

TOWARDS A DEPLETED ARGON TIME
PROJECTION CHAMBER WIMP SEARCH :
DARKSIDE PROTOTYPE ANALYSIS AND
PREDICTED SENSITIVITY

BEN MICHAEL LOER

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: CRISTIANO GALBIATI

NOVEMBER 2011

© Copyright by Ben Michael Loer, 2011.

All rights reserved.

Abstract

Scientists have now accumulated overwhelming evidence indicating that over 80% of the mass of the universe is in the form of dark matter, neutral particles with ultra-weak couplings to ordinary matter. One compelling candidate is a Weakly Interacting Massive Particle or WIMP, with mass on the order of 100 GeV. The signal of a WIMP interaction in a detector is a low energy (typically below ~ 100 keV) recoiling atomic nucleus. The expected rate is at most a few interactions per year per ton of target. The most critical issue for direct WIMP searches is reducing the background in the detector below this very low rate. Argon is a promising target because measurements of the scintillation pulse shape and the scintillation-to-ionization ratio allow the reduction of gamma-induced signals, the largest source of background, by a factor of 10^8 or better. One of the major drawbacks of argon is the presence of radioactive ^{39}Ar , which results in a decay rate of ~ 1 Bq/kg in natural argon. Because ^{39}Ar is produced primarily in the upper atmosphere, the ^{39}Ar fraction can be reduced significantly by obtaining the argon from underground wells.

Our collaboration, DarkSide, is developing a series of two-phase argon time projection chambers (TPCs) utilizing this depleted argon, along with passive shielding and active neutron and muon vetoes, for WIMP searches. I present results from a recent campaign of a 10 kg active mass prototype TPC that demonstrate the successful realization of many of the technical aspects necessary for a full-scale detector, in particular an electron-equivalent light yield of 4.5 photoelectrons per keV deposited, and a free electron lifetime in excess of 200 microseconds. Based on this successful prototype and Monte Carlo simulations, I then conclude that DarkSide-50, a 50 kg active mass TPC to be installed in LNGS, can likely acquire data background-free for three years, accumulating a fiducial exposure of ~ 100 kg-years and reaching a sensitivity to the WIMP-nucleon cross section of $\sim 3 \times 10^{-45}$ cm².

I also present some details on the analysis of Borexino data that resulted in the first real-time, spectroscopic measurement of ^7Be solar neutrino interactions.

Acknowledgements

First I would like to thank my adviser, Cristian Galbiati, as well as professors Peter Meyers and Frank Calaprice. Special thanks goes to Peter for his helpful and very detailed feedback on this work. I have been extremely fortunate to have the experience of these three distinguished professors to call upon, and their combined influence has made my stay at Princeton both enjoyable and productive. Moreover, they have assembled a group of fellow students whom I am glad to have had the opportunity to work with. My work has certainly proceeded much faster thanks to the willingness of my office-mates to interrupt their work and consider any problem I may have run into, and I hope that I have managed to be similarly helpful for them.

I would also like to acknowledge those whose efforts directly contributed to the work presented here. The design and construction of the DarkSide-10 detector was a complex, coordinated effort, and trying to name everyone involved would be a futile effort. In particular, the constant help and advice of Allan Nelson have been invaluable for this and every other hardware project undertaken. We are all indebted to Alex Wright and Tristen Hohman for their hours of tedious testing of the mechanical stability of detector components under cryogenic temperatures. The electric field apparatus was designed primarily by our collaborators Jeff Martoff, Christy Martin, and John Tatarowicz at Temple University. The procedures and algorithms used to analyze the DarkSide-10 data were developed with much help from Luca Grandi, Huajie Cao, and Jason Brodsky, and the results presented in Chapter 4 were achieved in great part due to their efforts along with those of Pablo Mosteiro, Emily Shields, and Alden Fan, who all carried out crucial analysis tasks and helped to draft the text describing the results. The Monte Carlo analysis of backgrounds would not be nearly so well-developed were it not for the efforts of Alex Wright and Pablo Mosteiro, both in compiling exhaustive lists of radioactive contaminations of a whole host of materials and in developing my initial software into a much more flexible and powerful package. The design, construction, and operation of Borexino are the result of a huge number of hours of

hard work on the part of a large team of dedicated scientists and engineers whom I could not hope to name in full. The analysis presented in Appendix A in particular is the result of the efforts of fellow Princeton students Kevin McCarty, Richard Saldanha, Alvaro Chavarria, Jingke Xu as well as collaborators from Virginia Tech and the University of Massachusetts at Amherst. I would also like to thank Alessandro Razeto, who worked to teach me the subtleties of the low-level Borexino data decoding.

Finally, but certainly not least, I would like to thank my friends and family. If anyone can get through the stress of everyday life, let alone graduate school, without the support of friends, it is surely not me. My parents have always been a font of guidance and encouragement, and never settled for anything less than the maximum of my potential; I would never have found myself in the place I am today without their influence. Finally, to my wife, Rachel, both family and best friend, who has shared every step of the way with me: thank you for your unending supply of love and patience, and for forcing me to keep writing during all those terrible late nights. I look forward to repaying the favor during your dissertation work.

Parts of the results of the measurement of ^{39}Ar levels in underground argon in Chapter 3 have been previously published in Ref. [1], and presented as a poster at the Neutrino 2008 conference in Christchurch, New Zealand [2]. Some of the results from the analysis of DarkSide-10 data in Chapter 4 were presented in a contributed talk at the 2011 APS April Meeting in Anaheim, CA [3]. The results of the measurement of the ^7Be solar neutrino rate with Borexino in Appendix A were previously published in Ref. [4].

This work was supported by grants 1004072, 0919363, 0811186, 0802646, and 0704220 from the National Science Foundation, and grant DE-FG02-91ER40671 from the Department of Energy.

To my wonderful wife, Rachel.

Contents

Abstract	iii
Acknowledgements	iv
List of Tables	xi
List of Figures	xiii
1 Introduction	1
1.1 The case for dark matter	1
1.1.1 Galactic Rotation Curves	2
1.1.2 The Cosmic Microwave Background	3
1.1.3 Big Bang Nucleosynthesis	8
1.2 Expected properties of dark matter	11
1.2.1 Dark Matter Candidates	11
1.2.2 Expected WIMP Direct Detection Signals	19
1.3 Current status of dark matter searches	34
1.3.1 Indirect detection	34
1.3.2 Direct detection	38
2 Detecting Dark Matter: The DarkSide Program	40
2.1 Backgrounds to direct dark matter searches	40
2.2 Argon as a dark matter target	43
2.3 Liquid Argon TPC	47

2.4	Background Rejection in an argon TPC	51
2.4.1	Surface alpha decays	51
2.4.2	Radiogenic Neutrons	53
2.4.3	Cosmogenic Muons and Neutrons	57
2.4.4	β/γ Radiation	59
2.4.5	Backgrounds Summary	66
3	Depleted Argon	68
3.1	The need for underground argon	68
3.2	Pressure Swing Adsorption	70
3.3	Prototype PSA Development	77
3.4	Prototype Testing	80
3.5	Oxygen Purification	80
3.6	Argon Extraction From Crude Helium	84
3.7	Measurements of ^{39}Ar fraction	86
3.8	Full-scale PSA and cryogenic distillation	89
4	10 kg Prototype	93
4.1	Detector description	96
4.2	Data taking conditions	101
4.2.1	Recirculation	101
4.2.2	Photomultipliers, DAQ, and trigger	102
4.2.3	Data sample	107
4.3	Low Level Analysis Description	110
4.3.1	ConvertData	110
4.3.2	SumChannels	110
4.3.3	BaselineFinder	111
4.3.4	EvalRois	116

4.3.5	PulseFinder	116
4.3.6	S1S2Evaluation	121
4.3.7	PositionRecon	121
4.4	Calibration and Efficiency of Algorithms and Cuts	122
4.4.1	Single Photoelectron Fit	122
4.4.2	Analysis Cuts	125
4.4.3	Energy and Fprompt variables	130
4.4.4	Moving Baseline	138
4.5	High Level Analysis	141
4.5.1	Light Yield	141
4.5.2	F90	156
4.5.3	Ionization (S2)	166
4.6	Conclusions	193
5	Monte Carlo Simulations	198
5.1	Full-scale DarkSide detectors: DarkSide-50	199
5.2	Simulation Details	202
5.2.1	Monte Carlo Software	202
5.2.2	Radioactive Background Rates and Spectra	210
5.3	Simulation Results	220
5.3.1	Analysis Procedure	220
5.3.2	Gammas	224
5.3.3	Radiogenic Neutrons	226
5.3.4	Cosmogenic Neutrons	228
5.4	Predicted Sensitivity	228
5.5	Conclusions	234

A	Measurement of the ^7Be solar neutrino flux with Borexino	235
A.1	Low level calibration	236
A.1.1	Laben Pulser Calibration	238
A.1.2	Laser and random calibration	246
A.1.3	Channel and Hit Rejection	247
A.2	Event reconstruction	250
A.2.1	Clustering	250
A.2.2	Energy reconstruction	251
A.2.3	Position Reconstruction	253
A.3	Cuts and Event Selection	256
A.3.1	Alpha-beta discrimination and the Gatti parameter	259
A.3.2	Radon daughter removal	260
A.3.3	Statistical alpha subtraction	261
A.3.4	^{85}Kr Tagging	263
A.4	Spectral Fit	264
A.4.1	^7Be Interaction Spectrum	265
A.4.2	Energy Response Function	268
A.4.3	Fit results	271
B	daqman File Format	273
B.1	Legacy format	273
B.2	New format	274
	Bibliography	276

