoils. To do so, we considered Models I, II, and III, as discussed in Chapter 6.3.

For each of these models, quenching factors were computed for each simulated energy deposition by integrating the quenching factors over the kinetic energy of the scintillating particle before and after the deposition. Stopping powers for recoiling

nuclei were computed using SRIM [96], and stopping powers for electron recoils were obtained using the NIST ESTAR tool [122]. As we discussed in Chapter 8.4.4, the ESTAR tool only provides stopping power data down to 10 keV; in order to integrate stopping powers down to lower energies, we assumed the stopping power varies roughly linearly between 0–10 keV as has been shown to be the case in some materials [185].

To speed up processing, the quenching factor integrals are precomputed for each particle that may scintillate in the LSV by integrating the given quenching model from 0 to E keV, where E varies from 0–10 MeV. These integrals are used as a lookup table, and the integral over a single deposition is calculated by subtracting the value of the integral before and after the deposition.

We accounted for the position-dependence of LSV depositions by multiplying each deposition by a position factor, determined from the position model generated by H. Qian, shown in Figure 7.4. Optical simulations done by H. Qian indicate that CALIS and the $^{241}\text{Am}^{13}\text{C}$ source holder have little effect on the overall optics of the LSV while they are inserted.

Early and late prompt coincidence spectra are generated by summing over all hits in a ± 200 ns window surrounding the time of the first deposition in the TPC, classifying all depositions that happen more than 12.8 ns prior to the TPC start time as early prompt, and those starting after as late prompt. Varying the size of this window by small amounts has little effect on the final results.

We found that the late prompt signal in the simulation was missing the low energy feature around ~ 100 PE. Further investigation revealed that GEANT4 does not account for the decay time of the second excited state of ^{19}F ; instead, γ -rays from this excited state decay are produced instantaneously. We also found that this second excited state is the only one present in the GEANT4 databases. To correct for the first error, we identified all depositions in the LSV that result from γ -rays that were produced in the Teflon with 190–200 keV by a neutron that recoiled off of a ^{19}F nucleus. A random time offset was chosen from an exponential distribution with a decay time of 89.3 ns for each event and was added to the deposition time of all electron and γ -ray depositions in the scintillator resulting from the identified γ -ray track.

Since the first excited state of ^{19}F has a lower cross section for being produced by an inelastic neutron scatter, we did not correct for it being missing in the GEANT4 libraries.

The delayed signal was identified by summing over depositions, multiplied by the appropriate quenching and position factors, between the end of the prompt window and the end of the acquisition window.

The prompt and delayed spectra were then convolved with a Gaussian response function accounting for the light yield and the SPE and Poisson variance of the signal, converting the spectrum into photoelectrons.

To compare the simulated spectra with data, we allowed each spectrum to have an overall scale factor accounting for the rate, which we treated as a fit parameter. Since the quenching model and light yield may both affect the overall amount of light collected during a scintillation event, we defined a light yield correction factor that we used to scale the x-axis to make the simulated 2.2 MeV peaks from neutrons




Figure 9.16: $^{241}\text{Am}^{13}\text{C}$ prompt coincidence spectrum in the LSV. (*Black*) Data taken during the calibration campaign. This spectrum is dominated by the neutron thermalization signal, which primarily consists of proton recoils; the kink around 700 PE is from the γ -ray produced by the first excited state of ^{40}Ar , and much of the tail at higher energies results from γ -rays produced by other inelastic scatters. (*Blue*) The best fit Monte Carlo simulation modeling the liquid scintillator quenching with Model II with $kB = 0$ and $C = 3.5 \times 10^{-5} (\text{cm/MeV})^2$. This fit has $\chi^2/\text{NDF}=1.7$.

capturing on ^1H line up with those in the data. Because this peak is made of relatively high energy electron recoils, it is quenched very little, and so these correction factors were typically around a couple percent. If the correction factors were larger than 1%, we reprocessed the simulation data using the corrected light yield. The observed prompt and delayed coincidence spectra with the best fit simulated spectra are shown in Figures 9.16 and 9.17, respectively.

We performed this fit with 14 values of kB ranging from 0–0.026 cm/MeV and 15 values of C ranging from $(0.3\text{--}4.5) \times 10^{-5} (\text{cm/MeV})^2$. The χ^2/NDF was computed for each set of parameters for both the the prompt and delayed spectra. We also determined the combined χ^2/NDF by summing the χ^2 for both fits for each set of parameters and dividing by the total number of degrees of freedom. These three distributions can be seen in Figure 9.18, where we have used Delaunay triangulation to interpolate between the sampled points and draw contours in the parameter space χ^2 map.

The best fit parameters were identified by finding the minima in Figure 9.18, and the uncertainty in the parameters was estimated by finding the values for each parameter that increased χ^2/NDF by one while the other parameter was held constant at its best fit value.




Figure 9.17: $^{241}\text{Am}^{13}\text{C}$ delayed coincidence spectrum in the LSV, with cuts to remove afterpulses from the prompt coincidence signal. (*Black*) Data taken during the calibration campaign, primarily resulting from the neutron capture signal. The peak at ~ 20 PE is from the signal produced by the $\alpha + {}^7\text{Li}$ produced by the neutron capture reaction on ^{10}B ; the peak ~ 240 PE is from the $\alpha + {}^7\text{Li} + \gamma\text{-ray}(478\text{ keV})$ produced by the neutron capture on ^{10}B . (*Blue*) The best fit Monte Carlo simulation modeling the liquid scintillator quenching with Model II with $kB = 0.022$ and $C = 5 \times 10^{-6} (\text{cm/MeV})^2$. This fit has $\chi^2/\text{NDF} = 1.5$.

Table 9.3: Best fit parameters for Model II and their χ^2/NDF .

Fit	kB [cm/MeV]	C [(cm/MeV) ²]	χ^2/NDF
Prompt	$0_{-0}^{+0.003}$	$3.5_{-0.8}^{+1.1} \times 10^{-5}$	1.7
Delayed	$0.022_{-0.003}^{+0.004}$	$0.5_{-0.3}^{+0.4} \times 10^{-5}$	1.5
Combined	$0.010_{-0.003}^{+0.004}$	$1.4_{-0.3}^{+0.6} \times 10^{-5}$	4.8




Figure 9.18: χ^2/NDF versus kB and C values in Quenching Model II for (top) the prompt $^{241}\text{Am}^{13}\text{C}$ signal, (middle) the delayed $^{241}\text{Am}^{13}\text{C}$ signal, and (bottom) both signals combined. These three fits have χ^2/NDF of 1.7, 1.5, and 4.8, respectively.

As can be seen in Figure 9.18, we were able to separately obtain good fits for the prompt and delayed spectra, with χ^2/NDF equal to 1.7 and 1.5, respectively. However, the parameters that optimize these fits are very different; if we require that our quenching model have the same parameters for all particles, the parameters that best fit the combination of the prompt and delayed spectra give us $\chi^2/\text{NDF}=4.8$. The best fit parameters for these three cases are summarized in Table 9.3. These parameters are consistent with a light yield of $0.63 \pm 0.03 \text{ PE/keV}_{ee}$. This light yield is slightly higher than the light yields reported in Chapter 8.4.4; this difference is due to a combination of electron recoil quenching and the position dependence, which we account for in this analysis. In particular, the spatially-averaged light yield near the cryostat is lower than this value, since the cryostat can block scintillation light.

As can be seen in Table 9.3, the proton quenching model is similar to what we saw in the prototype measurements reported in Chapter 6.3. Models I and III both produced proton recoil peaks that were too narrow and were unable to explain the breadth of the observed peak, favoring Model II as a description of the liquid scintillator response. However, if we try to extrapolate the parameters obtained by fitting the prompt spectrum, we find that the α and ${}^7\text{Li}$ from the neutron capture signal are too heavily suppressed, indicating that this model is not a complete description of the scintillator response.

In lieu of a single completely general model of liquid scintillator quenching, we build an effective model, using different parameters for different species of scintillating particles. We use the parameters obtained from the prompt spectrum to describe the quenching of proton recoils, while we use the parameters obtained from the delayed spectrum to describe heavier recoiling nuclei and electron recoils. We note that we have not verified this model's fidelity at reproducing signals from heavier energy nuclei and that we only have measurements of α and ${}^7\text{Li}$ recoils at the energies they pick up from the neutron capture reaction. However, as we can see in Figure 9.17, this reproduces the capture signal very well, which is the main reason we need it. Heavier nuclei that may recoil in the scintillator, like carbon, oxygen, and boron, are very heavily quenched and the maximum kinetic energy a neutron can lose to them in a single recoil is kinematically suppressed; these nuclei therefore contribute fairly little to the thermalization signal produced as the neutron slows down, and so uncertainty in their quenching factors has little effect on the total thermalization signal.

Figure 9.19 shows the quenching factor for several different types of nuclei depositing all of their energy into the scintillator according to the model presented here. Integrating Model II from the particle's initial energy down to 0 gives the total quenched energy; the quenching factors shown in this figure are calculated by dividing the quenched energy by the initial energy of each nucleus. When we apply this model to Monte Carlo simulations, we often want to calculate the quenched energy for individual recoils. To do so, we simply look up the quenching factor of the scintillating particle for its kinetic energy before the recoil, multiply its kinetic energy by the quenching factor, and subtract the quenching factor times the kinetic energy after the recoil.




Figure 9.19: Quenching factors for (*black*) electrons, (*blue*) hydrogen nuclei, (*light green*) helium nuclei, (*purple*) lithium nuclei, (*orange*) beryllium nuclei, (*teal*) boron nuclei, (*red*) carbon nuclei, (*pink*) nitrogen nuclei, and (*dark green*) oxygen nuclei depositing all of their energy in the LSV.

This model, which we have tuned on the $^{241}\text{Am}^{13}\text{C}$ calibration data, allows us to simulate particles interacting with the detector and reproduce their signal in the LSV with decent accuracy.

Chapter 10

Background Simulations and Neutron Vetoing Efficiency

We have now characterized and described the backgrounds present in DarkSide-50, presented a veto system to measure and remove these backgrounds, and provided a detailed description of the operation and performance of this veto system through a combination of calibration measurements and simulations.

The next step is to assess the efficiency of the neutron veto and determine whether or not it will allow an experiment such as DarkSide-50 to run background free for an extended dark matter campaign. To accomplish this goal, we simulate the backgrounds in GEANT4, using the G4DS package developed by D. Franco and our collaborators at the *Institut National de Physique Nucléaire et de Physique des Particules* in Paris.

In this chapter, we will discuss simulations and estimations of the background rates in DarkSide-50. While we will discuss backgrounds from γ -rays and cosmogenic neutrons, most of the focus will be on radiogenic neutrons, drawing heavily on our discussion of the rates and spectra discussed in Chapter [4.4](#).

10.1 Cosmogenic Neutron Backgrounds

Preliminary work exploring the cosmogenic neutron background was done by A. Empl et al. in [\[28\]](#). In this work, they performed FLUKA simulations of a field of muons passing through Hall C of LNGS, with the DarkSide-50 detector geometry built into the simulation. This simulation used the Phase-I LSV scintillator composition, and included Cherenkov light production in the WCV, but no optics in the LSV.

A number of muons equivalent to what we would expect from 34 years of running was simulated. Cosmogenic events in which at least one neutron deposited 0–1 MeV in the TPC were recorded, and all of these events were accompanied by a signal of at least 3 PE in the WCV or 1 MeV deposited in the LSV.

While this work provides compelling evidence that the WCV and LSV can reject all of the cosmogenic neutrons DarkSide-50 is exposed to during a three year campaign, a number of details make these results difficult to apply directly to the

DarkSide-50 experiment. The Phase-II LSV scintillator cocktail is different than the Phase-I cocktail used in this simulation. While the change in scintillator composition is likely small enough to not make a significant difference in the cosmogenic background, the simulation analysis does not accurately reflect the cuts made in DarkSide-50 to remove cosmogenic neutron backgrounds. In particular, the 3 PMT threshold in the WCV and the 1 MeV threshold in the LSV do not reflect the cuts made in the DarkSide-50 analysis.

More detailed simulations and analyses are currently being carried out by P. Singh to determine how the current DarkSide-50 veto conditions affect these results, so we can more accurately determine the efficiency of the outer detectors for vetoing cosmogenic neutrons.

10.2 (alpha,n) Neutron Backgrounds

As we saw in Chapter 4.4, the dominant sources of radiogenic neutrons are in the borosilicate glass stems of the TPC PMTs, the outer cryostat flange (especially the Viton o-ring), the inner cryostat flange, and the cryostat walls.

In order to determine the radiogenic neutron backgrounds present in DarkSide-50, we simulated ~ 100 million neutrons in each of these materials. The simulated neutrons were generated with energies ranging uniformly between 0–10 MeV, which covers nearly the entire radiogenic energy spectrum. These energy spectra can be seen in Figures 4.20–4.24.

10.2.1 TPC Event Reconstruction

Since this analysis requires detailed analysis of both the cuts made by both the LSV and the TPC, the TPC reconstruction had to be more detailed than we did for the $^{241}\text{Am}^{13}\text{C}$ simulation analysis. TPC recoils were identified by looking for energy depositions in the active liquid LAr constrained within a radius r of $0 \leq r \leq 17.77$ cm and a z -coordinate in the range $-21.44 \leq z \leq 14.31$ cm.