List of Tables

1.1	Galactic halo distribution parameters	29
3.1	Optimum PSA cycle for oxygen purification with Drierite bottles	82
3.2	Optimum PSA cycle for oxygen purification using 10' columns	84
3.3	Crude helium composition	85
4.1	PMT Characteristics	105
4.2	Statistics for focus run set.	109
4.3	Description of standard cuts for DS-10 analysis	127
4.4	Events cut in standard dataset runs	128
4.5	Cut rejection for runs with different field settings	130
4.6	Two-pulse Cut rejection for events with different field settings	131
4.7	Zero-field ^{22}Na fit results	143
4.8	Zero field ^{22}Na fits with all PMTs enabled	145
4.9	Light yield vs height	149
4.10	Comparison of the distribution of $F90$ for 50 p.e. wide bins from data and Monte Carlo.	165
4.11	Electron lifetime fits from run 1460	186
4.12	Electron lifetime fit results	193
5.1	DarkSide-50 Simulation geometry	204
5.2	Gammas emitted per decay of parent isotopes	206

5.3	Fit parameters as a function of site depth used in Equation (5.2) [5]	208
5.4	Assumed radioactive contamination rates of materials for DarkSide-50	213
5.5	Conversion between decay rate and concentration for ^{238}U , ^{232}Th , and ^{40}K [6].	214
5.6	(α, n) decay rates for detector materials	215
5.7	Rate of fission neutrons emitted per equilibrium decay of uranium and thorium, calculated from SOURCES4A.	215
5.8	Rate of cosmogenic neutrons produced in detector shielding materials.	215
5.9	Total background rates in detector materials	219
5.10	Simulated gamma background rates in DarkSide-50	225
5.11	Simulated radiogenic neutron background rates in DarkSide-50.	227
5.12	Simulated cosmogenic neutron background rates in DarkSide-50	229
A.1	Code used to determine gray counter shift	245
A.2	Fixed offset times applied to each crate in calculating hit times.	246

List of Figures

1.1	Galaxy rotation curves	4
1.2	WMAP CMB measurements	6
1.3	Big bang nucleosynthesis	10
1.4	Dark matter freeze-out process	17
1.5	Integration of WIMP velocity in Earth frame	25
1.6	WIMP velocity distribution in Earth frame	27
1.7	Differential WIMP-induced nuclear recoil spectrum	30
1.8	Integrated WIMP-induced nuclear recoil spectrum	31
1.9	Spectrum of annual modulation amplitude of WIMP-induced nuclear recoils	32
1.10	Integrated annual modulation of WIMP-induced nuclear recoils	33
1.11	WIMP detection processes	34
1.12	IceCube cross section sensitivity	35
1.13	High energy cosmic ray excesses measured by PAMELA, ATIC, Fermi, and HESS	36
1.14	WIMP annihilation cross section obtained from PAMELA and FERMI data	37
1.15	Current experimental WIMP direct detection cross section sensitivities. . . .	39
2.1	Argon scintillation quenching with applied electric field	42
2.2	Time projection chamber operation	48
2.3	Electron drift speed in liquid argon	50
2.4	Hamamatsu HR3998 PMT neutron emission spectrum	54

2.5	Total cross section for neutron scattering off hydrogen [7].	54
2.6	Neutron flux in underground sites	58
2.7	Neutron spallation rates for various nuclei	59
2.8	Variation of mean F_{90} vs energy	62
2.9	Event acceptance in nuclear recoil window based on ideal binomial distribution	64
2.10	Gamma leakage vs nuclear recoil acceptance	65
2.11	Ideal binomial F_{prompt} compared to Monte Carlo distribution	66
3.1	The adsorption isotherms of nitrogen, oxygen, and argon on zeolite NaX [8].	72
3.2	Example breakthrough curve	74
3.3	Basic six step PSA cycle employing pressurization with feed [9].	75
3.4	Idealized recovery of light component in PSA cycle	77
3.5	Schematic depiction of the prototype PSA apparatus.	78
3.6	Prototype PSA control software	80
3.7	The prototype PSA system, shown here partially disconnected after maintenance.	81
3.8	Asymmetric PSA output during oxygen separation	83
3.9	Crude helium PSA breakthrough curve	86
3.10	Gas proportional counters used for depleted argon measurement	87
3.11	Schematic depiction of the Princeton low-background liquid argon counter.	89
3.12	Process and instrumentation diagram of the two-stage industrial VPSA plant operating at the Doe Canyon site.	90
3.13	Schematic diagram of the cryogenic distillation column at FNAL.	92
4.1	DS-10 Cartoon Schematic	97
4.2	Photograph of the 10kg prototype inner detector during assembly.	98
4.3	DS-10 Assembly Diagram	98

4.4	Example event, showing the digitized signal from channel 7 for a gamma event that deposited ~ 65 keV in the active volume.	103
4.5	Diagram of the data acquisition and trigger setup for the second DarkSide-10 campaign.	106
4.6	Photograph of the diode laser, filters, and optical fiber used to calibrate the PMT response.	106
4.7	A typical scintillation event	111
4.8	Illustration of baseline finding algorithm. See text for details.	112
4.9	Summed signal waveform showing baseline oscillation and interpolation over S2 region.	115
4.10	Example trigger illustrating PulseFinder	119
4.11	Example laser response fit	124
4.12	Difference between single photoelectron PDFs	125
4.13	Fraction of events cut vs drift field	129
4.14	Fraction of events cut vs extraction field	129
4.15	Distribution of S1 and S2 pulse lengths	134
4.16	S1 fixed vs. S1 full	135
4.17	Integral of sums versus sum of integrals of the S1 pulse for events in ^{22}Na coincidence runs passing cuts.	136
4.18	S1 offset between IoS and SoI	137
4.19	$F90$ offset between IoS and SoI	138
4.20	Laser spectrum of events with moving baseline	139
4.21	Integral with fixed baseline vs moving baseline, with linear fit.	141
4.22	Zero field ^{22}Na S1 spectrum showing cuts	142
4.23	Zero-field ^{22}Na fit	142
4.24	Zero-field ^{22}Na fit, all tubes enabled	146
4.25	^{133}Ba spectrum and fit	147

4.26	^{22}Na spectrum with field	148
4.27	Light yield vs height	150
4.28	S1 bottom fraction	151
4.29	S1 bottom fraction vs drift time	151
4.30	Channel 1 laser occupancy vs time	152
4.31	Average quantum efficiency	153
4.32	S1 bottom fraction, no gas pocket	155
4.33	F_{90} Distributions	157
4.34	F_{90} vs S1, with and without drift field.	158
4.35	Mean F_{90} vs energy with exponential fit	159
4.36	Mean F_{90} vs energy, low energy	159
4.37	Electron stopping power in argon, obtained form NIST stopping power database [10].	160
4.38	F_{90} vs. stopping power	160
4.39	F_{prompt} estimate with no cuts	161
4.40	Mean F_{90} vs S1 for individual channels	162
4.41	F_{90} vs S1 for AmBe runs with and without drift field	163
4.42	F_{90} distribution with field compared to Monte Carlo	164
4.43	F_{90} distribution without field compared to Monte Carlo	165
4.44	Drift time distribution from run 1354.	166
4.45	Drift time vs source height	167
4.46	Electron drift speed in liquid argon vs. applied electric field [11].	168
4.47	Background drift time distribution	168
4.48	Collimated ^{22}Na trigger rate vs source height	169
4.49	^{22}Na coincidence drift time distribution	169
4.50	Setup of source and NaI for coincidence triggers	172
4.51	Photo of coincidence setup	172

4.52	^{22}Na coincidence trigger rate over time	173
4.53	Drift time distributions for various data acquisition periods	176
4.54	$\text{Log}_{10}(\text{S2})$ vs S1	177
4.55	S2/S1 Distribution	178
4.56	S2/S1 ^{22}Na leakage fraction	178
4.57	Drift-time correction for S2/S1	179
4.58	S2/S1 with $F90$ cuts	181
4.59	Fraction of events with low S2/S1 vs $F90$	181
4.60	S2/S1 vs $F90$	182
4.61	S1 vs S2 vs $F90$	182
4.62	S2 vs $F90$ vs S1	183
4.63	S2/S1 vs drift time	184
4.64	Example S2/S1 distributions from ^{22}Na runs in two drift time bins: 60–70 μs and 110–120 μs	185
4.65	Lifetime corrected for S1 Z-dependence	187
4.66	Electron lifetime vs time of run	189
5.1	DarkSide-50 Drawings	200
5.2	Cartoon representation of the detector volumes specified in <code>s4det</code> . Not to scale.	205
5.3	Gamma spectrum from ^{238}U , ^{235}U , and ^{232}Th decay chains	207
5.4	Decay chains for ^{238}U , ^{235}U and ^{232}Th , the most common contributors to intrinsic radioactive backgrounds [12].	211
5.5	Spectrum of cosmogenic neutrons assumed for LNGS.	216
5.6	Zenith angular distribution of cosmogenic neutrons (relative to straight down) assumed for LNGS.	216
5.7	(α, n) neutron spectrum in Hall C of LNGS	218
5.8	Quenching in liquid scintillator	223
5.9	Beta decay spectrum of ^{39}Ar	224

5.10	Likelihood for total neutron rate from Monte Carlo results	231
5.11	DarkSide-50 WIMP 3 year WIMP interaction sensitivity	233
A.1	Determining the width of the 10 MHz triangle wave in ADC counts	241
A.2	Value of ramp2 calculated from ramp1	242
A.3	Number of valid channels (separated into good timing and good charge reconstruction) for runs 7000-10000.	249
A.4	Comparison of the two energy variables in mach4, npe and nhits, after scaling and statistical correction. Thanks to A. Chavarria.	253
A.5	Fit of the ^{210}Po alpha peak over time, showing the stability of the nhits_corrected variable.	254
A.6	Effective scintillator emission PDFs for 1, 2 and ≥ 3 detected photoelectrons, including the smearing from the detector response function.	256
A.7	Expected spectrum from solar neutrinos in Borexino.	257
A.8	Alpha and beta scintillation pulse shapes	259
A.9	Spectrum of gatti parameter vs energy for events in the fiducial volume.	261
A.10	Example fits of gatti parameter distributions	262
A.11	Mean gatti parameter vs energy	263
A.12	Effect of cuts on Borexino spectrum	264
A.13	Predicted flux and spectrum from solar neutrinos according to the BS05(Op) model [13].	265
A.14	Neutrino-electron elastic scattering diagrams.	266
A.15	dE/dx curve for electrons in pseudocumene, the Borexino liquid scintillator.	269
A.16	Results of fit to spectrum of 41.3 ton-years of Borexino data [4].	271
B.1	Raw data format for a trigger from a single CAEN V1720 digitizer [14].	274

Chapter 1

Introduction

1.1 The case for dark matter

Over the past century, observers have accumulated compelling, and even overwhelming, evidence that indicates the existence of dark matter: some particle or set of particles beyond the standard model that acts on standard matter only through gravitation and possibly other ultra-weak processes. To thoroughly cover the history of these astronomical observations would itself be the work of an entire dissertation, and I will not attempt to do so here. Rather, in this section I will present a very small subset of the most compelling evidence thus far collected, and those measurements that have a direct bearing on parameters that will be required in later sections in order to predict the types of signatures we might observe from a direct interaction with dark matter particles. For a more in-depth review, the reader should consult references [15–18].

In this and the following sections, when discussing mass densities, it is convenient to work in units of the critical density:

$$\rho_{cr} = \frac{3H_0^2}{8\pi G} \tag{1.1}$$

where H_0 is Hubble's parameter, which determines the rate of the expansion of space today, and G is Newton's constant, ($6.67 \times 10^{-8} \text{ cm}^3 \text{g}^{-1} \text{sec}^{-2}$). The critical density is the amount of energy density that results in a flat universe, i.e., one with zero net curvature of space. In a flat spacetime, rays that start out parallel will remain that way. If the density is less than the critical density, then the universe is open, with a negative curvature, and initially parallel rays will eventually converge. Conversely, if the density is greater than the critical density, the universe is closed, with positive curvature, and initially parallel rays will diverge. The current best fit for the Hubble parameter is $70.4_{-1.4}^{+1.3} \text{ km s}^{-1} \text{ Mpc}^{-1}$ [19], so $\rho_{cr} = 9.32 \times 10^{-30} \text{ g/cm}^3$. Then, for a species i , we define the variable

$$\Omega_i \equiv \frac{\rho_i}{\rho_{cr}} \tag{1.2}$$

From various astronomical measurements, we know that the total density today is almost exactly the critical density [19], so Ω_i can also be interpreted as species i 's fractional contribution to the energy density of the universe.

1.1.1 Galactic Rotation Curves

Much of the earliest, and still most convincing, evidence for the existence of dark matter comes from measurements of galactic rotation curves; that is, the average rotational velocity of stars about the galactic center as a function of their distance from the center. The rotational velocity is generally measured by observing the Doppler shift of the neutral hydrogen 21 cm line. If a galaxy is composed primarily of luminous matter (stars and nebulae) with only a small fraction of its mass not visible to telescopes (i.e., with very weak or no electromagnetic emissions), it is a simple matter to calculate the predicted rotation curve. In particular, for regions far towards the edge of the galaxy, beyond the bulk of the (visible) mass, one expects the rotation curve to go roughly as $v \propto 1/\sqrt{r}$, as the acceleration due

to the gravity of the galaxy goes as $1/r^2$ and the necessary centripetal acceleration goes as v^2/r .

However, as can be seen in Figure 1.1, galactic rotation curves are roughly flat out to radii well past the main distribution of the visible mass [20]. The plots in Figure 1.1 also show the calculated contributions to the rotation curves due to the measured distributions of gas and stellar matter in the galaxies, and they fall well short of the measured rotational velocities. Unless one would propose a modified Newtonian dynamics (MOND) at large scales, there must be an unseen distribution of dark matter to make up the difference, shown by the dash-dot line in Figure 1.1. There is some uncertainty on the dark matter distribution within the bulge and especially within the galactic core, but outside the bulge, galaxies seem to exhibit a universal dark matter halo with a density profile of roughly $\rho \propto 1/r^2$ [20]. Moreover, the total mass of the dark matter must be 10-100 times that of the luminous matter.

1.1.2 The Cosmic Microwave Background

The strongest evidence indicating the existence of non-baryonic dark matter today comes from cosmological measurements, in particular measurements of the anisotropies of the cosmic microwave background (CMB) by the Wilkinson Microwave Anisotropy Probe (WMAP). The cosmic microwave background is an almost perfectly isotropic bath of low-temperature (2.73 K [21]) microwave photons emitted approximately 380,000 years after the big bang [19].

Prior to the time of emission of the CMB, or recombination, the universe was an ionized plasma. In this state, the universe is opaque to photons; i.e., the mean free path for a photon to travel before scattering is short. Therefore, the baryon and photon distributions are tightly coupled, and can be approximated as a single baryon-photon fluid. As the universe continues to expand and cool from the initial conditions of the big bang, the temperature eventually falls below atomic ionization energies (a few eV), at which point the ionized plasma forms into neutral atoms (recombination). Once there is no remaining free charge, the universe becomes transparent, and the photons decouple from the baryons and free-stream from that

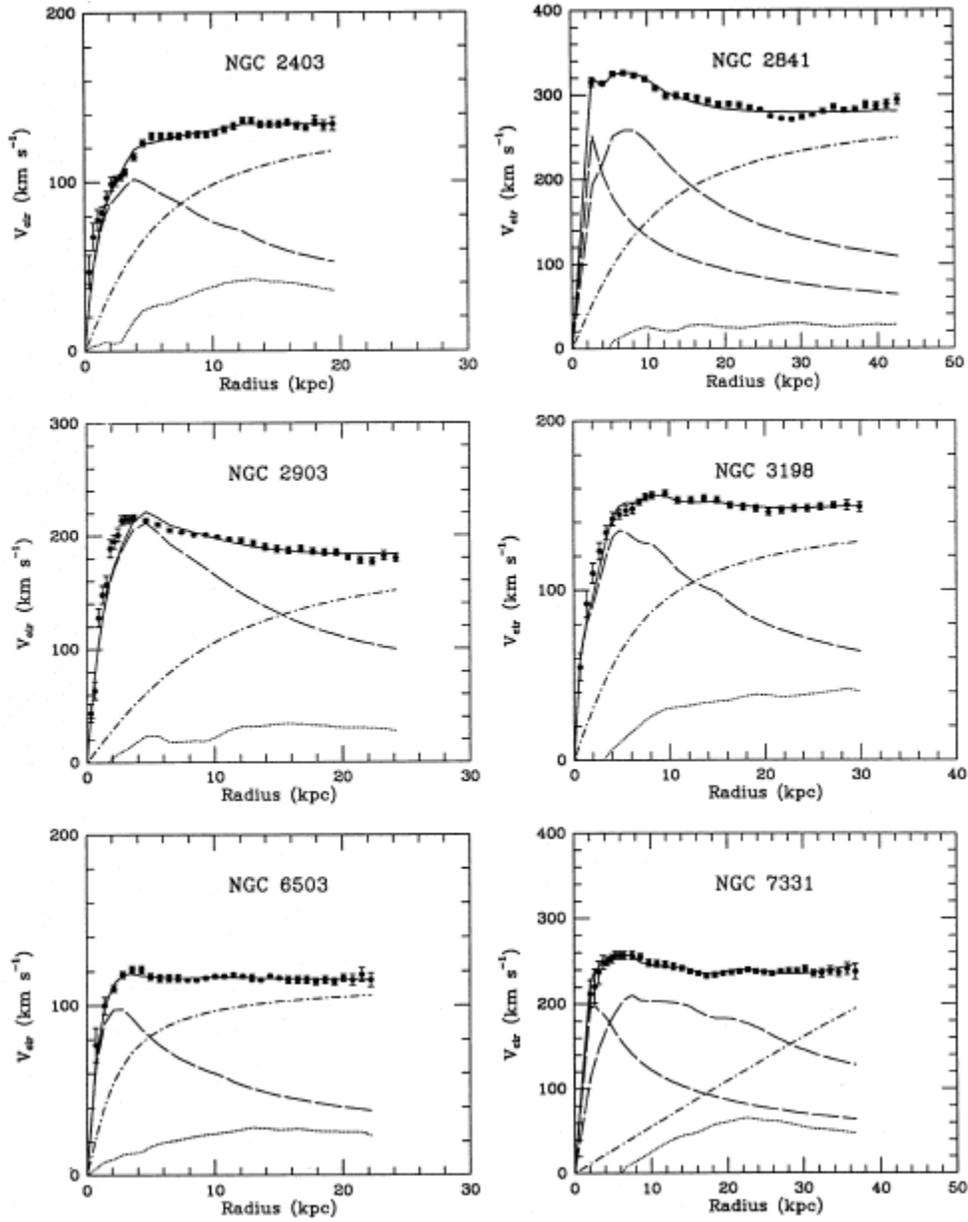


Figure 1.1: Measured rotation curves of several galaxies with three-parameter dark halo fits. The solid line in each figure is the total fit, with the individual contributions from luminous components, gas, and dark matter represented by the dashed, dotted, and dash-dotted lines, respectively [20].

time until we measure them today as the CMB. (The CMB has a temperature of 2.7 K rather than the $\sim 10^4$ K at the time of decoupling because of the continued expansion of space between that time and now.)