The physics underlying the LAr scintillation model is described in Chapter 2.3. When energy is deposited in the LAr, some number of ions and some number of excitons will be formed; we define the exciton-to-ion ratio as α . We assume that $\alpha = 0.21$ for electron recoils and $\alpha = 1$ for nuclear recoils, as measured by SCENE [92]. We also applied a Lindhard-Birks quenching factor to suppress the light yield of nuclear recoils, following the model measured by SCENE and described by Equations 2.4 and 2.5.

As illustrated in Figure 2.8, some of the ionized electrons produced by the energy deposition will recombine with argon ions. This recombination probability is described by the DSLight model tuned at 200 V/cm, designed and calibrated by D. Franco and P. Agnes. This model gives the recombination $R(T)$ as,

$$R(T) = p_0(1 - p_1 e^{p_2 T}) e^{p_3 T^{p_4}} + p_5 \quad (10.1)$$

Table 10.1: Recombination probability model parameters used in the model tuned by D. Franco and P. Agnes on electron recoil data.

p_0	p_1	p_2	p_3	p_4	p_5
0.296766	3.95496	-0.517812	-0.0138485	0.912436	0.661046

where T is the kinetic energy in keV of the recoiling nucleus and p_i are fit parameters. This model was tuned on electron recoil data in DarkSide-50 and DarkSide-10. These parameters are given in Table 10.1.

Since this model was tuned on electron recoils, it likely loses accuracy for nuclear recoils, which tend to have higher ionization densities and higher recombination probabilities. However, since we do not yet have a model that has been tuned separately on nuclear recoils, we use this model for the present analysis. Since we primarily use the recombination probability as a correction to the S1 and S2 light yields in the TPC, uncertainties in the final efficiency calculation due to uncertainties in the recombination model are likely small. This assumption is supported by the observation that small changes in the TPC cuts that depend on the recombination probability have small effects on the change in neutron veto efficiency.

We define the total number of quanta to be the total number of scintillation photons and ionization electrons produced by an energy deposition in the LAr. We calculate the number of quanta by multiplying the deposited energy by the Lindhard-Birks factor and dividing the energy by the mean energy of a quantum, which we take to be 19.5 eV [92].

We assume that the number of quanta produced in a scintillation event varies according to a Poisson distribution, and divide these quanta into excitons and ions according to a binomial distribution according to the exciton-to-ion ratio for the type of recoil. We assume that some fraction of the ions recombine to produce scintillation photons; this fraction is given by the recombination probability.

To get the quenched S1 energy, we multiply the number of photons by the mean quantum energy; to get the number of S1 photoelectrons, we multiply this number of the TPC light yield, which G. Koh has measured to be 8.8 PE/keV. We determine S2 by multiplying the number of electrons that do not recombine by the S2 yield, which has been measured to be ~ 24 PE/electron in DarkSide-50 by M. Wada.

We define the S1 start time as the time of the first deposition in the TPC. We define the S1-S2 separation time of each deposition by dividing the distance of each deposition from the gaseous argon layer by the electron drift velocity, which G. Koh has measured to be ~ 1 mm/ μ s. We then add the time difference between the deposition time and the S1 start time to the drift time. We define the S2 separation time as the difference between the maximum and minimum S1-S2 separation times for an event. G. Koh has found that the number of detected pulses can accurately be reproduced in Monte Carlo if we assume the S2-separation resolution is ~ 4 μ s. We therefore tag events with separation greater than this value as multiple scatter events.




Figure 10.1: Fraction of S1 photoelectrons that come from nuclear recoils for neutrons simulated between 0–10 MeV in the TPC PMT stems. For this analysis, nuclear recoils were defined as those with an S1 nuclear fraction above 0.8.

In lieu of simulating the time profile of a scintillation event in the LAr, we approximate pulse shape discrimination cuts by considering the fraction of S1 photons that come from nuclear recoils. The distribution of this variable is shown in Figure 10.1. For this analysis, we define nuclear recoils to be those for which the fraction of S1 resulting from nuclear recoils is greater than 0.8. As discussed in [12], the S1 nuclear recoil acceptance is approximately constant above 95 PE, where $f_{90} \sim 0.75$ for nuclear recoils and $f_{90} \sim 0.3$ for electron recoils, and the pulse shape discrimination cut requires $f_{90} \gtrsim 0.7$. This is comparable to the cut we make in the simulation data; however, Figure 10.1 shows that small changes in this cut value have relatively little effect on the number of events that pass it.

In the DarkSide-50 analysis presented in [12], a maximum energy of 460 PE was placed on the WIMP-search window. We therefore rejected all events with an S1 signal greater than this value. Furthermore, while the WIMP-search window extended down to 0 PE, the pulse shape discrimination cut has approximately 0% nuclear recoil acceptance if $S1 \lesssim 40$ PE. Above ~ 95 PE, the acceptance is approximately constant at $\sim 90\%$. Between these energy scales, the acceptance varies approximately linearly.

To account for the changing nuclear recoil acceptance in this analysis, we assume that the probability of a nuclear recoil passing the energy cuts is 0 if $S1 < 40$ PE, the probability is 1 if $S1 > 95$ PE, and varies linearly between these two values. For events with S1 between 40 and 95 PE, we draw a number from a uniform random distribution and say that the event passes the energy cut with the interpolated probability.

As mentioned in Chapter 9.3, there is a bug in GEANT4 that causes γ -rays from the second excited state of ^{19}F to be produced immediately, rather than with the

appropriate decay constant. We applied the same correction described in this chapter here.

10.2.2 LSV Event Reconstruction

The LSV event reconstruction performed here is very similar to the reconstruction discussed for the $^{241}\text{Am}^{13}\text{C}$ simulation analysis discussed in Chapter 9.3. We use the same light yield position correction map shown in Figure 7.4, and we quench recoils in the liquid scintillator according to the model shown in Figure 9.19 with a light yield of 0.63 PE/keV_{ee} .

We account for Poisson variance and SPE variance in the light yield by drawing a random number from a Gaussian distribution with a mean given by the product of the quenched energy and the light yield, and variance accounting for the average PMT SPE variance of 14% and Poisson statistics, as explained in Equation 6.4.

For each event, we subtract the TPC S1 start time from the time of each deposition in the LSV, adding an additional delay for γ -rays from the second excited state of ^{19}F as appropriate. All deposition between 50 ns before the S1 start time to 250 ns after the S1 start time are recorded as part of the prompt signal. Additionally, we identify the 500 ns window before the S1 start time with the most photoelectrons for the preprompt signal, and the 500 ns window after the S1 start time with the most photoelectrons for the delayed signal.

To simulate the analysis reported in [12], we produced a veto signal for events with a simulated prompt signal greater than 1 PE, a preprompt signal greater than 3 PE, or a delayed signal greater than 6 PE.

This analysis does not account for the time it takes photoelectrons to be detected after the deposition times. To estimate this effect, we modeled the amount of time it takes to detect the first photoelectron when N PE were detected in total. The LSV has a photocathode coverage of $p = \sim 6.98\%$, meaning that a single photon that does not get absorbed by the scintillator has a probability of p of hitting a photocathode rather than the reflector. Neglecting effects from absorption by the scintillator, the probability of a single photoelectron having been detected after k additional reflections can be approximated by a binomial distribution where the probability per photoelectron is p . The probability of the first photoelectron being detected after k reflections is described by a geometry distribution, where the probability of at least one photoelectron being detected is described by the aforementioned binomial distribution. The distribution on the number of reflections undergone by the first photon to be detected $P(k)$ is then approximately,

$$P(k) = (1 - p)^{Nk} - (1 - p)^{N(k+1)} \quad (10.2)$$

Assuming that the average distance that a photon travels between reflections is the average chord length of a sphere, we can estimate the average time between reflections by dividing this distance by the speed of light in the scintillator. By drawing k reflections from the distribution given in Equation 10.2, we can estimate the amount of time it takes the first photoelectron from the deposition to be detected.




Figure 10.2: Probability of neutrons from the PMT stems surviving cumulative TPC cuts as a function of neutron energy. Each curve represents the probability of neutrons surviving each given cut as well as the cuts above it. (*Black*) All neutrons simulated at each energy bin (normalized to 1 for each bin), (*blue*) neutrons depositing any nonzero amount of energy in the active LAr, (*green*) those neutrons that produce nuclear recoils and pass the pulse shape discrimination cut, (*purple*) those neutrons that also pass the multiple scatter cut, (*orange*) those neutrons that pass fiducial z cuts, and (*teal*) those neutrons that also leave signals in the WIMP energy window.

Overall, we expect this estimation to be conservative; since photons can be absorbed by the scintillator, we expect the distribution of k to be biased towards lower values where the photons are less likely to be absorbed.

While we found that implementing this delay slightly changed the efficiencies of the prompt and delayed signal, this change did not have an overall effect on the simulated vetoing efficiency. This is most likely because photons produced by prompt depositions that are delayed outside of the prompt window will still be detected as delayed coincidence signals. We therefore decided to neglect this effect for the rest of this analysis to speed up processing time.

10.2.3 Results

The first goal in analyzing the simulation data was to determine the probability of a neutron surviving all TPC cuts and all TPC and LSV cuts as a function of the energy with which the neutron was generated in each of the dominant (α , n) sources in DarkSide-50, as identified in Chapter 4.4.3. These sources include the TPC PMT stems, the inner and outer cryostat flanges, and the cryostat walls.




Figure 10.3: Probability of neutrons from the inner cryostat flange surviving cumulative TPC cuts as a function of neutron energy. Each curve represents the probability of neutrons surviving each given cut as well as the cuts above it. (*Black*) All neutrons simulated at each energy bin (normalized to 1 for each bin), (*blue*) neutrons depositing any nonzero amount of energy in the active LAr, (*green*) those neutrons that produce nuclear recoils and pass the pulse shape discrimination cut, (*purple*) those neutrons that also pass the multiple scatter cut, (*orange*) those neutrons that pass fiducial z cuts, and (*teal*) those neutrons that also leave signals in the WIMP energy window.

Figures 10.2–10.5 show the probability of a neutron generated with a given energy between 0–10 MeV passing each set of TPC cuts, as simulated in this analysis. As can be seen in these figures, a great reduction in the neutron background comes from geometric factors alone, as $\lesssim 10\%$ of neutrons of any energy produce any signal at all in the TPC. While this factor does not vary drastically with neutron energy, we find that the probability of a neutron leaving a signal in the TPC increases slowly with the neutron’s initial energy.

Adding a cut on the fraction of S1 light deposited in the TPC by nuclear recoils—our proxy for pulse shape discrimination—further decreases the neutron event rate. This cut tends to affect higher energy neutrons more strongly than it does lower energy neutrons, likely because higher energy neutrons are more likely to excite inelastic states in the nuclei of the detector, which may then produce γ -rays that can mix electron and nuclear recoils within an event. At low energies, this cut further reduces the neutron background rate by a factor of ~ 2.5 , while at higher energies it decreases the rate closer to a factor of ~ 6 .

The multiple scatter cut made on the S2-separation time decreases the neutron background relatively uniformly in energy by a factor ~ 2 . We see a similar reduction




Figure 10.4: Probability of neutrons from the outer cryostat flange surviving cumulative TPC cuts as a function of neutron energy. Each curve represents the probability of neutrons surviving each given cut as well as the cuts above it. (*Black*) All neutrons simulated at each energy bin (normalized to 1 for each bin), (*blue*) neutrons depositing any nonzero amount of energy in the active LAr, (*green*) those neutrons that produce nuclear recoils and pass the pulse shape discrimination cut, (*purple*) those neutrons that also pass the multiple scatter cut, (*orange*) those neutrons that pass fiducial z cuts, and (*teal*) those neutrons that also leave signals in the WIMP energy window.

factor in the fiducial z cut, which removed events in the top and bottom 3.63 cm of the active LAr volume.

Finally, the energy cut reduces the neutron background significantly for neutrons below ~ 1000 keV, where kinematic constraints make neutron recoils in the WIMP energy window rare. At higher neutron energies, we see this cut approximately flattens out, reducing the neutron background by a factor of ~ 3 .

In Figures 10.6–10.9, we see the effects of each of the LSV cuts acting on neutrons that pass all of the TPC cuts. These plots show events that pass all TPC cuts as a reference point, then the effects of individually adding each LSV cut, and finally the probability of a neutron of a given energy passing all TPC and all LSV cuts. Neutrons that pass all of these cuts are potential nuclear recoil backgrounds that may mimic a WIMP event in DarkSide-50.

At energies below ~ 2500 keV, we see relatively little effect from the preprompt cut. This is to be expected because most neutrons that scatter in the TPC do so before scattering in the LSV. However, at higher initial neutron energies, we find that the preprompt cut is more likely to reject potential background events. While some of these events are due to neutrons that recoil off of a nucleus in the liquid scintillator




Figure 10.5: Probability of neutrons from the cryostat walls surviving cumulative TPC cuts as a function of neutron energy. Each curve represents the probability of neutrons surviving each given cut as well as the cuts above it. (*Black*) All neutrons simulated at each energy bin (normalized to 1 for each bin), (*blue*) neutrons depositing any nonzero amount of energy in the active LAr, (*green*) those neutrons that produce nuclear recoils and pass the pulse shape discrimination cut, (*purple*) those neutrons that also pass the multiple scatter cut, (*orange*) those neutrons that pass fiducial z cuts, and (*teal*) those neutrons that also leave signals in the WIMP energy window.

and then enter the TPC, most of these events result from high energy neutrons that excite nuclei in the TPC or cryostat and produce inelastic scatter γ -rays that may interact with the LSV before the neutron scatters in the TPC. Overall, however, the effects of this cut on (α, n) neutrons is small; this cut was primarily intended to remove cosmogenic or delayed neutron events in which an LSV signal may precede the TPC signal.

We see a much more dramatic effect for the prompt cut, which vetoes events based on their thermalization signal or by any γ -rays that may be created by neutron inelastic scatters. The effectiveness of this cut does not vary dramatically with the initial energy of the neutron. At most neutron energies, this cut reduces the neutron background by a factor of ~ 10 .

The delayed cut gives us an even more dramatic reduction in the neutron background. The effectiveness of this cut is mostly independent of the initial neutron energy, since the thermal neutron capture signal is independent of the initial kinetic energy of the neutron. However, we do see a kink around 3.5 MeV where this cut's effectiveness decreases slightly. This small decrease may be due to fast neutron capture reaction, such as those on ^{56}Fe and ^{19}F that become energetically accessible for neu-




Figure 10.6: Probability of neutrons from the PMT stems surviving LSV cuts as a function of neutron energy. (*Black*) Neutrons passing all TPC cuts, (*blue*) neutrons passing all TPC cuts and the LSV prompt cut, (*green*) neutrons passing all TPC cuts and the LSV preprompt cut, (*purple*) neutrons passing all TPC cuts and the LSV delayed cut, and (*orange*) neutrons passing all TPC cuts and all three LSV cuts.

trons in this energy range. Nevertheless, this cut remains a strong tool for rejecting neutron backgrounds.

Overall, we see that the LSV cuts together can suppress neutron backgrounds by a factor of 100–1000, depending on the initial energy of the neutron and the material in which it was created. Since these cuts are not independent of each other, their rejection efficiencies do not compound independently. However, Figures 10.6–10.9 show us how they complement each other. Ultimately, we see that the delayed cut is the most powerful cut. However, at lower energies below ~ 3.5 MeV, we find that the prompt signal significantly enhances the vetoing efficiency. At higher energies, we tend to see a much smaller gain in rejection power from adding the prompt cut to the delayed cut. This smaller gain is likely because fast neutrons that capture on the detector materials and under go a (n, α) or (n, p) reaction will generally produce neither a prompt nor a delayed signal in the LSV.

Ultimately, Figures 10.6–10.9 tell us the probability of a neutron generated in a given material passing all TPC and LSV cuts as a function of the neutron’s initial energy. As discussed in Chapter 4.4.3, we have calculated the (α, n) yields and neutron spectra expected from each of these materials based on the measurements we have made of the radioactive contamination in them. Figures 4.20–4.24 tell us the number of neutrons we expect from each of these sources at a given energy per year. Since these neutron spectra and the cut efficiency histograms are both binned




Figure 10.7: Probability of neutrons from the inner cryostat flange surviving LSV cuts as a function of neutron energy. (*Black*) Neutrons passing all TPC cuts, (*blue*) neutrons passing all TPC cuts and the LSV prompt cut, (*green*) neutrons passing all TPC cuts and the LSV preprompt cut, (*purple*) neutrons passing all TPC cuts and the LSV delayed cut, and (*orange*) neutrons passing all TPC cuts and all three LSV cuts.

in 100 keV bins, we can combine both sets of plots to determine the initial energies of neutrons that pass all analysis cuts, and ultimately determine the number of neutron background events we expect to see over the course of a year.

This computation can be described by,

$$S'_i = S_i \times p_i \quad (10.3)$$

where S'_i the number of neutrons in the i^{th} energy bin that pass all cuts per year, S_i is the number of neutrons in the i^{th} energy bin that are produced in the given material per year, and p_i is the probability of a neutron whose initial energy falls in the i^{th} energy bin passing all cuts. The total number of neutrons that pass all cuts over the course of a year can then be calculated by $\sum_i S'_i$.

In this analysis, we consider the 6 dominant neutron sources discussed in Chapter 4.4.3: the PMT stems, the PMT hardware, the inner flange of the TPC cryostat (including the hardware in the flange), the outer flange of the cryostat (including the hardware and the Viton o-ring in the flange), the walls of the inner and outer cryostats, and the multilayer insulation between the walls of the inner and outer cryostats.

Since the PMT stems and hardware are in approximately the same location, we use Figure 10.6 for the survival probabilities in both of these materials. We also assume that the survival probability for neutrons originating in the multilayer insula-




Figure 10.8: Probability of neutrons from the outer cryostat flange surviving LSV cuts as a function of neutron energy. (*Black*) Neutrons passing all TPC cuts, (*blue*) neutrons passing all TPC cuts and the LSV prompt cut, (*green*) neutrons passing all TPC cuts and the LSV preprompt cut, (*purple*) neutrons passing all TPC cuts and the LSV delayed cut, and (*orange*) neutrons passing all TPC cuts and all three LSV cuts.

tion is between those for neutrons originating in the inner and outer cryostats. Since Figure 10.9 includes neutrons generated in both cryostats, we use these probabilities for neutrons generated in the multilayer insulation as well. Finally, we assume neutrons generated in the inner flange have survival probabilities as given by Figure 10.7, and neutrons generated in the outer flange have survival probabilities given by Figure 10.8.

As we can see in Figure 4.24, the neutron spectrum for the PMT hardware extends out to higher energies we simulated here. This is due to the high energy neutrons generated by the ${}^9\text{Be}(\alpha, n){}^{12}\text{C}$ reaction on the ${}^9\text{Be}$ in some of the electronics components. Since we observed fairly little variation in neutron vetoing efficiency with the neutron energy at these higher energies, we assumed that all neutrons with energies above 10 MeV are vetoed as efficiently as 10 MeV neutrons.

Figures 10.10–10.15 show the number of neutrons per year that we expect to be generated in each of these materials, as well as the number that we expect to pass the TPC cuts, the TPC cuts along with each of the individual LSV cuts, and all of the TPC and LSV cuts. The total rates from these plots are summarized in Table 10.2.

As we can see from these figures, the dominant source of background neutrons comes from neutrons generated in the PMT stems. This is largely due to the high levels of the upper ${}^{238}\text{U}$ chain present in the borosilicate glass in the stems, as discussed




Figure 10.9: Probability of neutrons from the cryostat walls surviving LSV cuts as a function of neutron energy. (*Black*) Neutrons passing all TPC cuts, (*blue*) neutrons passing all TPC cuts and the LSV prompt cut, (*green*) neutrons passing all TPC cuts and the LSV preprompt cut, (*purple*) neutrons passing all TPC cuts and the LSV delayed cut, and (*orange*) neutrons passing all TPC cuts and all three LSV cuts.




Figure 10.10: Initial energy of neutrons generated in the PMT stems for (*black*) all neutrons, (*teal*) neutrons passing all TPC cuts, (*green*) neutrons passing all TPC cuts and the LSV preprompt cut, (*blue*) neutrons passing all TPC cuts and the LSV prompt cut, (*purple*) neutrons passing all TPC cuts and the LSV delayed cut, and (*orange*) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.




Figure 10.11: Initial energy of neutrons generated in the PMT hardware for (*black*) all neutrons, (*teal*) neutrons passing all TPC cuts, (*green*) neutrons passing all TPC cuts and the LSV preprompt cut, (*blue*) neutrons passing all TPC cuts and the LSV prompt cut, (*purple*) neutrons passing all TPC cuts and the LSV delayed cut, and (*orange*) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.

in Chapter 3.2, and the relatively high (α, n) yield of the boron in it, as discussed in Chapter 4.4.3.

The next largest contribution comes from the flange on the outer cryostat; the fluorine in the Viton o-ring gives it a very high (α, n) yield which combines with the relatively high levels of radioactivity measured in the o-ring to make it a hot neutron source. However, as we can see in Figure 10.4, these background rate from these neutrons is naturally suppressed by the distance from the flange to the TPC.

The third largest source of neutron backgrounds comes from the PMT hardware. The proximity of the neutrons to the TPC makes them relatively more likely to leave a signal in the TPC to begin with, and the high neutron yield of ^9Be makes the electronics produce a comparatively large number of neutrons, despite their low mass. It is worth noting that the (α, n) yield calculations for some of the electronics components conservatively assumed that they were entirely made of beryllium; in reality, the components are likely made of some combination of beryllium and metals such as copper and zinc. However, since we did not have access to the full chemical composition of most of these materials, we conservatively assumed that they are primarily beryllium. This indicates that the estimate for the neutron background from the PMT hardware is likely conservative.

Compared to the three aforementioned sources, we found that the background rate of neutrons coming from the cryostat walls, multilayer insulation, and inner cryostat




Figure 10.12: Initial energy of neutrons generated in the inner cryostat flange for (black) all neutrons, (teal) neutrons passing all TPC cuts, (green) neutrons passing all TPC cuts and the LSV preprompt cut, (blue) neutrons passing all TPC cuts and the LSV prompt cut, (purple) neutrons passing all TPC cuts and the LSV delayed cut, and (orange) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.

flange were negligible. Despite the high mass of the cryostat stainless steel, they are clean enough to not pose a significant background risk to the experiment.

As can be seen in Table 4.13, there is a $\sim 7.4\%$ uncertainty in the neutron production rate of the PMT stems, a $\sim 4.1\%$ uncertainty in the neutron production rate of the cryostat flanges, and a $\sim 8.6\%$ uncertainty in the neutron production rate of the PMT hardware. This source of uncertainty is due to the uncertainty in the radioisotope contamination and component masses.

As discussed in Chapter 4.4.3, we estimate that a 1σ confidence interval around the predicted (α, n) yields extends 2% above the predicted value and 61.5% below this value.

Furthermore, varying the S2-separation and pulse shape discrimination thresholds in the TPC analysis of the Monte Carlo has no significant effect on the number of neutrons that pass all cuts. We found that increasing the S1 energy threshold or changing the turning point in the lower bound of the full S1 acceptance by ± 20 PE varied the number of events passing all cuts by a 2–20%; however, since the energy reconstruction model has been confirmed by SCENE and DarkSide-50 data, we do not believe this contributes significantly to the total uncertainty. Additionally, varying the prompt cut threshold by 2 PE changed the number of neutrons surviving all cuts by $\sim 5\%$, while varying the delayed and preprompt cut thresholds by 2 PE changed the number of surviving neutron backgrounds by less than 1%; since the energy




Figure 10.13: Initial energy of neutrons generated in the outer cryostat flange for (*black*) all neutrons, (*teal*) neutrons passing all TPC cuts, (*green*) neutrons passing all TPC cuts and the LSV preprompt cut, (*blue*) neutrons passing all TPC cuts and the LSV prompt cut, (*purple*) neutrons passing all TPC cuts and the LSV delayed cut, and (*orange*) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.

Table 10.2: Neutron background predictions from the dominant (α , n) sources in DarkSide-50 for one year. Columns show the total number of neutrons produced in one year, the number expected to pass all TPC cuts, the number expected to pass the TPC cuts and each of the three LSV cuts, and all TPC and all LSV cuts. Neutron sources included here are the PMT stems, PMT hardware, cryostat walls, multilayer insulation (MLI), the inner cryostat flange, and the outer cryostat flange. Uncertainties are discussed in the text.

Source	Total	TPC cuts	Pass all TPC cuts and LSV...			
			Preprompt	Prompt	Delayed	All
PMT Stem	1474	11	8.9	1.2	4.8×10^{-2}	2.3×10^{-2}
PMT HW	34	0.22	0.18	2.6×10^{-2}	1.6×10^{-3}	1.0×10^{-3}
Cryo. Walls	8.5	0.043	0.038	5.1×10^{-3}	2.0×10^{-4}	7.0×10^{-5}
MLI	5.3	0.024	0.018	2.7×10^{-3}	1.7×10^{-4}	1.0×10^{-4}
Inner Flange	5.3	0.0043	0.0030	5.5×10^{-4}	2.6×10^{-5}	7.7×10^{-6}
Outer Flange	794	0.71	0.47	8.7×10^{-2}	3.5×10^{-3}	1.1×10^{-3}
Total	2320	12	10	1.3	5.3×10^{-2}	2.5×10^{-2}




Figure 10.14: Initial energy of neutrons generated in the cryostat walls for (*black*) all neutrons, (*teal*) neutrons passing all TPC cuts, (*green*) neutrons passing all TPC cuts and the LSV preprompt cut, (*blue*) neutrons passing all TPC cuts and the LSV prompt cut, (*purple*) neutrons passing all TPC cuts and the LSV delayed cut, and (*orange*) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.

reconstruction model was tested in the $^{241}\text{Am}^{13}\text{C}$ calibration data, we decided that uncertainty from the LSV energy reconstruction uncertainty is small compared to other sources.

Based on these simulations and the results show in Table 10.2, we draw the following conclusions regarding the (α, n) neutron backgrounds:

- We expect an overall rate of $0.025 \pm 0.001^{+0.001}_{-0.015}$ neutrons/year passing all TPC and LSV cuts, accounting for uncertainties in the neutron production rate and (α, n) yields, respectively.
 - Over the course of a three year campaign, we expect $0.075 \pm 0.003^{+0.003}_{-0.045}$ neutron background events in DarkSide-50.
 - These neutron backgrounds will primarily be produced by the (α, n) reaction in the borosilicate glass of the PMT stems.
- We expect 12 neutrons to pass all TPC cuts per year, before applying neutron veto cuts
 - This rate is consistent with the 1 (α, n) neutron event seen by DarkSide-50 in the 120 days of WIMP search data published so far, which passed all TPC cuts but was vetoed by the LSV.