Because the CMB photons were tightly coupled to the baryons before decoupling and have been largely unperturbed since then, they provide a “snapshot” of the state of the universe at the time of recombination. How can a single snapshot give us information about the composition and evolution of the universe up to the present day? At the risk of oversimplifying, I like to consider the analogy of the universe as a musical instrument.

If we pluck the string of an instrument, we set up a series of perturbations in the string. For a sufficiently “sharp” pluck, we introduce these perturbations more-or-less uniformly across all frequencies; i.e., the spectrum of these perturbations in frequency space will be approximately flat. Most of those frequencies will quickly decay away, but some will resonate with the instrument and persist for long times: the fundamental frequency, whose wavelength is twice the length of the string, and a harmonic series with wavelengths $\lambda_n = n\pi/L$. The relative amplitudes of each of the harmonics (the instrument’s *tone*), are determined by the physical properties of the instrument. So, by examining only the power spectrum of the string at some point in time after plucking, we can determine the properties of our instrument, and distinguish, say, a violin from a guitar.

In much the same way, the snapshot of the universe provided by the CMB allows us to determine the properties of the universe in the time leading up to recombination. The universe is “plucked” shortly after the big bang during the epoch of inflation, a period in which the size of the universe expanded by a factor of 10^{30} over a period of some $\sim 10^{-35}$ seconds [18]. The evidence for the inflationary epoch comes from the large-scale uniformity of the universe, (better than a part in 10^5), even on scales much larger than the horizon (i.e., separated by a longer distance than light could have traveled since the beginning of the universe). After inflation, the universe is largely uniform, but seeded with perturbations,

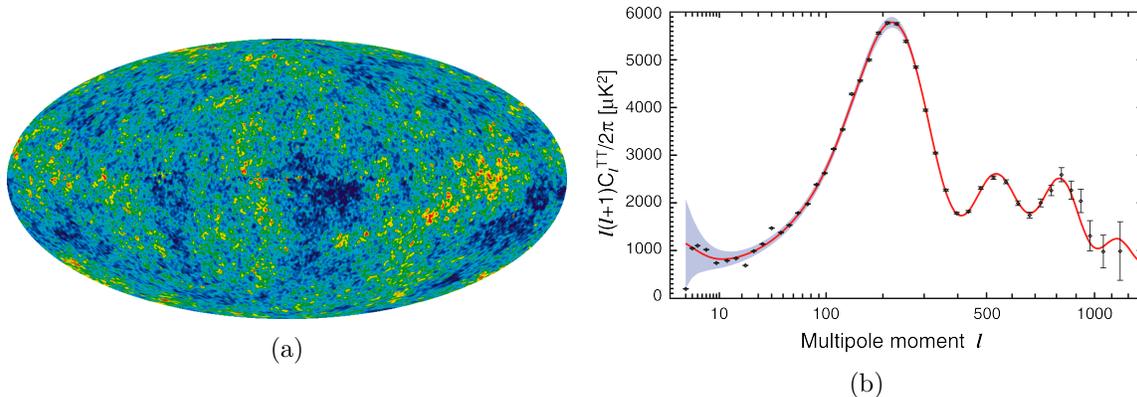


Figure 1.2: WMAP seven-year measurements of the cosmic microwave background (CMB): full sky map (a) and temperature anisotropy spectrum decomposed into spherical harmonics, with the best-fit prediction from Λ CDM cosmology (b) [19, 22].

regions of small over- or under-density, that have roughly uniform amplitudes for all length-scales.

After inflation, the overdense regions create a gravitational potential and begin to grow as they attract more of the surrounding matter. At the same time, the high pressure of the photons causes a net force pushing outward from the overdensities. These opposing forces, tightly coupled in the baryon-photon fluid, result in acoustic oscillations. Dark matter, on the other hand, feels the gravitational attraction, but is coupled to the photon density only indirectly through weak interactions with the baryons.

Figure 1.2 shows the map of the CMB sky measured by WMAP and the power spectrum of temperature anisotropies in spherical harmonics. The first peak in the power spectrum, the fundamental frequency in our instrument analogy, is determined by the sound horizon: the distance a sound wave in the primordial plasma would travel in the time from inflation to recombination. The time at which recombination occurs, and therefore the location of the first peak and the spacing of subsequent peaks, is strongly influenced by the total matter content of the universe. The temperature of matter scales as a^{-3} , while the temperature of a photon gas scales as a^{-4} , where a is the characteristic or “comoving” length scale of the expanding universe [18]. Therefore, as the matter content of the universe is increased,

recombination happens at a later time, and so the sound horizon increases and the peaks in the CMB power spectrum move to smaller angular scales. On the other hand, an increase in the baryon density (and therefore electron density given that the universe is neutral) increases the inertial element of the baryon-photon fluid, lowering the sound speed and sound horizon. However, the change in sound speed is small compared to the change in cooling power with added mass, so we can generally say that more mass leads to larger peak spacing.

Two features in the power spectrum allow us to disentangle the baryon density from the total matter density. As we increase the baryon density, the baryon-photon fluid is more influenced by gravity, attracted toward the initial regions of over-density (which are increased by the infalling matter). The pressure of the photons is then less able to smooth over the gravity wells, and the underdensities from oscillations are damped. Since the anisotropy spectrum in Figure 1.2(b) is a *power* spectrum (i.e., amplitude squared), the even peaks correspond to underdense regions. So, with more baryons, underdensities are smaller, and the ratio of heights between the odd and even peaks gets larger. The second indicator of baryon densities appears at the smallest angular scales. Perturbations at these very small scales are smoothed over by the random diffusion of photons in the plasma. As the baryon density increases, the photon mean free path and diffusion length decrease, so the damping moves to smaller scales [18].

Of course, all of the preceding discussion greatly simplifies an enormously complex analysis. Nevertheless, using similar, but more rigorously defined, arguments, cosmologists have constructed a model of the universe that describes the shape of the CMB anisotropy power spectrum with only 6 free parameters that are determined by fitting to the measured data. The excellent agreement between the model and data, which is shown in Figure 1.2(b), is one of the great scientific successes of this century. From the WMAP data alone, we infer that “normal” baryonic matter makes up only $4.49 \pm 0.28\%$ of the universe. Non-baryonic dark matter makes up $22.2 \pm 2.6\%$, and dark energy, a negative pressure that accelerates the expansion of space at large scales, makes up the majority of the universe at $73.4 \pm 2.9\%$ [19].

The overdensities present at the time of recombination, once there is no longer significant pressure since photons are decoupled, will grow approximately linearly until the present time. This echo of the CMB anisotropies has been measured in the two-point correlation function of galaxy distributions in data from the Sloan Digital Sky Survey (SDSS) [23]. Combining this and other astronomical measurements with the WMAP CMB spectrum allows for a further reduction in the uncertainties of the cosmological parameters. The current best measurements for the baryon, dark matter, and dark energy densities are respectively $\Omega_b = 0.0456 \pm 0.0016$, $\Omega_{CDM} = 0.227 \pm 0.014$, and $\Omega_\Lambda = 0.728^{+0.015}_{-0.016}$, with Ω , the total energy density, equal to $1.0023^{+0.0056}_{-0.0054}$ [19].

For the dark matter density, it is common to use the variable Ω_{CDM} , where CDM indicates “cold,” i.e., non-relativistic, dark matter. Dark matter that was hot, or relativistic, when it decoupled from equilibrium with the primordial plasma (see Section 1.2.1 for an explanation of this process) would have had a high pressure that would dampen overdensities and slow the early growth of large-scale structure in the universe, as well as leaving other tell-tale clues in the CMB. This is an important property, since it largely rules out as a candidate for dark matter any particle that would be relativistic at decoupling.

1.1.3 Big Bang Nucleosynthesis

Measurements of the ratios of light element abundances to hydrogen in the interstellar medium can be used to provide an independent measurement of the total baryon density. This does not measure the dark matter density directly, but does confirm the baryon density estimates from the CMB and reiterate the fact that only a small amount of the total energy density of the universe can be baryonic.

To understand this measurement, we should again consider the universe at an early stage, shortly after the end of inflation. At this point, the universe is still very hot and dense, with all the familiar standard particles, plus dark matter and grand unification scale particles, all constantly being created and annihilated in equilibrium. The universe continues to expand

and cool, and all the heavier particles decouple and quickly decay away. Eventually, when the universe has cooled to a temperature of a few MeV, the universe is composed primarily of stable baryons (protons and neutrons), leptons and anti-leptons undergoing annihilation and pair production in equilibrium, dark matter, and photons (radiation), which make up the bulk of the energy density. At these temperatures, above the nuclear binding energy, virtually no bound atomic nuclei exist, since, as soon as one forms, it is almost immediately broken up by a high-energy photon.

Above a few MeV, neutrons and protons have about the same abundance. As the temperature drops far below the neutron-proton mass splitting (~ 1.3 MeV), the equilibrium neutron to proton ratio is suppressed by the Boltzmann factor $e^{(m_n - m_p)/T}$. If neutrons were to remain in equilibrium, they would eventually all be converted to protons. However, because neutron to proton conversion is a weak process, eventually the process drops out of equilibrium as the universe expands. Neutrons are not stable, so even without multiparticle interactions, they would still eventually all decay if not bound into stable nuclei. Because of the high radiation density, this binding cannot happen to any significant degree until the temperature of the plasma has dropped significantly below the binding energy of the deuteron, 2.22 MeV. The deuterons in turn would all convert into the more tightly bound ${}^4\text{He}$, but this reaction is only possible while the density and temperature are high enough to overcome the Coulomb barrier. ${}^3\text{He}$ and some trace amounts of slightly heavier elements such as Lithium are also formed at this time. Figure 1.3(a) shows how the fractions of the different species evolve as a function of temperature.

Once the temperature falls below around keV, all of these processes have essentially halted; all of the neutrons not bound into nuclei have decayed away, and all the baryons left consist of free protons, ${}^4\text{He}$, and less than 1% of other nuclei. Since we have measurements for all of the relevant cross sections for these processes, and we know how the overall density and temperature evolve with time from general relativity and the Friedmann equations [25, 26], the relative amounts of each of the elements remaining after this Big Bang nucleosynthesis

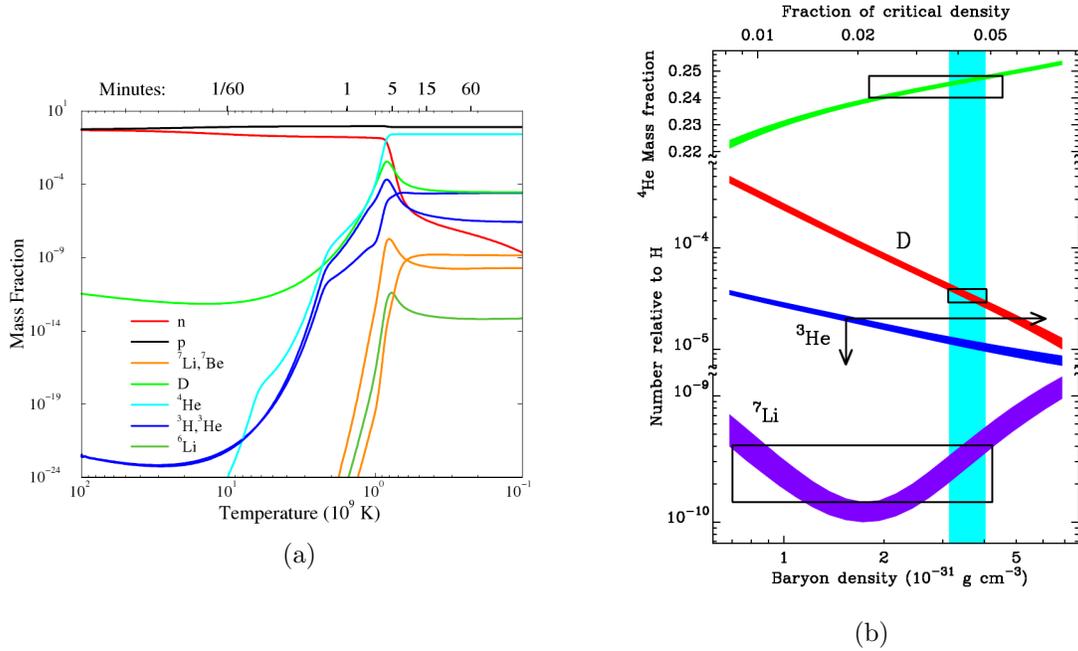


Figure 1.3: (a): Evolution of the relative amounts of protons, neutrons, and light elements during Big Bang nucleosynthesis. (b): Light element abundances from Big Bang nucleosynthesis as a function of total baryon density, with current limits. [24]

(BBN) are uniquely determined by the total baryon density. Therefore, a measurement of the relative primordial densities of elements (before the onset of stellar nucleosynthesis) can determine the total average baryon density of the universe.

The primordial element fractions can be measured by observing the relative amounts of absorption as light emitted from distant galaxies passes through interstellar dust. Figure 1.3(b) shows how the element abundances depend on the baryon density, and shows the current (as of 2000) limits on these values. The hydrogen/deuterium ratio is currently the best measured, and puts the most stringent BBN constraints on the total baryon density, at $\Omega_b = 0.040 \pm 0.004$ (95% C.L.) [27, 28]. This number is in very good agreement with the results derived from CMB measurements, and serves as further, independent evidence that baryons in truth account for only a small fraction of the makeup of the universe.

1.2 Expected properties of dark matter

1.2.1 Dark Matter Candidates

From the scale of anisotropies in the cosmic microwave background and large-scale structures in the visible universe today, we know that dark matter makes up some 23% of the present day energy density and 83% of the total mass. Here, by “dark matter” we refer specifically to non-baryonic dark matter, as opposed to non-luminous baryonic matter. We know that a sizable fraction of baryonic matter is not in luminous bodies like stars, but in harder to detect distributions such as interstellar gas and dust, as well as massive compact halo objects (MACHOs) such as neutron and white dwarf stars. These “dark” mass distributions might be tempting suspects for dark matter, but all the cosmological evidence indicates otherwise. In addition to the CMB anisotropy spectrum, the baryon density is independently measured by the ratios of light elements created during big bang nucleosynthesis, and indicate a baryonic matter density $\Omega_b \approx 0.04$.

In the standard model of particle physics, baryons and electrons are the only stable particles with mass. Since the universe is neutral, the baryon and electron number densities must be equal, so electrons cannot account for the dark matter either. Since we have ruled out the familiar standard model, we must look beyond the standard model for dark matter candidates.

Neutrinos

Although neutrinos in the original standard model are massless, the neutrinos have been observed to oscillate between the different neutrino flavors by multiple experiments, a phenomenon only possible if neutrinos do in fact have mass. Since they do not couple directly to photons but only weakly to leptons and quarks, massive neutrinos would seem at first glance to be a promising candidate for dark matter, and not require proposing a new particle.

In the early universe, neutrinos are held in equilibrium with the universe via annihilation to and from other fermions. Since they are highly relativistic, the number (and energy) density for each species of neutrinos in equilibrium will be roughly the same as for photons. The cross section for the annihilation reaction, once the temperature is well below the mass of the W and Z, scales roughly as $G_F^2 s$, where G_F is the Fermi constant ($\sim 1.2 \times 10^{-5} \text{ GeV}^{-2}$) and s is the square of the total energy in the center of momentum system (CMS). Knowing this rate, it can be shown that neutrinos (at least the three almost massless generations of the standard model) decouple from the primordial plasma at temperatures of a few MeV [29]. In this case, they are still highly relativistic, which contradicts evidence from structure formation indicating that dark matter must have been cold at the time of decoupling. If most of the dark matter was “hot,” or relativistic, during early structure formation, the large pressure would cause them to stream away from and dampen any primordial overdensities, preventing the large-scale structures observable today.

Moreover, knowing roughly the number density of neutrinos at the time of decoupling, the sum of the masses of the three neutrino families would have to be ~ 50 meV [29], but the experimental limits for the neutrino mass are less than a few eV, and likely much smaller based on the mass splittings [30].

An alternative would be a fourth generation of neutrino, which could be massive enough to be “cold” at the time of decoupling. We know from the width of the Z resonance measured in collider experiments that any new neutrino must have a mass greater than $M_Z/2 = 45$ GeV. For neutrino masses between ~ 45 GeV and ~ 1 TeV, the annihilation cross section ($\nu\bar{\nu} \rightarrow e^+e^-$) is large enough that the surviving relic density must be small ($\Omega < 0.01$), so neutrinos with standard model Weak couplings would have to have masses above the TeV scale to account for a significant portion of the observed dark matter relic density. Dirac neutrinos of these masses with standard model couplings have been ruled out experimentally, but sterile Majorana neutrinos or Dirac neutrinos with weaker couplings to the standard model Z boson than the 3 light species remain viable candidates [31]. With these masses and couplings, we

can lump these hypothetical neutrinos into a broader category of weakly interacting massive particles (WIMPs), which will be described later in this section.

Axions

The axion is a very light pseudo-scalar boson first proposed by Peccei and Quinn to account for the smallness of strong CP violations [32]. CP-violating complex phases occur naturally in the calculation of quantum chromodynamics (QCD) interactions. This CP violation is not *a priori* forbidden, but have not been observed experimentally. In particular, the magnetic moment of the neutron is at least nine orders of magnitude smaller than that predicted by strong CP violation [29]. Peccei and Quinn proposed that a spontaneously broken U(1) symmetry, resulting in a Goldstone boson, the axion, could account for this discrepancy.

Although not technically stable (the axion can decay to two photons), the decay rate scales as $\sim m_a^5$, so the lifetime would be longer than the age of the universe for masses below ~ 10 eV [29]. Such an axion would survive as a relic from early expansion and could be a viable dark matter candidate. Axions can be produced by photon interaction with atomic nuclei in stellar cores, but, being ultra-weakly interacting, would then escape the star, accelerating the cooling of the star. The observed evolution of red giant stars and Supernova 1987a then place further limits on the axion mass of $m_a < \sim 0.01$ eV [33].

Although the very light axion mass would seem to disqualify it as a source of cold dark matter by the same arguments about light neutrinos, the very weak coupling to ordinary matter would prevent axions from ever obtaining thermal equilibrium with the early universe. Therefore, they would be produced during inflation as a cold boson condensate. To obtain the proper relic density, axions would require masses on the μeV to meV range. Axions of this mass range are considered to be viable dark matter candidates, and are the focus for several experimental searches [34, 35].