Figure 10.15: Initial energy of neutrons generated in the multilayer insulation for (black) all neutrons, (teal) neutrons passing all TPC cuts, (green) neutrons passing all TPC cuts and the LSV preprompt cut, (blue) neutrons passing all TPC cuts and the LSV prompt cut, (purple) neutrons passing all TPC cuts and the LSV delayed cut, and (orange) neutrons passing all TPC cuts and all LSV cuts, after 1 year of running.

- We predict that the LSV will veto (α, n) neutron events that pass TPC cuts with an efficiency of 99.8%.

Since this analysis does not account for γ -rays that accompany some (α, n) reactions, it should be noted that this analysis is conservative, since coincident γ -rays will enhance the vetoing efficiency. Furthermore, if these coincident γ -rays scatter in the TPC, they may push the event outside of the f_{90} or energy WIMP search box.

10.3 Fission Neutron Backgrounds

As discussed in Chapter 4.4.2, we expect a total of ~ 250 fission neutrons to be produced each year in all of DarkSide-50, all from the fission of ^{238}U . The highest concentration of ^{238}U is found in the borosilicate glass of the PMTs. If we assume that all of the fission neutrons come from the PMT stem and multiply the fission spectrum in Figure 3.12 by the survival probabilities shown in Figure 10.6, we find that 1.9 neutrons will pass all TPC cuts per year, and 4.4×10^{-3} will pass all TPC and LSV cuts.

The fission neutron spectrum after all TPC cuts, after the TPC cuts plus each individual LSV cut, and after all TPC and LSV cuts can be seen in Figure 10.16.




Figure 10.16: Initial energy of fission neutrons generated in the PMT stems for (*black*) all neutrons, (*teal*) neutrons passing all TPC cuts, (*green*) neutrons passing all TPC cuts and the LSV preprompt cut, (*blue*) neutrons passing all TPC cuts and the LSV prompt cut, (*purple*) neutrons passing all TPC cuts and the LSV delayed cut, and (*orange*) neutrons passing all TPC cuts and all LSV cuts.

This rate of neutrons passing all TPC cuts is consistent with the one suspected fission event seen in the 120 published live days of DarkSide-50. The signal that this event produced in the LSV is shown in Figure 8.16.

In reality, fission neutrons often come in higher multiplicities than do (α, n) neutrons; Table 4.9 shows that the average multiplicity of ^{238}U fission neutrons is 2.01. Fission events with higher multiplicities are far more likely to be vetoed; if more than one neutron scatters in the TPC, they will produce a multiple scatter event, and the probability that neither neutron leaves a signal in the LSV suppressed by an extra factor of the LSV neutron suppression power. Furthermore, $\sim 4\text{--}9$ γ -rays with a mean energy of ~ 1 MeV are typically produced in these fission events. These γ -rays can provide additional veto signals and further suppress fission neutron backgrounds.

Preliminary simulations of fission neutrons in GEANT4 have indicated that fission neutrons are suppressed by the LSV by about two orders of magnitude more than (α, n) neutrons are. This is due to the aforementioned signals produced by the coincident γ -rays and multiple neutrons produced in a fission event. However, because we found that the fission background rate is negligible compared to the (α, n) background rate even without this additional vetoing power, we did not perform a more detailed analysis of the fission neutron backgrounds.

10.4 Gamma-ray Backgrounds

As discussed in Chapter 4, there are many radioisotopes that produce γ -rays that may scatter off of electrons in the TPC.

Most of the DarkSide-50 detector is built out of high-Z elements to minimize the number of neutrons that are blocked from reaching the LSV by the detector materials. While high-Z elements tend to not block neutrons, their abundance of electrons makes them block γ -rays much more effectively.

Preliminary simulations have shown that ^{60}Co and ^{40}K decays in the cryostat and TPC PMTs are typically suppressed by a factor of $\sim 1 \times 10^{-4}$ after applying TPC energy, fiducial, and multiple scattering cuts as well as the LSV prompt cut. However, the LSV specifically only reduces the γ -ray background by a rough factor ~ 2 .

In one year of running, these preliminary simulations indicated that $\sim 2.5 \times 10^5$ ^{60}Co and ^{40}K γ -rays would pass the aforementioned cuts.

Since pulse shape discrimination can easily reject all of these electron recoils, we did not simulate these backgrounds in further detail. We do not expect to see any significant background in DarkSide-50 from γ -ray electron recoils.

Chapter 11

Conclusions

As we discussed in Chapter 1.3.3, in order for a dark matter detector to maximize its sensitivity, it needs to minimize its background and reduce the number of WIMP-nucleus recoils it needs to see to claim a detection. Many dark matter experiments employ powerful tools such as pulse shape discrimination to distinguish between electron and nuclear recoils. Since WIMPs are expected to produce nuclear recoils, nuclear recoil backgrounds are the most dangerous type of background.

Nuclear recoils can arise from two sources. They can result from surface backgrounds on the inner surfaces of the detector, or they can result from neutrons scattering in the bulk of the detector.

Surface backgrounds

Since radon is a gas found naturally in the air in small concentrations, any material that has been exposed to air has also been exposed to radon. Radon may decay and deposit its daughters on the surfaces of these materials. Since the radon decay chain includes α -emitting isotopes, an experiment may expect to find an elevated rate of α -decays on the inner surfaces of the detector.

Surface background events occur when an α -decay on the inner surface of the detector sends a nucleus recoiling into the sensitive volume of the detector. In Chapter 4.3.3, we measured that the recoiling lead nuclei from these surface decays are quenched by a factor of $(19.7 \pm 0.9)^{-1}$, heavily enough that they fall below the WIMP search energy range. However, the additional light produced by the α from the decay scintillating in the TPB wavelength shifter can bring these events back into this energy range.

Since these recoiling nuclei do not travel far in the detector, they are conventionally removed very efficiently by fiducializing the detector. Because fiducialization can be used to remove these events, we do not consider them a significant background in DarkSide-50. However, fiducializing the detector lowers the exposure of the experiment and makes it less sensitive to dark matter recoils. We found the pulse shape of α scintillation in TPB differs from that of nuclear recoils in LAr enough to allow us to discriminate between potential WIMP events and surface backgrounds. We also found that pTP could be used in place of TPB in a future experiment to enhance

this discrimination. C. Stanford is currently working on quantifying the effects of the discriminant; however, it is likely that the use of such a signal will allow DarkSide-50 to fiducialize less and therefore increase its dark matter sensitivity.

Neutron backgrounds

Neutrons can be produced from five different origins:

1. Prompt cosmogenic muon interactions (as discussed in Chapter 4.4.1)
2. β -delayed cosmogenic muon interactions (as discussed in Chapter 4.4.1)
3. Fission reactions from ^{232}Th , ^{238}U , and ^{235}U decay chains in the detector materials (as discussed in Chapter 4.4.2)
4. Radiogenic β -delayed neutron emission from ^{210}Tl in the ^{238}U decay chain in the detector materials (as discussed in Chapter 3.1.8)
5. (α, n) reactions from α -emitters interacting with medium and light nuclei in the detector materials (as discussed in Chapter 4.4.3)

Cosmogenic prompt sources As we discussed in Chapter 10.1, preliminary simulations done by A. Empl of 34 years' worth of cosmogenic muons showed that no cosmogenic neutrons survived rudimentary cuts in both the LSV and WCV.

While some details of these simulations do not reflect the final implementation of DarkSide-50, these simulations provide strong evidence that the DarkSide-50 outer detector veto system will effectively remove all cosmogenic neutron backgrounds.

A more detailed analysis using the actual cuts implemented in DarkSide-50 is currently being developed by P. Singh. However, the strength of the results found by A. Empl lead us to believe that the final cosmogenic vetoing efficiency will not change dramatically when the simulation is updated.

Cosmogenic β -delayed sources As we discussed in Chapter 10.1, we roughly expect ~ 1.43 β -delayed neutron precursors to be produced by muons outside of the cryostat over the course of a three year campaign. These neutrons have a comparable energy scale to the (α, n) neutrons. Since they are all generated outside the cryostat, we may conservatively estimate that their probability of passing TPC cuts is comparable to the probability of neutrons generated in the cryostat doing so. In reality, it is presumably much more unlikely that these neutrons will pass all TPC cuts, since they may lose significant amounts of energy in the LSV before reaching the cryostat.

Nevertheless, Table 10.2 indicates that we may expect $\lesssim 7.2 \times 10^{-3}$ of these neutrons to pass all TPC cuts during the course of a three year campaign, and $\lesssim 1.2 \times 10^{-5}$ to pass all TPC and LSV cuts. We therefore conclude that these neutrons produce a negligible background contribution to DarkSide-50.

Radiogenic fission sources As we showed in Chapter 10.3, we expect the background of fission neutrons surviving all TPC and LSV cuts to be negligible.

In a simplistic model, where we ignore fission events that produce multiple neutrons or γ -rays, we found that fission reactions in DarkSide-50 produced single scatter nuclear recoils in the WIMP search energy range at a rate of ~ 2 events/year. However, the prompt and delayed LSV cuts brought the total fission background down by a few additional orders of magnitude.

Preliminary simulations that factor in these γ -rays and additional neutrons show that these backgrounds can be further dramatically reduced by the LSV to a low enough that these backgrounds become negligible compared to the predicted (α, n) backgrounds.

Radiogenic β -delayed sources As we estimated in Chapter 3.1.8, we expect 2.1 ± 0.6 neutrons to be produced by this process over the course of a three year campaign of DarkSide-50. These neutrons will have energy around ~ 201 – 260 keV. Due to the low energy of these neutrons, they can deposit at most ~ 25 keV into the LAr with a single nuclear recoil.

Figure 10.2 shows that ~ 200 keV neutrons coming from the PMT stems—the dominant source of ^{210}Pb —have a 1 in 10^5 chance of passing all TPC cuts. We therefore conclude that neutrons from this source provide a negligible contribution to the overall nuclear recoil background rate.

Radiogenic (α, n) sources Since (α, n) neutrons can be produced by any α -emitting isotope, these neutrons can be produced at a higher rate than fission neutrons. Furthermore, since they are often produced by themselves, without accompanying γ -rays or other neutrons, they pose the most insidious risks to dark matter experiments.

We therefore studied these in the greatest detail. In Chapter 4.4.3, we developed a calculator to determine (α, n) yields in various material based on the TALYS software, and we compiled a detailed list of (α, n) yields for many materials commonly used in dark matter experiments in Tables 4.10 and 4.11. We used these yields to build an (α, n) neutron background model for DarkSide-50, which is summarized in Table 4.13.

We simulated these neutrons in Chapter 10.2 and determined that the LSV vetoes (α, n) neutrons that pass all TPC cuts with an efficiency of 99.8%.

11.1 Impact of Nuclear Recoil Backgrounds

Based on the backgrounds considered in this document and the discussion in Chapter 1.3.3 and Figure 1.12, we draw the following conclusions:

- Without the LSV we could expect 14 radiogenic neutrons to pass all TPC cuts per year.
- Without the WCV and LSV, we would expect ~ 8.8 cosmogenic neutrons to produce nuclear recoil backgrounds per year.

- We predict that DarkSide-50 will see $0.025 \pm 0.001^{+0.001}_{-0.015}$ nuclear recoil backgrounds/year and $0.075 \pm 0.003^{+0.003}_{-0.045}$ in three years.
- Nuclear recoil backgrounds predominantly come from the (α, n) reaction in the borosilicate glass of the Hamamatsu R11065 PMT stems.
 - Eliminating this source of backgrounds by using cleaner PMTs or other photodetector technology such as silicon photomultipliers would have decreased the neutron background by a factor of ~ 10 .
- DarkSide-50 would need to see 3 WIMP recoils to claim an observation with 3σ significance.
- DarkSide-50 would need to see 5 WIMP recoils to claim a discovery with 5σ significance.

This allows DarkSide-50 to achieve a very high degree of WIMP sensitivity, as illustrated in Figure 1.14.

This background estimate was built upon the model of DarkSide-50 as discussed in Chapter 2, and the background model developed in Chapters 3 and 4. This level of background is unprecedented in WIMP experiments, and was only achievable by the active shielding provided by the neutron veto system, consisting of the LSV and WCV, as were introduced in Chapter 5. DarkSide-50 is the first dark matter experiment to use such a veto system, though a similar veto system could be built for future dark matter detectors as well, such as the neutron veto being designed for the LZ experiment [82].

The prototype studies and simulations that guided the development of the LSV are discussed in Chapter 6 and Chapter 7, including tests of various quenching models describing the response of the scintillator to energy depositions. This quenching model was verified and tuned for the DarkSide-50 LSV through a set of calibration campaigns discussed in Chapter 9, where we found a model that is quadratic in the stopping power—given by Model II in Chapter 6.3—describes the quenching of individual nuclear species well. However, we found that a single set of parameters could not be used to describe the quenching of all nuclei, indicating that this model is not complete. Instead, we built an effective model that uses one set of parameters to describe proton quenching, and another set of parameters to describe the quenching of heavier nuclei and electrons. In order to perform these calibrations, we developed an $^{241}\text{Am}^{13}\text{C}$ neutron calibration source; the design and predicted spectrum of which is described in Chapter 9.3.

Lastly, the implementation of the DarkSide-50 WCV and LSV, along with a description of the data acquisition and analysis procedures, is discussed in Chapter 8. In this chapter, we describe the performance of the DarkSide-50 outer detectors in detail as well as all of the cuts we made using these detectors to efficiently reject neutron backgrounds in the TPC.

11.2 Additional Backgrounds

In the analysis presented in this document, we addressed surface backgrounds and neutron backgrounds in great detail in order to show how an experiment such as DarkSide-50 can use an active veto system to run background free for an extended campaign. While these are the dominant nuclear recoil backgrounds in most WIMP experiments, it is worth drawing attention to a few additional sources of background that we have neglected.