WIMPs

Perhaps the most interesting cold dark matter candidate is a general class of weakly interacting massive particle or WIMP. The majority of direct dark matter detection searches have focused on WIMPs, and so, too, will the remainder of this work. Most searches focus on WIMPs with masses roughly in the GeV–TeV range and couplings to ordinary matter at the weak scale or below. Particles that could be considered WIMPs appear naturally in many theories that attempt to address physics beyond the standard model, which is why the WIMP is such a favored cold dark matter candidate. One class of beyond-standard-model theories, supersymmetry (SUSY), is widely favored as a likely successor to the standard model. Supersymmetry predicts that, for every particle of the standard model, there exists a supersymmetric partner, or sparticle, with TeV-scale mass and opposite spin statistics (i.e., a boson’s superpartner is a fermion, and vice versa). In the minimum supersymmetric model (MSSM) and most extensions, spontaneous symmetry breaking causes the W, Z, and Higgs superpartners to mix, forming a series of neutral particles (neutralinos), the lightest of which is stable, and is the prototypical WIMP candidate [36].

Neutralinos in particular, and WIMPs in general, are favored dark matter candidates because of the so-called “Weak miracle:” if one assumes approximately weak-scale cross sections for the cold dark matter annihilation process, one obtains a dark matter relic density (i.e., remaining density after decoupling from the plasma of the early universe) consistent with cosmological measurements. Moreover, the resulting relic density varies only logarithmically with the mass of the WIMP, allowing a wide range of possible candidates. The basic outline of the calculation is as follows:

The annihilation reaction rate for a dark matter particle X in the plasma of the early universe is given by

$$\Gamma_- = \int \sigma \frac{d\Phi}{dv} dv = \langle \sigma \Phi \rangle = n \langle \sigma v \rangle \quad (1.3)$$

where σ is the total annihilation cross section, Φ is the flux of antiparticles, n is the antiparticle number density, v is the average speed of the antiparticle flux, and $\langle \rangle$ denotes the thermal average. Equation (1.3) is the average annihilation rate *per particle*, so the total change in particle number density due to annihilations (assuming that the particle and antiparticle densities are equal) is

$$\frac{dn}{dt} = -\langle \sigma v \rangle n^2 \quad (1.4)$$

In many theories, dark matter is its own antiparticle, so the assumption of equal densities is automatically enforced.

Similarly, the rate of X -creating reactions depends on the flux of particles that can annihilate into X particles:

$$\Gamma_+ = \int \sigma \frac{d\Phi_1}{dv} dv = n_1 \langle \sigma v \rangle e^{-m_X/T} \quad (1.5)$$

where m_X is the mass of a single dark matter particle. The exponential term arises because heavy dark matter particles can only be created in collision where the total energy is greater than twice the X mass. Since the particles are in thermal equilibrium, their energies will follow a Boltzmann distribution, with an exponential tail at higher energies. This expression, once we substitute for the density n_1 of X -creating particles, becomes the number density of X particles in thermal equilibrium (assuming the Maxwell-Boltzmann approximation) [18]:

$$n_{\text{eq}} = g \int \frac{d^3p}{(2\pi)^3} e^{-E/T} = \begin{cases} g \left(\frac{mT}{2\pi}\right)^{3/2} e^{-m/T} & m \gg T \\ g \frac{T^3}{\pi^2} & m \ll T \end{cases} \quad (1.6)$$

where g is the degrees of freedom of the particle (2 for a spin-1/2 particle, etc.).

Finally, the number density is reduced by the expansion of space, expressed by the Hubble parameter H . Thus, the total rate of change in X particle density is

$$\frac{dn}{dt} = -3Hn - \langle\sigma v\rangle(n^2 - n_{\text{eq}}^2) \quad (1.7)$$

We are most interested in the evolution of the ‘‘co-moving’’ number density, that is, the change in dark matter particle density ignoring the expansion of space. A convenient way to express this is to normalize to the entropy density s , which, due to entropy conservation, changes only with the expansion of space:

$$\frac{ds}{dt} = -3Hs \quad (1.8)$$

where s is given by

$$s = 2\pi^2 h_{\text{eff}} T^3 / 45 \quad (1.9)$$

and h_{eff} counts the effective relativistic degrees of freedom and is a function of temperature [16]. If we introduce the variable $Y \equiv n/s$, we can combine equations (1.7) and (1.8) to obtain

$$\frac{dY}{dt} = -s\langle\sigma v\rangle(Y^2 - Y_{\text{eq}}^2). \quad (1.10)$$

To further simplify, we can track the density not as a function of absolute time but rather temperature in units of the dark matter particle mass $x \equiv m_X/T$. Combining this with Equation (1.9), Equation (1.10) becomes

$$\frac{dY}{dx} = -\frac{2\pi^2 h_{\text{eff}}}{45H} \frac{mT^2}{x^2} \langle\sigma v\rangle (Y^2 - Y_{\text{eq}}^2) \quad (1.11)$$

The Friedmann equation relates the rate of expansion of space (i.e., the Hubble parameter) to the energy density:

$$H^2 = \frac{8\pi}{3M_P^2} \rho, \quad (1.12)$$

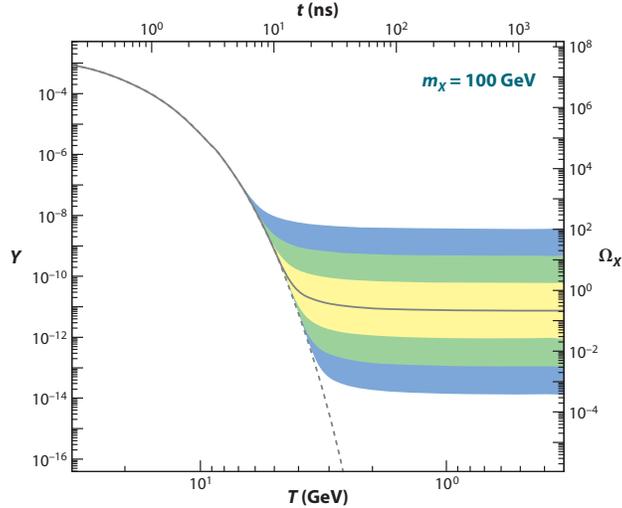


Figure 1.4: Numerical solution to Equation (1.14) for a 100 GeV particle [38]. The central gray solid line is for an annihilation cross section that yields the correct relic density; the shaded regions are for cross sections that differ from the central value by 10, 10^2 , and 10^3 from this value. Larger annihilation cross sections result in lower surviving relic density.

where $M_P = 1,22 \times 10^{19}$ GeV is the Planck mass. In the early universe where the energy density is dominated by photons [37],

$$\rho = \frac{\pi^2}{g_{eff}} T^4 \quad (1.13)$$

where g_{eff} is the effective degrees of freedom for the energy of the photon plasma. Combining Equations (1.11), (1.12), and (1.13), we obtain

$$\frac{dY}{dx} = -\sqrt{\frac{\pi g_* M_P^2}{45}} \frac{m_X}{x^2} \langle \sigma v \rangle (Y^2 - Y_{eq}^2) \quad (1.14)$$

where $g_* \equiv h_{eff}^2/g_{eff}$ parameterizes the degrees of freedom.

The freeze-out, or decoupling, process is evident on examination of Equation (1.14), for which Figure 1.4 shows a numerical solution for a particle with mass $m_X = 100$ GeV. At early times, the temperature is large, so $x \ll 1$, making dY/dx large compared to the $(Y^2 - Y_{eq}^2)$ term, and so the abundance closely tracks the equilibrium value. As the temperature decreases to the X mass and below, the equilibrium abundance is suppressed

exponentially by the Boltzmann factor e^{-x} . At very late times, $T \ll m$ and so $x \gg 1$, and $dY/dt \rightarrow 0$, and the comoving X density becomes constant. The final dark matter relic abundance depends on the freeze-out temperature, at which the annihilation rate becomes less than the expansion rate and the abundance no longer closely tracks the equilibrium value.

Equation 1.14 can only be solved numerically, but Dodelson [18] provides a convenient method to approximate the final abundance analytically. First, for convenience we define the dimensionless constant

$$\lambda \equiv \sqrt{\frac{\pi g_*}{45}} M_P m_x \langle \sigma v \rangle. \quad (1.15)$$

We've already said that, at large temperatures, Y_{eq} drops off exponentially, so we can approximate Equation (1.14) by

$$\frac{dY}{dx} \simeq -\frac{\lambda Y^2}{x^2} \quad (1.16)$$

This can be integrated analytically from the epoch of freeze-out to the present time where $x \rightarrow \infty$ to yield

$$\frac{1}{Y_0} - \frac{1}{Y_f} = \frac{\lambda}{x_f} \quad (1.17)$$

where the subscript f denotes the value at freeze-out. Numerical solutions show that the freeze-out temperature, x_f , varies only slightly with particle mass; generally we can assume $x_f \sim 20$ [37]. At this temperature, annihilation reactions are still taking place at a high rate, although not fast enough to track the equilibrium abundance. Therefore the abundance at freeze-out Y_f is generally much larger than the present day abundance Y_0 , so we can approximate Equation (1.17) as

$$Y_0 \simeq \frac{x_f}{\lambda}. \quad (1.18)$$

The fraction of the critical density contributed by dark matter is then

$$\Omega_X = \frac{m_x s_0 Y_0}{\rho_c} = \frac{m_x s_0 x_f}{\rho_c \lambda} = \sqrt{\frac{45}{\pi g_*}} \frac{s_0}{M_P} \frac{1}{\langle \sigma v \rangle}. \quad (1.19)$$

Plugging in numerical values for all of the terms in Equation (1.19) yields [18]

$$\Omega_X \simeq \frac{10^{-36} \text{cm}^2}{\langle \sigma v \rangle}. \quad (1.20)$$

Equation (1.20) tells us that the dark matter relic abundance to first order is insensitive to the mass of the dark matter particle and inversely proportional to the annihilation cross section. The so-called weak miracle is that, without any *a priori* assumptions about the annihilation process, we find that cross-sections on the order of 10^{-36} cm^2 , consistent with approximately weak-scale interactions, yield a dark matter relic abundance consistent with cosmological measurements. This weak-scale result is consistent with similar particles proposed in many beyond-standard-model theories such as supersymmetry and its extensions, and is one of the main reasons leading to the favored status of the WIMP hypothesis.

For the remainder of this work, we will concern ourselves only with WIMP-type dark matter.

1.2.2 Expected WIMP Direct Detection Signals

As we have seen in the previous sections, we can loosely define a class of Weakly Interacting Massive Particles, or WIMPs, that are well-motivated both theoretically and by astronomical observations. In order to fit the evidence so far accumulated, in particular to obtain the correct relic abundance known by cosmology, WIMPs must have an annihilation cross section at approximately the weak scale, in the neighborhood of 10^{-36} cm^2 , mass approximately in the GeV-TeV range, and be “cold”, i.e., non-relativistic. How can we search for evidence of these WIMPs through something other than their gravitational interactions? One obvious way is to look for evidence of their continuing annihilation; and the status of some experiments attempting to do just that will be discussed in Section 1.3.1. Or we can invert the process and attempt to create WIMPs in high energy accelerators. Finally, we can turn the annihilation diagram on its side and look for WIMPs to scatter off of some target, which will

be the subject the bulk of this work. In this section, we will determine the expected rate and spectrum of signals due to WIMPs scattering in targets of various materials, making as few assumptions as possible about the actual interaction.

By crossing symmetry, we know that all of these processes should have approximately the same cross section. The exact WIMP mass is usually left as a somewhat free parameter in calculations of this sort, but one important statement can be made for any WIMP mass in the GeV or above scale: WIMPs can only transfer a negligible amount of energy to an electron due to the simple kinematics of large difference in masses. Therefore, if it is possible to observe a WIMP directly interacting with some target material, the signature will be a recoiling atomic nucleus. This signal would be extremely similar to that of elastic neutron scattering, a fact which has both good and bad consequences, as we will see in later chapters. Here, with a few well founded assumptions about the local distribution of WIMP velocities, we will now calculate the spectrum of these elastic nuclear recoils.

Basic Rates

We can derive the expected rate of nuclear recoils from WIMP interactions as follows. In very basic terms, the differential rate of recoils at a given energy E_R per kilogram of target mass is given by

$$\frac{dR}{dE_R} = \frac{N_0}{A} \int_0^\infty \frac{d\sigma(E_R, v_D)}{dE_R} \frac{d\Phi_D(v_D)}{dv_D} dv_D \quad (1.21)$$

where N_0 is Avogadro's number, $6.02 \times 10^{26} \text{ kg}^{-1}$, A is the atomic mass of the target in AMU, σ is the per-nucleus scattering cross section, v_D is the velocity of the incident dark matter particle in the target's frame, and $\Phi(v_D)$ is the flux of dark matter particles with velocity v_D through the target.

We can expand $\frac{d\sigma}{dE_R} = \frac{d\sigma}{d\cos\theta} \frac{d\cos\theta}{dE_R}$, where θ is the angle of the recoiling nucleus in the center of momentum (CM) frame with respect to the initial vector of the DM particle. If the scattering is spin-independent and with small momentum transfer, then it should be isotropic in the CM frame, with the only angular dependence being the solid angle factor.

Then $\frac{d\sigma}{d\cos\theta} = \sigma_0$, roughly constant for a given nucleus, and we can derive from the kinematics that

$$\frac{d\cos\theta}{dE_R} = \frac{M_T}{2v_D^2\mu^2}; \quad v_D \geq v_{min} \quad (1.22)$$

where $M_T = 0.932A$ GeV is the mass of the target nucleus, μ is the reduced mass of the WIMP-nucleus system, and v_{min} is the minimum velocity of the incident WIMP necessary to produce a recoil of energy E_R , given by the case of forward scattering of the nucleus ($\cos\theta = 1$). For elastic scattering, v_{min} is given by

$$v_{min} = \sqrt{\frac{M_T E_R}{2\mu^2}} \quad (1.23)$$

If there is a constant number density of dark matter particles $n = \rho_D/M_D$, where ρ_D is the mass density of dark matter per unit volume and M_D is the mass of a single dark matter particle, then the differential flux of dark matter with a given velocity v_D in the Earth's rest frame is given by

$$\frac{d\Phi_D}{dv_D} = \frac{\rho_D}{M_D} v_D f(\mathbf{v}_D, \mathbf{v}_E, v_{esc}) \quad (1.24)$$

where f is the velocity distribution function, \mathbf{v}_E is the Earth's velocity relative to the halo rest frame, and v_{esc} is the galactic escape velocity. We will derive a particular form for f in section 1.2.2. Combining equations (1.21), (1.22), and (1.24), we get

$$\frac{dR}{dE_R} = \frac{N_0 \sigma_0 \rho_D}{A} \frac{M_T}{2\mu^2 M_D} \int_{v_{min}}^{\infty} \frac{f(\mathbf{v}_D, \mathbf{v}_E, v_{esc})}{v_D} dv_D \quad (1.25)$$

To compare rates and spectra across different experiments using different targets, it is necessary to express the rate in terms of the WIMP-*nucleon* cross section, rather than the WIMP-*nucleus* cross section as used up to this point. To see how these two cross sections relate, we need to take a step back and look at the amplitude of the scattering WIMP wavefunction. If the total amplitude (i.e., matrix element) for a WIMP to scatter off a given nucleon i located at position \mathbf{x}_i from momentum-state \mathbf{p} into state \mathbf{p}' (by transferring

momentum $\mathbf{q} = \mathbf{p}' - \mathbf{p}$ to the nucleus) is $f_i(\mathbf{p}, \mathbf{p}')$, then the total amplitude for scattering from all nucleons adds together with associated phase factors:

$$F(\mathbf{p}, \mathbf{p}') = \sum_{i=1}^A f_i(\mathbf{p}, \mathbf{p}') e^{i\mathbf{q} \cdot \mathbf{r}_i} \quad (1.26)$$

The cross section depends on the square of the total scattering amplitude,

$$|F(\mathbf{p}, \mathbf{p}')|^2 = \sum_{i=1}^A |f_i(\mathbf{p}, \mathbf{p}')|^2 + \sum_{i,j \neq i}^A f_i(\mathbf{p}, \mathbf{p}') f_j^*(\mathbf{p}, \mathbf{p}') e^{i\mathbf{q} \cdot (\mathbf{r}_i - \mathbf{r}_j)} \quad (1.27)$$

Again, here we are assuming a spin-independent interaction.

If the momentum transfer is small, i.e. the quantity

$$qr_n = |\mathbf{q} \cdot \max(\mathbf{r}_i - \mathbf{r}_j)| \ll 1$$

(here r_n is approximately the size of the nucleus), then the phase factors in the second term in equation (1.27) reduce to unity, and the scattering amplitudes from each individual nucleon add coherently. In other terms if the momentum transfer is small, the de Broglie wavelength of the momentum transfer can be large compared to the size of the nucleus, and all the nucleons are probed at once. Since q , the total momentum transfer, is given by $\sqrt{2M_T E_R}$ and a typical size for a nucleus is a few fm, the coherence condition holds well for recoil energies up to ~ 50 keV. As the momentum transfer increases, the difference in the phase factors becomes more significant, and the coherence effect becomes less pronounced. We can model the gradual effect of loss of coherence by the use of a nuclear form factor, $F(q)$. A convenient parameterization is given by [39]

$$F(q) = 3 \frac{j_1(qr_n)}{qr_n} e^{-(qs)^2/2} \quad (1.28)$$

where j_1 is a spherical Bessel function of the first kind; r_n is the nuclear radius and is approximately $1.14A^{1/3}$ fm; and s is the nuclear skin thickness, typically ~ 0.9 fm [40].

Finally, we can further simplify equation (1.27) by assuming that the scattering amplitude for any given nucleon depends only on whether it is a proton or a neutron. Then, if we define $\sigma_n \propto |f_n(\mathbf{p}, \mathbf{p}')|^2$ as the cross section for WIMP scattering off of a neutron, and f_p/f_n as the ratio of the scattering amplitudes between neutrons and protons, then we obtain for the WIMP-nucleus cross section

$$\sigma_0 = \sigma_n \left(Z \frac{f_p}{f_n} + (A - Z) \right)^2 F^2(q) \frac{\mu^2}{\mu_n^2} \quad (1.29)$$

where μ_n is the reduced mass of a WIMP-nucleon system (treating the masses of the proton and neutron as equal). The factor μ^2/μ_n^2 comes from the different sum over possible final states based on the target mass. We could of course define σ_0 similarly in terms of the proton-WIMP cross section.

If, as is often assumed, $f_p/f_n = 1$, then the term in parentheses in equation (1.29) reduces to A^2 , which is the usually-cited scaling factor. In this case, or indeed so long as neither of the scattering amplitudes is negative, heavier elements can expect to see a higher rate of WIMP-induced nuclear recoils for a given σ_n and target mass, at least in the low recoil-energy range where the scattering is coherent.