11.2.1 Cherenkov-Mixed Electron Recoils

In Chapters 4.1.2 and 4.2.3, we acknowledge another class of potential backgrounds that may mimic nuclear recoils, despite being entirely electronic in origin. These backgrounds are caused by Cherenkov light being produced by high energy electrons in the TPC Teflon or fused silica producing a very fast pulse in the TPC PMTs, while a slower pulse is caused by an electron recoil from a γ -ray Compton scattering in the LAr. These high energy electrons may be produced by β -decay of contaminants in the Teflon or fused silica, or by γ -rays Compton scattering in these materials and ionizing an electron. Since β -decays are often accompanied by γ -rays, either of these processes are likely to produce a Cherenkov signal in coincidence with an electron recoil.

In DarkSide-50, we see some evidence that the ITO coating on the fused silica may greatly reduce such events from occurring in the fused silica, since ITO is opaque to much of the Cherenkov light produced. Cherenkov events in the fused silica can be further identified by looking for events in which a disproportionately high amount of light is seen in a single PMT.

Such techniques, however, do not help suppress mixed Cherenkov events in the Teflon, where the S1 light distribution may be similar to that of an ordinary scintillation event. While some of these events may have prompt coincidence signals in the LSV where the γ -ray enters the liquid scintillator after scattering the LAr, enough γ -rays are blocked by the stainless steel of the cryostat that we may not be able to rely on the LSV vetoing all such events.

While DarkSide-50 has not yet seen any background events survive all TPC and LSV cuts yet, G. Koh and X. Xiang are currently investigating them to quantify the background that they pose. They are also studying these events to see if other features of these events such as their pulse shape may be used to discriminate between them and actual nuclear recoils.

11.2.2 Neutrino-Induced Nuclear Recoils

Neutrino backgrounds in dark matter experiments are discussed in detail by J. Billard and E. Figueroa-Feliciano [86]. While these studies mostly focus on dark matter experiments with liquid xenon or germanium targets, they show that neutrino-induced nuclear recoils start to become significant for tonne-scale detectors, where we see they

reach a rate of ~ 0.01 event/tonne-year for xenon detectors with an energy threshold of ~ 10 keV.

While these backgrounds are irrelevant for an experiment such as DarkSide-50, where the neutron-induced nuclear recoil rate is negligible for a 50 kg LAr detector, it may become relevant as detectors reach the ~ 10 tonne scale. Since neutrino scattering cross sections are so small, it is unlikely that a veto design such as the one we have discussed here would offer much reduction in their background rate.

It is worth noting that neutrino-induced electron recoils tend to happen at significantly higher rates than do nuclear recoils. This means that a LAr experiment, which can reject electron recoils with very high efficiency have little risk of experiencing significant background from neutrino-electron scattering events. Liquid xenon detectors can typically reject electron recoils by a factor of ~ 200 [86], and therefore risk electron recoil backgrounds becoming significant before nuclear recoil backgrounds do.

11.3 Larger Detectors

As experiments grow larger and increase their exposure, it becomes progressively more important that the detector be able to effectively reject background events.

While we have not performed a full simulation of a larger detector here, we have done preliminary studies on the effectiveness of rejecting various neutron backgrounds in a LAr TPC. While more precise simulations must be done to better understand the effects of scaling up the size of the detector, we have found that a 3.8 tonne LAr TPC with at least 1 m of liquid scintillator completely surrounding it can maintain a high neutron rejection efficiency.

In particular, we found that as the size of the detector grows, a higher fraction of neutrons are rejected by multiple scattering and fiducial cuts, while the neutron veto rejects slightly fewer of the neutrons as a result. Nevertheless, we predict that a boron-loaded liquid scintillator veto could be a key component of a large background-free LAr dark matter detector.

11.4 Final Remarks

By better understanding neutron backgrounds and using a two-part neutron veto system such as the one presented here, it is possible for a dark matter experiment to maximize its WIMP sensitivity by running background free for an extended campaign.

As experiments grow more and more sensitive, we face many challenges pushing limits on the WIMP-nucleus scattering cross section down progressively further. In order for an experiment to convincingly detect such a rare event, it is important that the backgrounds in the experiment be understood as well as possible and be reduced to their lowest obtainable levels. As we seek to further understand what dark matter is, it is important that we be able to confidently characterize dark matter with what little data we may have available from future experiments.

We have shown the neutron veto design presented here can very effectively remove neutron backgrounds from a dark matter experiment, allowing these experiments to search for dark matter with much greater sensitivity.

Appendix A

Neutron Yield Calculation and Measurement Comparisons

As discussed in Chapter 4.4.3, the (α, n) neutron yield calculator presented in this document needs to be validated against data and prior calculation tools before it can be trusted so that its uncertainties and reliability can be better understood.

While Chapter 4.4.3 contains some general validation of the calculations performed by the code, in this section we will present more direct side-by-side comparisons for α energies between 5 and 11 MeV incident upon lithium, beryllium, boron, carbon, nitrogen, oxygen, fluorine, sodium, magnesium, aluminum, silicon, iron, cobalt, nickel, copper, and zinc targets.

We will compare calculations made by this work to those made using the SOURCES4A tool, wherever adequate data exists in the SOURCES4A libraries. We will also compare these calculations to measurements from several experiments. For elements ranging from lithium to iron, we draw heavily upon the compilation work done by Heaton et al [31]. We also include measurements by Bair and Gomez del Campo [33] and Roughton et al [152] when they were not cited in Heaton's compilation. The measurements taken by Bair and Gomez del Campo were performed by irradiating targets in an α beam and then counting the neutrons with a graphite neutron detector. The measurements reported by Roughton et al were done by irradiating a thick target in an α beam and then measuring radioactive decays of (α, n) daughter nuclei immediately afterwards; this method allows measurements to be taken independently of the neutron's energy, though they are constrained by the production of measurable radioisotopes from the (α, n) process.

For heavier metals, we drew heavily on the measurements taken by Stelson and McGowan [32], using a very similar apparatus as was used by Bair and Gomez del Campo. These measurements focus on heavier metals, though only those up to zinc are of particular interest to a dark matter experiment, since the heavier metals are found in much lower masses, and their threshold α energies for undergoing the (α, n) reaction become higher than we find in the uranium and thorium decay chains. Many of the measurements of the heavier elements, as reported by Stelson and McGowan as well as by Roughton et al distinguish between different isotopes, as well. To show the measurements for the different isotopes of each element, we plot the different graphs

separately but in the same plot, and label the groups of graphs with the isotope they represent.

Measurements from these references typically report uncertainties in the range of 10–20%.

Equation 4.15 can be used to estimate the Coulomb barrier for each element. The Coulomb barrier gives a sense of scale in the energy range for when the (α, n) process “turns on”—even though the reaction can occur at lower α energies of the α tunnels through the barrier, we tend to find the rate increasing rapidly around the point where the α energy surpasses this barrier. Measurements of the parameter r_0 in Equation 4.15 vary from 1.3–1.5 fm, and so we represent the Coulomb barrier with a shaded region to show the upper and lower bounds of where we might expect it to lay based on the range of r_0 values. Additionally, different isotopes of the same element will all have different Coulomb barriers due to their different nuclear sizes. In the case of elements where we plot (α, n) yields of individual isotopes, we expand the range of the shaded region to include the range of Coulomb barriers from all of the naturally present isotopes.

For elements lighter than sodium, the Coulomb barrier was too low to appear on the plots, so these barriers are not shown.

Comparisons in (α, n) calculations and measurements are shown in Figures A.1, A.2, A.3, and A.4. As can be seen in these figures, the work presented here generally over-estimates the (α, n) yields in the benchmarks, making it a conservative calculation, and reproduces general shapes observed in the measured neutron yield curves. It should be noted that the frequent close agreement between SOURCES4A predictions and measurements is not a coincidence: SOURCES4A is heavily data-driven and frequently uses (α, n) cross section measurements taken from the same references shown here. In particular, many of the (α, n) cross sections used by SOURCES4A come from the works done by Bair and Gomez del Campo.

In addition to comparing absolute yields, it is useful to compare the neutron energy spectra generated by this calculator to benchmarks. Unfortunately, the measurements mentioned above all report the absolute neutron yield but do not report the spectrum of the neutrons emitted by this process. Instead, we rely on SOURCES4A as benchmark for the spectra, keeping in mind its limitations in available data for various isotopes at different energies.

We choose as representative elements boron, carbon, oxygen, fluorine, silicon, iron, chromium, and nickel. These elements were chosen for their abundance in the parts of DarkSide-50 that we have calculated have the greatest total (α, n) reaction rate. In particular, boron, carbon, oxygen, and silicon are abundant in the borosilicate glass of the PMTs, while oxygen and carbon are found in the Mylar multilayer insulation, and carbon and fluorine are abundant in the Viton o-ring of the outer cryostat and the Teflon reflector. Meanwhile, iron, nickel, and chromium are the primary components of the stainless steel of the inner and outer cryostats.

These spectra are shown in Figures A.5, A.6, A.7, A.8, A.9, A.10, A.11, and A.12. As can be seen in these figures, these spectra mostly agree with each other; in most cases, they share most of their common features and similar general shape. This agreement makes sense, since both calculators base their spectra on equations similar




Figure A.1: (α, n) yields calculated by the present work as compared to a series of benchmarks for lithium, beryllium, boron, and carbon targets. See text for a description of the reference sources. (*Blue solid dots*) Calculations made by the present work, (*green squares*) measurements reported by Heaton et al, (*empty cyan circles*) measurements made by Blair and Gomez del Campo, (*purple triangles*) calculations made by the SOURCES4A code.

to Equation 4.21. However, we generally see more defined features in the spectra generated by the calculator presented in this work; this is likely because TALYS uses more detailed model to predict different nuclear structures than SOURCES4A does, and therefore has more values of Q_m .

The spectra start to diverge more dramatically for the metals, where we see major differences in the spectral shapes, especially around 5 MeV where SOURCES4A predicts a total rate of 0 neutrons per α , and this calculator predicts a positive (α, n) yield. These discrepancies are generally due to elements where SOURCES4A is missing data on key isotopes that have significant (α, n) yields compared to the other naturally present isotopes. This effect is especially seen in chromium, for which SOURCES4A lacks data for ^{52}Cr , ^{53}Cr , and ^{64}Cr , and nickel, for which SOURCES4A is missing data for ^{58}Ni and ^{61}Ni .

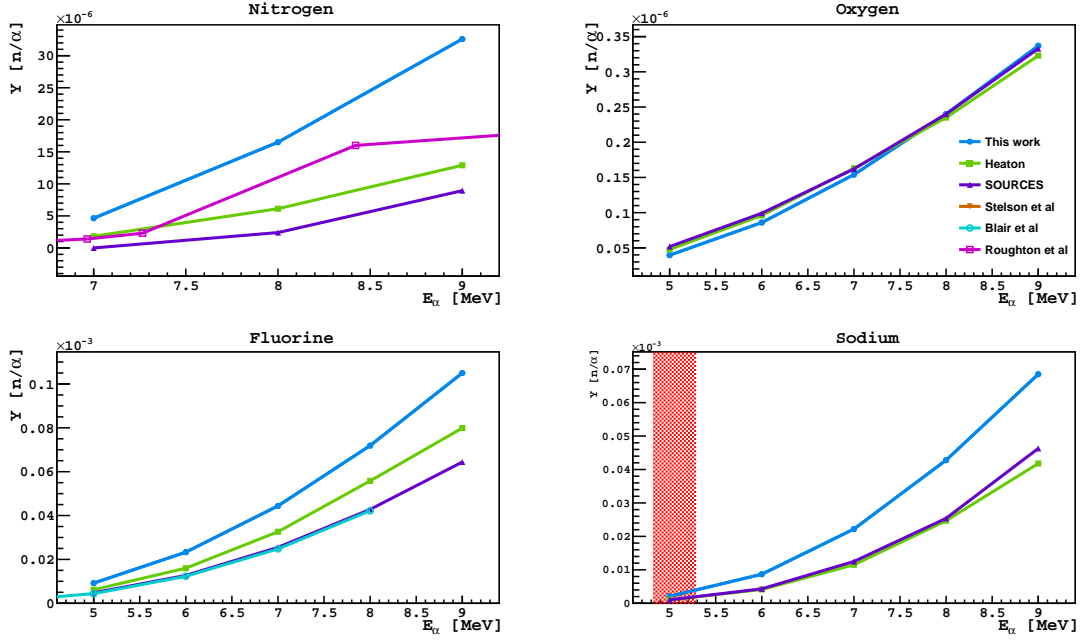


Figure A.2: (α, n) yields calculated by the present work as compared to a series of benchmarks for nitrogen, oxygen, fluorine, and sodium targets. See text for a description of the reference sources. (*Blue solid dots*) Calculations made by the present work, (*green squares*) measurements reported by Heaton et al, (*purple triangles*) calculations made by the SOURCES4A code, (*empty light purple squares*) measurements by Roughton et al. Regions highlighted in red show where the Coulomb barrier for a given nucleus falls.

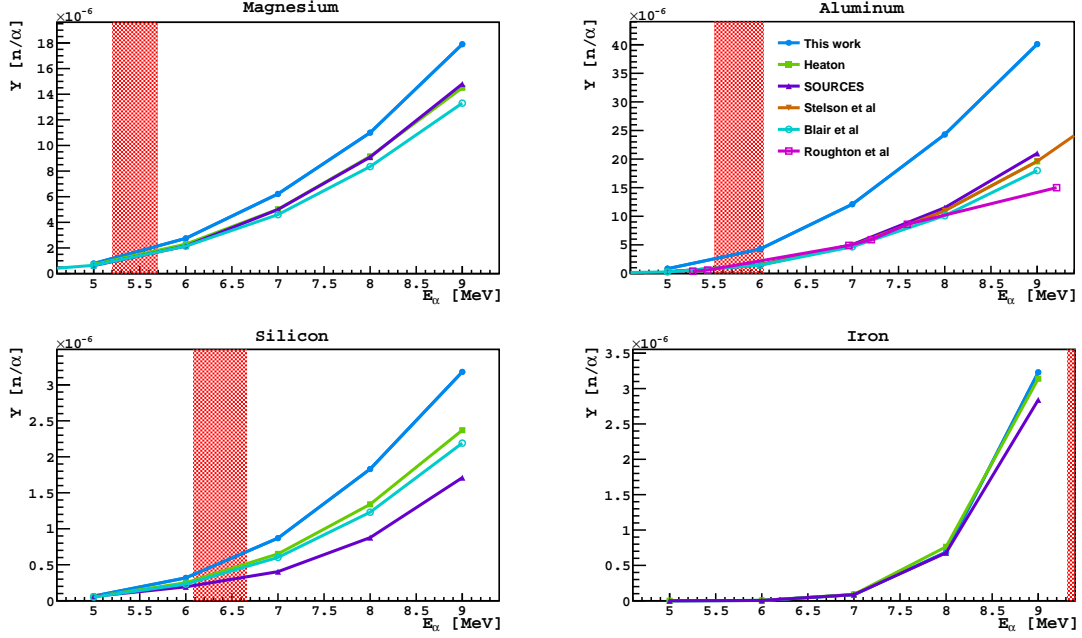


Figure A.3: (α, n) yields calculated by the present work as compared to a series of benchmarks for magnesium, aluminum, silicon, and iron targets. See text for a description of the reference sources. (*Blue solid dots*) Calculations made by the present work, (*green squares*) measurements reported by Heaton et al, (*empty cyan circles*) measurements made by Bair and Gomez del Campo, (*purple triangles*) calculations made by the SOURCES4A code, (*empty light purple squares*) measurements by Roughton et al, (*upside-down orange triangles*) measurements reported by Stelson and McGowan. Regions highlighted in red show where the Coulomb barrier for a given nucleus falls.




Figure A.4: (α, n) yields calculated by the present work as compared to a series of benchmarks for cobalt, nickel, copper, and zinc targets. See text for a description of the reference sources. (*Blue solid dots*) Calculations made by the present work, (*green squares*) measurements reported by Heaton et al, (*empty cyan circles*) measurements made by Bair and Gomez del Campo, (*purple triangles*) calculations made by the SOURCES4A code, (*empty light purple squares*) measurements by Roughton et al, (*upside-down orange triangles*) measurements reported by Stelson and McGowan. Regions highlighted in red show where the Coulomb barrier for a given nucleus falls.




Figure A.5: (α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick boron target, as calculated both by (*green*) SOURCES4A and (*blue*) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.




Figure A.6: (α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick carbon target, as calculated both by (*green*) SOURCES4A and (*blue*) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.




Figure A.7: (α, n) spectrum for 5, 6, 7, and 8 MeV α incident on a thick oxygen target, as calculated both by (*green*) SOURCES4A and (*blue*) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.




Figure A.8: (α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick fluorine target, as calculated both by (*green*) SOURCES4A and (*blue*) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.




Figure A.9: (α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick silicon target, as calculated both by (green) SOURCES4A and (blue) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.




Figure A.10: (α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick iron target, as calculated both by (green) SOURCES4A and (blue) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.




Figure A.11: (α, n) spectrum for 5, 6, 7, and 8 MeV α incident on a thick chromium target, as calculated both by (green) SOURCES4A and (blue) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.

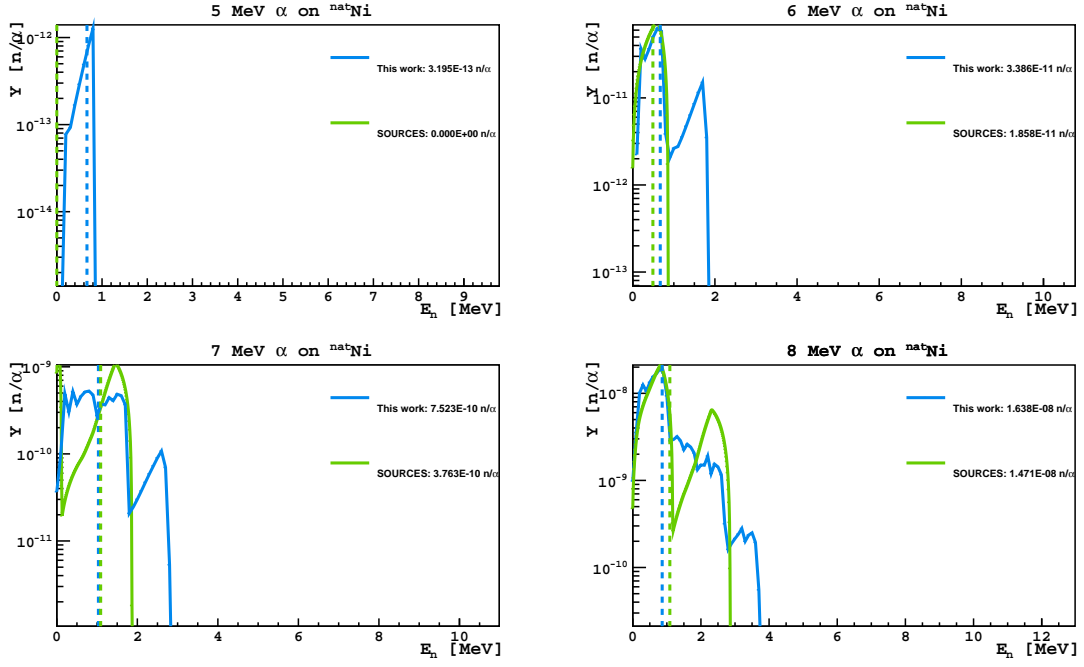


Figure A.12: (α, n) spectrum for 5, 6, 7, and 8 MeV α s incident on a thick nickel target, as calculated both by (*green*) SOURCES4A and (*blue*) this work. Dashed vertical lines show the means of the spectrum of the matching color. For comparison, the spectrum generated by SOURCES4A was scaled to have the same integral as the spectrum calculated by the present work. However, the total rates calculated by each tool prior to scaling are shown in the legends.

Bibliography

- [1] A. L. Fitzpatrick, et al., *The effective field theory of dark matter direct detection*, [Journal of Cosmology and Astroparticle Physics](#), **2013**, 004 (2013), ISSN 1475-7516.
- [2] M. I. Gresham and K. M. Zurek, *Effect of nuclear response functions in dark matter direct detection*, [Physical Review D](#), **89** (2014), ISSN 1550-7998, 1550-2368.
- [3] J. A. Cameron and B. Singh, *Nuclear data sheets for $a=40$* , [Nuclear Data Sheets](#), **102**, 293 (2004).
- [4] W. Wilson et al., *Sources-4c: A code for calculating (α , n) spontaneous fission and delayed neutron sources and spectra*, american nuclei society, [Radiation Protection and Shielding Division](#), (2002).
- [5] D. Tilley, H. Weller, C. Cheves, and R. Chasteler, *Energy levels of light nuclei $a=18-19$* , [Nuclear Physics A](#), **595**, 1 (1995).
- [6] Y. Dong and H. Junde, *Nuclear data sheets for $a=52$* , [Nuclear Data Sheets](#), **128**, 185 (2015).
- [7] M. B. Chadwick, et al., *ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data*, [Nucl. Data Sheets](#), **112**, 2887 (2011), ISSN 0090-3752.
- [8] M. Giammarchi, *Water purification in Borexino*, [Int. J. Mod. Phys. A](#), **29**, 1442008 (2014), ISSN 0217-751X, 1793-656X.
- [9] K. G. Begeman, A. H. Broeils, and R. H. Sanders, *Extended rotation curves of spiral galaxies - Dark haloes and modified dynamics*, [Monthly Notices of the Royal Astronomical Society](#), **249**, 523 (1991), ISSN 0035-8711.
- [10] M. Markevitch, *Constraints on physics of gas and dark matter from cluster mergers* (2005), talk given at “Six Years of Science with Chandra” symposium.
- [11] P. A. R. Ade, et al. (The Planck Collaboration), *Planck 2013 results. I. Overview of products and scientific results*, [arXiv:1303.5062 \[astro-ph\]](#), (2013).

- [12] P. Agnes, et al. (The DarkSide Collaboration), *Results from the first use of low radioactivity argon in a dark matter search*, [Physical Review D](#), **93**, 081101 (2016).
- [13] S. Westerdale, E. Shields, and F. Calaprice, *A Prototype Neutron Veto for Dark Matter Detectors*, [arXiv:1509.02782 \[hep-ex, physics:physics\]](#), (2015), [arXiv:1509.02782](#).
- [14] Hamamatsu Corporation, *R11065 Photomultiplier Tube Data Sheet*, Hamamatsu Corporation, (2009) .
- [15] B. Singh and J. A. Cameron, *Nuclear data sheets for $a=39$* , [Nuclear Data Sheets](#), **107**, 225 (2006).
- [16] B. Singh and J. Chen, *Nuclear data sheets for $a=85$* , [Nuclear Data Sheets](#), **116**, 1 (2014).
- [17] E. Browne and J. Tuli, *Nuclear data sheets for $a=60$* , [Nuclear Data Sheets](#), **114**, 1849 (2013).
- [18] J. Purcell and C. Sheu, *Nuclear data sheets for $a=3$* , [Nuclear Data Sheets](#), **130**, 1 (2015).
- [19] F. Ajzenberg-Selove, *Energy levels of light nuclei $A=1315$* , [Nuclear Physics A](#), **523**, 1 (1991), ISSN 0375-9474.
- [20] F. E. Wietfeldt, et al., *Further studies on the evidence for a 17-keV neutrino in a 14 doped germanium detector*, [Physical Review C](#), **52**, 1028 (1995).
- [21] K. Abusaleem, *Nuclear data sheets for $a=228$* , [Nuclear Data Sheets](#), **116**, 163 (2014).
- [22] E. Browne and J. Tuli, *Nuclear data sheets for $a=234$* , [Nuclear Data Sheets](#), **108**, 681 (2007).
- [23] M. Shamsuzzoha Basunia, *Nuclear data sheets for $a=210$* , [Nuclear Data Sheets](#), **121**, 561 (2014).
- [24] J. Chen and F. Kondev, *Nuclear data sheets for $a=209$* , [Nuclear Data Sheets](#), **126**, 373 (2015).
- [25] E. Browne and J. Tuli, *Nuclear data sheets for $a=231$* , [Nuclear Data Sheets](#), **114**, 751 (2013).
- [26] R. Francini, et al., *Tetraphenyl-butadiene films: VUV-Vis optical characterization from room to liquid argon temperature*, [Journal of Instrumentation](#), **8**, C09010 (2013), ISSN 1748-0221.

- [27] J. M. Dixon, M. Taniguchi, and J. S. Lindsey, *PhotochemCAD 2: A Refined Program with Accompanying Spectral Databases for Photochemical Calculations*, [Photochemistry and Photobiology](#), **81**, 212 (2005), ISSN 1751-1097.
- [28] A. Empl, E. V. Hungerford, R. Jasim, and P. Mosteiro, *A Fluka study of underground cosmogenic neutron production*, [JCAP](#), **2014**, 064 (2014).
- [29] C. Galbiati and J. F. Beacom, *Measuring the cosmic ray muon-induced fast neutron spectrum by (n,p) isotope production reactions in underground detectors*, [Phys. Rev. C](#), **72**, 025807 (2005).
- [30] N. E. Holden and M. S. Zucker, *A reevaluation of the average prompt neutron emission multiplicity (nubar) values from fission of uranium and transuranium nuclides*, National Nuclear Data Center, Brookhaven National Laboratory (1984).
- [31] R. Heaton, H. Lee, P. Skensved, and B. C. Robertson, *Neutron production from thick-target (α , n) reactions*, [Nucl. Inst. Meth. A](#), **276**, 529 (1989).
- [32] P. H. Stelson and F. K. McGowan, *Cross Sections for (α ,n) Reactions for Medium-Weight Nuclei*, [Physical Review](#), **133**, B911 (1964).
- [33] J. K. Bair and J. Gomez del Campo, *Neutron Yields from Alpha-Particle Bombardment*, [Nuclear Science and Engineering](#), **71**, 18 (1979), ISSN 00295639.
- [34] M. Johnson, *Scintillator Purification and Study of Light Propagation in a Large Liquid Scintillation Detector*, Ph.D. thesis, Princeton University (1998).
- [35] I. Berلمان, *Handbook of fluorescence spectra of Aromatic Molecules* Elsevier (2012), ISBN 9780323161671.
- [36] I. B. Nemchenok, et al., *Liquid scintillator based on linear alkylbenzene*, [Physics of Particles and Nuclei Letters](#), **8**, 129 (2011), ISSN 1547-4771, 1531-8567.
- [37] H.-L. Xiao, et al., *Study of absorption and re-emission processes in a ternary liquid scintillation system*, [Chinese Physics C](#), **34**, 1724 (2010), ISSN 1674-1137.
- [38] C. A. Burke, M. T. Lunnon, and H. W. Lefevre, *$Li^7(p,n^0)Be^7$ angular distributions to $E_p=3.8$ MeV*, [Physical Review C](#), **10**, 1299 (1974).
- [39] J. Hong, et al., *The scintillation efficiency of carbon and hydrogen recoils in an organic liquid scintillator for dark matter searches*, [Astroparticle Physics](#), **16**, 333 (2002).
- [40] P. Agnes, et al., *The veto system of the DarkSide-50 experiment*, [Journal of Instrumentation](#), **11**, P03016 (2016), ISSN 1748-0221.
- [41] Hamamatsu Corporation, *Large photocathode area photomultiplier tubes*, Hamamatsu Corporation, (2008) .

- [42] P. Agnes, et al. (the DarkSide Collaboration), *Low radioactivity argon dark matter search results from the DarkSide-50 experiment*, [arXiv:1510.00702 \[astro-ph, physics:hep-ex, physics:physics\]](#), (2015).
- [43] A. D. Vijaya and A. Kumar, *The neutron spectrum of Am-Be neutron sources*, [Nuclear Instruments and Methods](#), **111**, 435 (1973), ISSN 0029-554X.
- [44] T. W. Bonner, A. A. Kraus, J. B. Marion, and J. P. Schiffer, *Neutrons and Gamma Rays from the Alpha-Particle Bombardment of Be 9 , B 10 , B 11 , C 13 , and O 18*, [Physical Review](#), **102**, 1348 (1956), ISSN 0031-899X.
- [45] M. Basunia, *Nuclear data sheets for a= 241*, [Nuclear Data Sheets](#), **107**, 3323 (2006).
- [46] G. Aad, et al. (The ATLAS Collaboration), *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, [Physics Letters B](#), **716**, 1 (2012), ISSN 0370-2693.
- [47] S. Chatrchyan, et al. (The CMS Collaboration), *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, [Physics Letters B](#), **716**, 30 (2012), ISSN 0370-2693.
- [48] D. N. Spergel, et al., *First-Year Wilkinson Microwave Anisotropy Probe(WMAP) Observations: Determination of Cosmological Parameters*, [Ap. J. Supp. Ser.](#), **148**, 175 (2003).
- [49] T. Alexander, et al. (The DarkSide Collaboration), *Light yield in DarkSide-10: A prototype two-phase argon TPC for dark matter searches*, [Astropart. Phys.](#), **49**, 44 (2013).
- [50] P. Agnes, et al. (The DarkSide Collaboration), *First results from the DarkSide-50 dark matter experiment at Laboratori Nazionali del Gran Sasso*, [Phys. Lett. B](#), **743**, 456 (2015).
- [51] F. Zwicky, *Die Rotverschiebung von extragalaktischen Nebeln*, [Helv. Phys. Acta](#), **6**, 110 (1933).
- [52] V. C. Rubin and W. K. Ford, Jr., *Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions*, [The Astrophysical Journal](#), **159**, 379 (1970), ISSN 0004-637X, 1538-4357.
- [53] V. C. Rubin, et al., *Motion of the Galaxy and the local group determined from the velocity anisotropy of distant SC I galaxies. I - The data*, [The Astronomical Journal](#), **81**, 687 (1976), ISSN 00046256.
- [54] K. C. Freeman, *On the Disks of Spiral and so Galaxies*, [The Astrophysical Journal](#), **160**, 811 (1970), ISSN 0004-637X, 1538-4357.

- [55] J. P. Ostriker and P. J. E. Peebles, *A Numerical Study of the Stability of Flattened Galaxies: or, can Cold Galaxies Survive?*, [The Astrophysical Journal](#), **186**, 467 (1973), ISSN 0004-637X, 1538-4357.
- [56] M. Moniez, *Review of results from EROS Microlensing search for Massive Compact Objects*, [arXiv:0901.0985 \[astro-ph\]](#), (2009), arXiv: 0901.0985.
- [57] L. Wyrzykowski, et al., *The OGLE view of microlensing towards the Magellanic Clouds III. Ruling out subsolar MACHOs with the OGLE-III LMC data*, [Monthly Notices of the Royal Astronomical Society](#), **413**, 493 (2011), ISSN 0035-8711, 1365-2966.
- [58] D. Clowe, et al., *A Direct Empirical Proof of the Existence of Dark Matter*, [The Astrophysical Journal Letters](#), **648**, L109 (2006), ISSN 1538-4357.
- [59] W. Hu, *Intermediate guide to the acoustic peaks and polarization* (2001).
- [60] C. Conroy, R. H. Wechsler, and A. V. Kravtsov, *Modeling Luminositydependent Galaxy Clustering through Cosmic Time*, [The Astrophysical Journal](#), **647**, 201 (2006), ISSN 0004-637X, 1538-4357.
- [61] M. Pospelov and A. Ritz, *Electric dipole moments as probes of new physics*, [Annals of Physics](#), **318**, 119 (2005), ISSN 0003-4916.
- [62] P. A. R. Ade, et al. (The Planck Collaboration), *Planck 2013 results. I. Overview of products and scientific results*, [Astronomy & Astrophysics](#), **571**, A1 (2014), ISSN 0004-6361, 1432-0746.
- [63] J. L. Feng, *Dark Matter Candidates from Particle Physics and Methods of Detection*, [Annual Review of Astronomy and Astrophysics](#), **48**, 495 (2010).
- [64] R. N. Mohapatra and G. Senjanovi, *Neutrino Mass and Spontaneous Parity Nonconservation*, [Physical Review Letters](#), **44**, 912 (1980), ISSN 0031-9007.
- [65] J. Schechter and J. W. F. Valle, *Neutrino masses in $SU(2) \otimes U(1)$ theories*, [Physical Review D](#), **22**, 2227 (1980), ISSN 0556-2821.
- [66] T. Asaka and S. Eijima, *Direct search for right-handed neutrinos and neutrinoless double beta decay*, [Progress of Theoretical and Experimental Physics](#), **2013**, 113B02 (2013), ISSN 2050-3911.
- [67] K. Petraki and R. R. Volkas, *REVIEW OF ASYMMETRIC DARK MATTER*, [International Journal of Modern Physics A](#), **28**, 1330028 (2013), ISSN 0217-751X, 1793-656X.
- [68] G. Bertone, D. Hooper, and J. Silk, *Particle dark matter: evidence, candidates and constraints*, [Phys. Rep.](#), **405**, 279 (2005).
- [69] K. Hagiwara, et al., *Review of Particle Properties*, [Physical Review D](#), **66** (2002), ISSN 0556-2821, 1089-4918.

- [70] B. A. Dobrescu and I. Mocioiu, *Spin-Dependent Macroscopic Forces from New Particle Exchange*, [Journal of High Energy Physics](#), **2006**, 005 (2006), ISSN 1029-8479, arXiv: hep-ph/0605342.
- [71] J. B. Muoz, E. D. Kovetz, and Y. Ali-Hamoud, *Heating of baryons due to scattering with dark matter during the dark ages*, [Physical Review D](#), **92** (2015), ISSN 1550-7998, 1550-2368.
- [72] M. Farina, et al., *Implications of Xenon100 and LHC results for Dark Matter models*, [Nuclear Physics B](#), **853**, 607 (2011), ISSN 05503213.
- [73] O. Buchmueller, et al., *Likelihood functions for supersymmetric observables in frequentist analyses of the CMSSM and NUHM1*, [The European Physical Journal C](#), **64**, 391 (2009), ISSN 1434-6044, 1434-6052.
- [74] K. Garrett and G. Duda, *Dark Matter: A Primer*, 1006.2483, (2010) , adv.Astron.2011:968283,2011.
- [75] J. D. Lewin and P. F. Smith, *Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil*, [Astroparticle Physics](#), **6**, 87 (1996), ISSN 0927-6505.
- [76] G. R. Knapp, S. D. Tremaine, and J. E. Gunn, *The global properties of the Galaxy. I - The H I distribution outside the solar circle*, [The Astronomical Journal](#), **83**, 1585 (1978), ISSN 00046256.
- [77] F. J. Kerr and D. Lynden-Bell, *Review of galactic constants*, [Monthly Notices of the Royal Astronomical Society](#), **221**, 1023 (1986), ISSN 0035-8711, 1365-2966.
- [78] J. A. R. Caldwell and I. M. Coulson, presentation at [The Outer Galaxy \(1988\)](#), edited by L. Blitz and F. J. Lockman, vol. 306, pp. 68–72.
- [79] A. K. Drukier, K. Freese, and D. N. Spergel, *Detecting cold dark-matter candidates*, [Physical Review D](#), **33**, 3495 (1986), ISSN 0556-2821.
- [80] E. I. Gates, G. Gyuk, and M. S. Turner, *The Local Halo Density*, [The Astrophysical Journal Letters](#), **449**, L123 (1995), ISSN 1538-4357.
- [81] E. Aprile (for the The XENON1T Collaboration), presentation at [LNGS Scientific Committee Meeting April 2015 \(2015\)](#).
- [82] D. C. Malling, et al., presentation at [arXiv:1110.0103 \[astro-ph\] \(2011\)](#), arXiv: 1110.0103.
- [83] D. S. Akerib, et al. (The LUX Collaboration), *First Results from the LUX Dark Matter Experiment at the Sanford Underground Research Facility*, [Phys. Rev. Lett.](#), **112**, 091303 (2014).
- [84] E. Aprile, et al. (The XENON100 Collaboration), *Dark Matter Results from 225 Live Days of XENON100 Data*, [Phys. Rev. Lett.](#), **109**, 181301 (2012).

- [85] M. G. Boulay and t. D. Collaboration, *DEAP-3600 Dark Matter Search at SNOLAB*, [Journal of Physics: Conference Series](#), **375**, 012027 (2012), ISSN 1742-6596.
- [86] J. Billard, E. Figueroa-Feliciano, and L. Strigari, *Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments*, [Phys. Rev. D](#), **89**, 023524 (2014).
- [87] Southwest Stainless & Alloy, [ALLOY INFORMATION AND SPECS](#).
- [88] S. A. M. Brand, [The Red Book](#).
- [89] D. M. Mei, Z. B. Yin, L. C. Stonehill, and A. Hime, *A model of nuclear recoil scintillation efficiency in noble liquids*, [Astroparticle Physics](#), **30**, 12 (2008), ISSN 0927-6505.
- [90] T. Doke, et al., *Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles*, [Japanese Journal of Applied Physics](#), **41**, 1538 (2002), ISSN 0021-4922, 1347-4065.
- [91] J. Lindhard, V. Nielsen, M. Scharff, and P. V. Thomsen, *Integral equations governing radiation effects*, [Det Kgl. Danske Viden.](#), **33**, 10:1 (1963).
- [92] H. Cao, et al. (The SCENE Collaboration), *Measurement of scintillation and ionization yield and scintillation pulse shape from nuclear recoils in liquid argon*, [Phys. Rev. D](#), **91**, 092007 (2015).
- [93] A. Hitachi, T. Doke, and A. Mozumder, *Luminescence quenching in liquid argon under charged-particle impact: Relative scintillation yield at different linear energy transfers*, [Phys. Rev. B](#), **46**, 11463 (1992).
- [94] D. C. Lorents, *The physics of electron beam excited rare gases at high densities*, [Physica B+C](#), **82**, 19 (1976), ISSN 0378-4363.
- [95] P. Sorensen, *Atomic limits in the search for galactic dark matter*, [Physical Review D](#), **91**, 083509 (2015).
- [96] J. F. Ziegler, M. D. Ziegler, and J. P. Biersack, *SRIM—The stopping and range of ions in matter (2010)*, [Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms](#), **268**, 1818 (2010).
- [97] J. B. Birks and F. A. Black, *Deterioration of Anthracene under α -Particle Irradiation*, [Proceedings of the Physical Society. Section A](#), **64**, 511 (1951), ISSN 0370-1298.
- [98] J. Xu, et al., *A study of the trace ^{39}Ar content in argon from deep underground sources*, [Astropart. Phys.](#), **66**, 53 (2015).
- [99] H. O. Back, et al., *First Large Scale Production of Low Radioactivity Argon From Underground Sources*, [arXiv:1204.6024v2](#) (2012).

- [100] P. Benetti, et al. (The WArP Collaboration), *Measurement of the specific activity of ^{39}Ar in natural argon*, [Nucl. Inst. Meth. A](#), **574**, 83 (2007).
- [101] W. M. Haynes, [CRC Handbook of Chemistry and Physics](#), CRC Press, LLC (2015), 96th ed.
- [102] A. Santamarina, et al., *The jeff-3.1. 1 nuclear data library*, [JEFF report](#), **22**, 2 (2009).
- [103] W. Maneschg, et al., *Measurements of extremely low radioactivity levels in stainless steel for GERDA*, [Nucl. Instr. Meth. Phys. Res. A](#), **593**, 448 (2008), ISSN 0168-9002.
- [104] International Atomic Energy Agency, [Mexico informs IAEA of theft of dangerous radioactive source](#) (2015).
- [105] Petco, [Recalls: Stainless steel bowl notice](#) (2012).
- [106] M. Wieser, *Atomic weights of the elements 2005 (iupac technical report)*, [Pure and Applied Chemistry](#), **78**, 2051 (2006).
- [107] O. Serot, C. Wagemans, and J. Heyse, presentation at [AIP Conference Proceedings](#) (2005), vol. 769, pp. 857–860.
- [108] E. S. Kim, C. H. Oh, and M. Patterson, *Icapp 2011*, , (2011).
- [109] G. Alimonti, et al. (The Borexino Collaboration), *Measurement of the ^{14}C abundance in a low-background liquid scintillator*, [Phys. Lett. B](#), **422**, 349 (1998).
- [110] F. P. Calaprice and B. R. Holstein, *Weak magnetism and the beta spectra of ^{12}B and ^{12}n* , [Nuclear Physics A](#), **273**, 301 (1976), ISSN 0375-9474.
- [111] V. V. Kuzminov and N. J. Osetrova, *Precise measurement of ^{14}C beta spectrum by using a wall-less proportional counter*, [Physics of Atomic Nuclei](#), **63**, 1292 (2000), ISSN 1063-7788, 1562-692X.
- [112] H. Behrens, et al., *Allowed β -transitions, weak magnetism and nuclear structure in light nuclei*, [Annals of Physics](#), **115**, 276 (1978), ISSN 0003-4916.
- [113] E. Aprile, et al., *Lowering the radioactivity of the photomultiplier tubes for the XENON1t dark matter experiment*, [The European Physical Journal C](#), **75**, 1 (2015), ISSN 1434-6044, 1434-6052.
- [114] D. Wilkinson, *Evaluation of beta-decay*, [Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment](#), **275**, 378 (1989), ISSN 01689002.
- [115] D. Wilkinson, *Evaluation of beta-decay: II. Finite mass and size effects*, [Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment](#), **290**, 509 (1990), ISSN 01689002.

- [116] D. Wilkinson, *Evaluation of beta-decay Part III. The complex gamma function*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, **335**, 305 (1993), ISSN 01689002.
- [117] D. Wilkinson, *Evaluation of beta-decay Part IV. The complex gamma function; practicalities*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, **365**, 203 (1995), ISSN 01689002.
- [118] D. Wilkinson, *Evaluation of beta-decay Part V. The Z-independent outer radiative corrections for allowed decay*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, **365**, 497 (1995), ISSN 01689002.
- [119] D. Wilkinson, *Evaluation of beta-decay, Part VI: The Z-dependent outer radiative corrections for allowed decay*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, **401**, 275 (1997), ISSN 01689002.
- [120] D. Wilkinson, *Evaluation of beta-decay, Part VII: The Z-independent outer radiative correction for unique-forbidden decay*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, **406**, 89 (1998), ISSN 01689002.
- [121] N. Mott and H. Massey, *The theory of atomic collisions*, clarendon (1965).
- [122] M. J. Berger, J. Coursey, M. Zucker, and J. Chang, *Stopping-power and range tables for electrons, protons, and helium ions*, NIST Physics Laboratory (1998).
- [123] F. Kondev, *Nuclear data sheets for a= 206*, Nuclear Data Sheets, **109**, 1527 (2008).
- [124] H. Paul, presentation at AIP Conference Proceedings (2013), vol. 1525, pp. 309–313.
- [125] P. Sigmund and A. Schinner, *Shell correction in stopping theory*, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms, **243**, 457 (2006), ISSN 0168-583X.
- [126] T. Pollmann, M. Boulay, and M. Kuniak, *Scintillation of thin tetraphenyl butadiene films under alpha particle excitation*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, **635**, 127 (2011), ISSN 0168-9002.
- [127] G. Bellini, et al. (The Borexino Collaboration), *Muon and cosmogenic neutron detection in Borexino*, JINST, **6**, P05005 (2011), ISSN 1748-0221.

- [128] G. Bellini, et al. (The Borexino Collaboration), *Cosmogenic Backgrounds in Borexino at 3800 m water-equivalent depth*, [JCAP](#), **2013**, 049 (2013).
- [129] M. Ambrosio, et al. (The MACRO Collaboration), *Measurement of the residual energy of muons in the Gran Sasso underground laboratories*, [Astroparticle Physics](#), **19**, 313 (2003), ISSN 0927-6505.
- [130] M. G. Boulay (The DEAP Collaboration), *DEAP-3600 Dark Matter Search at SNOLAB*, [J. Phys.: Conf. Ser.](#), **375**, 012027 (2012).
- [131] F. E. Gray, et al., *Cosmic ray muon flux at the Sanford Underground Laboratory at Homestake*, [Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment](#), **638**, 63 (2011), ISSN 0168-9002.
- [132] C. Galbiati, et al., *Cosmogenic ^{11}C production and sensitivity of organic scintillator detectors to pep and CNO neutrinos*, [Phys. Rev. C](#), **71**, 055805 (2005).
- [133] N. Agafonova, et al. (The OPERA Collaboration), *Measurement of the atmospheric muon charge ratio with the OPERA detector*, [The European Physical Journal C](#), **67**, 25 (2010), ISSN 1434-6044, 1434-6052.
- [134] H. Schopper, ed., *Shielding Against High Energy Radiation*, vol. 11 of *Landolt-Brnstein - Group I Elementary Particles, Nuclei and Atoms*, [Springer-Verlag \(Berlin/Heidelberg \(1990\)\)](#), ISBN 978-3-540-52125-9.
- [135] P. Catalano, G. Cavinato, F. Salvini, and M. Tozzi, *Analisi strutturale nei laboratori dell'infn del gran sasso d'italia*, [MEMORIE DELLA SOCIETA'GEOLOGICA ITALIANA](#), **35**, 647 (1986).
- [136] M. Birch, et al., *Evaluation of Beta-Delayed Neutron Emission Probabilities and Half-Lives for $Z = 228$* , [Nuclear Data Sheets](#), **128**, 131 (2015), ISSN 0090-3752.
- [137] T. Bjrnstad, et al., *The decay of 8He* , [Nuclear Physics A](#), **366**, 461 (1981), ISSN 0375-9474.
- [138] I. Mukha, et al., *Observation of the $^{11}\text{Li}(d)$ decay*, [Physics Letters B](#), **367**, 65 (1996), ISSN 0370-2693.
- [139] A. Empl and E. V. Hungerford, *A FLUKA Study of β -delayed Neutron Emission for the Ton-size DarkSide Dark Matter Detector*, [arXiv:1407.6628 \[astro-ph, physics:physics\]](#), (2014), arXiv: 1407.6628.
- [140] D. E. Alburger and D. R. Goosman, *Delayed neutrons in the decay of ^{13}B* , [Physical Review C](#), **10**, 935 (1974).
- [141] D. E. Alburger and D. H. Wilkinson, *Beta decay of ^{16}C and ^{17}N* , [Physical Review C](#), **13**, 835 (1976).

- [142] D. Halliday et al., *Introductory nuclear physics* Wiley (1955).
- [143] D. Clark, *U235: a gamma ray analysis code for uranium isotopic determination*, (1997) .
- [144] H. Lemmel, *Physics of neutron emission in fission*, (1989) .
- [145] T. E. Valentine, *Evaluation of prompt fission gamma rays for use in simulating nuclear safeguard measurements*, *Annals of Nuclear Energy*, **28**, 191 (2001), ISSN 03064549.
- [146] M. J. Carson, et al., *Neutron background in large-scale xenon detectors for dark matter searches*, *Astroparticle Physics*, **21**, 667 (2004), ISSN 0927-6505.
- [147] A. Koning, S. Hilaire, and S. Goriely, *Talys-1.6*, *Nuclear Reaction Program*, (2013).
- [148] K. Shibata, et al., *Japanese evaluated nuclear data library version 3 revision-3: Jendl-3.3*, *Journal of Nuclear Science and Technology*, **39**, 1125 (2002).
- [149] D. Mei, C. Zhang, and A. Hime, *Evaluation of (α, n) induced neutrons as a background for dark matter experiments*, *Nucl. Inst. Meth. A*, **606**, 651 (2009).
- [150] P. De Bivre and P. D. P. Taylor, *Table of the isotopic compositions of the elements*, *International Journal of Mass Spectrometry and Ion Processes*, **123**, 149 (1993), ISSN 0168-1176.
- [151] J. K. Tuli, *Evaluated nuclear structure data file*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **369**, 506 (1996), ISSN 0168-9002.
- [152] N. A. Roughton, et al., *Thick-target measurements and astrophysical thermonuclear reaction rates: Alpha-induced reactions*, *Atomic Data and Nuclear Data Tables*, **28**, 341 (1983), ISSN 0092-640X.
- [153] Materion Brush Performance Alloys, *PROthermTM Alloy Data Sheet*.
- [154] Carpenter Technology Corporation, *Kovar[®] Alloy*.
- [155] T. I. C. of America, *Solder Quotation. Solder SAC 305*, private communication.
- [156] Pure Technologies, *ALPHA-Lo[®] Material Solder* (2016).
- [157] A. Wright, P. Mosteiro, B. Loer, and F. P. Calaprice, *A highly efficient neutron veto for dark matter experiments*, *Nucl. Inst. Meth. A*, **644**, 18 (2011).
- [158] J. H. Hubbell and S. M. Seltzer, *Tables of X-ray mass attenuation coefficients and mass energy-absorption coefficients 1 keV to 20 MeV for elements Z= 1 to 92 and 48 additional substances of dosimetric interest*, (1995) .

- [159] J. B. Birks, *Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations*, [Proceedings of the Physical Society. Section A](#), **64**, 874 (1951), ISSN 0370-1298.
- [160] J. B. Birks and F. W. K. Firk, *The Theory and Practice of Scintillation Counting*, [Physics Today](#), **18**, 60 (1965), ISSN 0031-9228.
- [161] M. A. Islam, T. J. Kennett, and W. V. Prestwich, *Thermal neutron capture in silicon*, [Physical Review C](#), **41**, 1272 (1990), ISSN 0556-2813, 1089-490X.
- [162] G. Alimonti, et al. (The Borexino Collaboration), *A large-scale low-background liquid scintillation detector: the counting test facility at Gran Sasso*, [Nucl. Instr. Meth. A](#), **406**, 411 (1998).
- [163] The Daya Bay Collaboration, *The muon system of the Daya Bay Reactor antineutrino experiment*, [Nucl. Instr. Meth. Phys. Res. A](#), **773**, 8 (2015), ISSN 01689002, arXiv: 1407.0275.
- [164] F. Elisei, et al., *Measurements of liquid scintillator properties for the Borexino detector*, [Nucl. Instr. Meth. Phys. Res. A](#), **400**, 53 (1997), ISSN 0168-9002.
- [165] Bal Seal Engineering, Inc., *Chemical Compatibility Chart: Metallic Materials Used in Bal Seal Products*, TR-60C, (1999) .
- [166] Labsphere, *A Guide to Reflectance Materials and Coatings*, 1, (2015) .
- [167] M. Janecek, *Reflectivity Spectra for Commonly Used Reflectors*, [IEEE Trans. Nucl. Sci.](#), **59**, 490 (2012), ISSN 0018-9499.
- [168] B. Loer, [daqman](#) (2014).
- [169] J. Xu, et al., *Scintillation efficiency measurement of Na recoils in NaI(Tl) below the DAMA/LIBRA energy threshold*, [Physical Review C](#), **92**, 015807 (2015).
- [170] S. Agostinelli, et al., *Geant4—a simulation toolkit*, [Nucl. Instr. Meth. A](#), **506**, 250 (2003).
- [171] S. Sjøe, *Relativistic two-body reaction kinematics* (2010).
- [172] R. L. Craun and D. L. Smith, *Analysis of response data for several organic scintillators*, [Nuclear Instruments and Methods](#), **80**, 239 (1970), ISSN 0029-554X.
- [173] S. Westerdale, et al., *The DarkSide-50 outer detectors*, [Journal of Physics: Conference Series](#), **718**, 042062 (2016), ISSN 1742-6596.
- [174] P. Agnes, et al., *The Electronics and Data Acquisition System for the DarkSide-50 Veto Detectors*, [arXiv:1606.03316 \[hep-ex, physics:physics\]](#), (2016), arXiv: 1606.03316.

- [175] CAEN, [A1536 and AG536 User Manual](#), (2015) .
- [176] CAEN, [SY4527 and SY4527LC User Manual](#), (2015) .
- [177] C. Green, et al., *The art framework*, [JPCS](#), **396**, 022020 (2012), ISSN 1742-6588, 1742-6596.
- [178] G. Bellini, et al. (The Borexino Collaboration), *Cosmic-muon flux and annual modulation in Borexino at 3800 m water-equivalent depth*, [JCAP](#), **2012**, 015 (2012).
- [179] Hentzen, [Sterilkote®](#), website (2014).
- [180] Easterline Souriau, [Underwater connector: M series](#), website.
- [181] W.L. Gore & Associates, [Tetra-etch® fluorocarbon etchant](#), website (2015).
- [182] G. Alimonti, et al. (The Borexino Collaboration), *The liquid handling systems for the Borexino solar neutrino detector*, [Nucl. Instr. Meth. Phys. Res. A](#), **609**, 58 (2009), ISSN 0168-9002.
- [183] J. Benziger, et al., *The fluid-filling system for the Borexino solar neutrino detector*, [Nucl. Inst. Meth. A](#), **608**, 464 (2009).
- [184] B. E. Lehmann, S. N. Davis, and J. T. Fabryka-Martin, *Atmospheric and sub-surface sources of stable and radioactive nuclides used for groundwater dating*, [Water Resour. Res.](#), **29**, 2027 (2010).
- [185] K. O. Al-Ahmad and D. E. Watt, *Stopping power of low energy electrons (<10 keV) in solid polyethylene*, [Journal of Physics D: Applied Physics](#), **17**, 1899 (1984), ISSN 0022-3727.
- [186] J. Liu, et al., *Neutron calibration sources in the Daya Bay experiment*, [Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment](#), **797**, 260 (2015), ISSN 01689002.