Substituting equation (1.29) into equation (1.25), the differential rate of recoils as a function of the WIMP-neutron cross section is given by

$$\frac{dR}{dE_R} = \frac{N_0 \sigma_n \rho_D M_T}{2A \mu_n^2 M_D} \left(Z \frac{f_p}{f_n} + (A - Z) \right)^2 F^2(q) \int_{v_{min}}^{\infty} \frac{f(\mathbf{v}_D, \mathbf{v}_E, v_{esc})}{v_D} dv_D \quad (1.30)$$

In order to eliminate a free parameter for further discussion, from this point forward unless specified we will take the commonly held assumption that $f_p = f_n$, that is, WIMPs interact identically with protons and neutrons, and the cross section depends only on A^2 .

Velocity Distribution Function

We can derive a specific form for the velocity distribution function of dark matter particles in the Earth's rest frame via the following. At the time of decoupling, the WIMP gas was in thermal equilibrium with the surrounding universe, and then cooled adiabatically with the expansion of the universe. As local anisotropies gave rise to the current galaxy, WIMPs with velocities higher than the escape velocity would quickly leave the system, but, because of the low interaction rate, no large-scale evaporative cooling could take place during this process. So, in the rest frame of the galactic halo, we assume a Maxwellian velocity distribution, truncated above the escape velocity:

$$f_H(\mathbf{v}) = k \exp\left(-\frac{v^2}{v_0^2}\right), \quad v \leq v_{esc}; \quad 0 \text{ otherwise} \quad (1.31)$$

where v_0 is given by the local rotational velocity, approximately 230 km/s [40]. k is a normalization constant:

$$1/k = \int_0^{v_{esc}} \exp\left(-\frac{v^2}{v_0^2}\right) dv = \pi^{3/2} v_0^3 \left[\operatorname{erf}(v_{esc}/v_0) - \frac{2v_{esc}}{\pi^{1/2}v_0} \exp\left(-\frac{v_{esc}^2}{v_0^2}\right) \right] \quad (1.32)$$

To find the velocity distribution in the Earth's frame, we boost this distribution with the Earth's velocity \mathbf{v}_E . Because the largest velocity involved, v_{esc} , is much less than the speed of light (estimates of the upper limit for v_{esc} are typically on the order of 600 km/s [40], $\sim 10^{-3}c$), we can use Galilean relativity with negligible loss of accuracy. For velocities $v_D \leq v_{esc} - v_E$ (where again, v_D is the dark matter velocity in the earth frame), the cumulative distribution function, $P(V \leq v_D)$ is equal to the integral of the halo distribution f_H over a sphere of radius v_D , centered at \mathbf{v}_E (Figure 1.5, left). Choosing our reference frame such that \mathbf{v}_E lies along the $+\hat{z}$ direction, and using cylindrical coordinates $z = v_{H,z}$ and

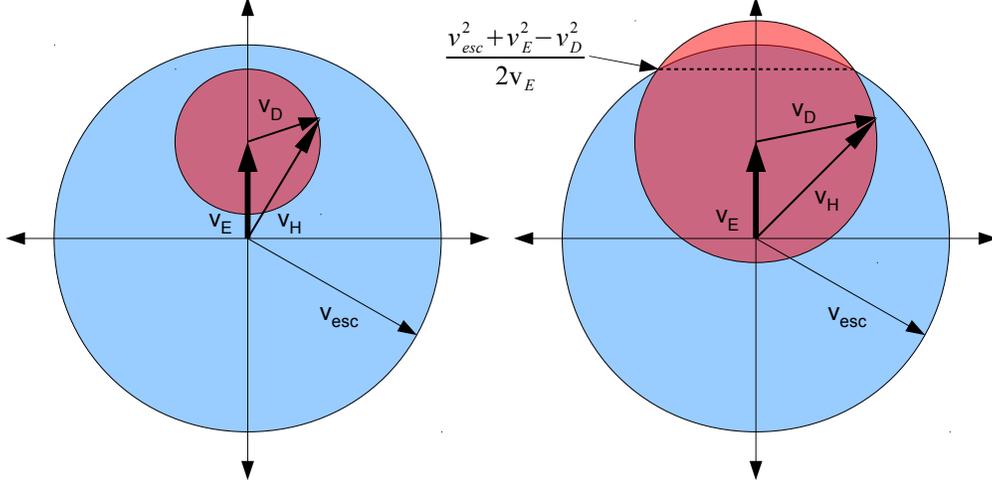


Figure 1.5: Integration volume defining $f(v_D)$ in the Earth's frame. The light blue circle (which should properly be a sphere) represents the cutoff at v_{esc} of the halo velocity distribution. The smaller circle has a radius v_D . The integration region defining $P(V \leq v_D)$ is the intersection of the two areas. Left: $v_D \leq v_{esc} - v_E$, right, $v_{esc} - v_E < v_D \leq v_{esc} + v_E$.

$r = \sqrt{v_{H,x}^2 + v_{H,y}^2}$, this function is

$$P(V \leq v_D) = 2\pi k \int_{v_E - v_D}^{v_E + v_D} \int_0^{\sqrt{v_D^2 - (z - v_E)^2}} r \exp\left(-\frac{r^2 + z^2}{v_0^2}\right) dr dz \quad (1.33)$$

The distribution function is then given by

$$\begin{aligned} f(v_D) &= \frac{\partial}{\partial v_D} P(V \leq v_D) \\ &= 2\pi k \frac{\partial}{\partial v_D} \int_{v_E - v_D}^{v_E + v_D} \int_0^{\sqrt{v_D^2 - (z - v_E)^2}} r \exp\left(-\frac{r^2 + z^2}{v_0^2}\right) dr dz \\ &= 2\pi k \frac{v_0^2}{2} \frac{\partial}{\partial v_D} \int_{v_E - v_D}^{v_E + v_D} \exp\left(-\frac{z^2}{v_0^2}\right) - \exp\left(-\frac{v_D^2 - v_E^2}{v_0^2}\right) \exp\left(-\frac{2zv_E}{v_0^2}\right) dz \\ &= 2\pi k \frac{v_0^2}{2} \frac{\partial}{\partial v_D} \left[\frac{\sqrt{\pi} v_0}{2} \left(\operatorname{erf}\left(\frac{v_D - v_E}{v_0}\right) + \operatorname{erf}\left(\frac{v_D + v_E}{v_0}\right) \right) - \right. \\ &\quad \left. \frac{v_0^2}{2v_E} \left(\exp\left(-\frac{(v_D - v_E)^2}{v_0^2}\right) - \exp\left(-\frac{(v_D + v_E)^2}{v_0^2}\right) \right) \right] \\ f(v_D) &= \frac{\pi k v_0^2 v_D}{v_E} \left[\exp\left(-\frac{(v_D - v_E)^2}{v_0^2}\right) - \exp\left(-\frac{(v_D + v_E)^2}{v_0^2}\right) \right]; \quad v_D \leq v_{esc} - v_E \quad (1.34) \end{aligned}$$

If $v_D > v_{esc} - v_E$, then we have the case shown in Figure 1.5, right, where the integral is to be taken over the volume defined by the intersection of one sphere of radius v_{esc} centered at the origin and one sphere of radius v_D centered at \mathbf{v}_E . In this case, the cumulative distribution function is given by

$$\begin{aligned}
P(V \leq v_D) = & 2\pi k \int_{v_E - v_D}^{\frac{v_{esc}^2 + v_E^2 - v_D^2}{2v_E}} \int_0^{\sqrt{v_D^2 - (z - v_E)^2}} r \exp\left(-\frac{r^2 + z^2}{v_0^2}\right) dr dz + \\
& 2\pi k \int_{\frac{v_{esc}^2 + v_E^2 - v_D^2}{2v_E}}^{v_{esc}} \int_0^{\sqrt{v_{esc}^2 - z^2}} r \exp\left(-\frac{r^2 + z^2}{v_0^2}\right) dr dz
\end{aligned} \tag{1.35}$$

And the distribution function is then given by

$$f(v_D) = \frac{\pi k v_0^2 v_D}{v_E} \left[\exp\left(-\frac{(v_D - v_E)^2}{v_0^2}\right) - \exp\left(-\frac{v_{esc}^2}{v_0^2}\right) \right]; \quad v_{esc} - v_E < v_D \leq v_{esc} + v_E \tag{1.36}$$

Finally, combining equations (1.34) and (1.36), we can write

$$f(v_D) = \frac{\pi k v_0^2 v_D}{v_E} \begin{cases} \exp\left(-\frac{(v - v_E)^2}{v_0^2}\right) - \exp\left(-\frac{(v + v_E)^2}{v_0^2}\right) & \text{if } v_D \leq v_{esc} - v_E \\ \exp\left(-\frac{(v - v_E)^2}{v_0^2}\right) - \exp\left(-\frac{v_{esc}^2}{v_0^2}\right) & \text{if } v_{esc} - v_E < v_D \leq v_{esc} + v_E \\ 0 & \text{if } v_D > v_{esc} + v_E \end{cases} \tag{1.37}$$

Figure 1.6 is a plot of $f(v_D)$ from Equation (1.37) for the parameters as given in table 1.1.

Full Rate Expression

Before we finish producing the final rate expression, we can derive some basic order-of-magnitude expressions. First, the total flux of WIMPs at earth is given by the integral of Equation (1.24):

$$\Phi_D = \rho_D / M_D \bar{v}_D \tag{1.38}$$

From Figure 1.6, $\bar{v}_D \approx 360$ km/s. Then, for $\rho_D = 0.3$ GeV/cm³ [40] and a 100 GeV WIMP, the total WIMP flux is on the order of 10^5 cm⁻² s⁻¹. The total interaction rate per atom of

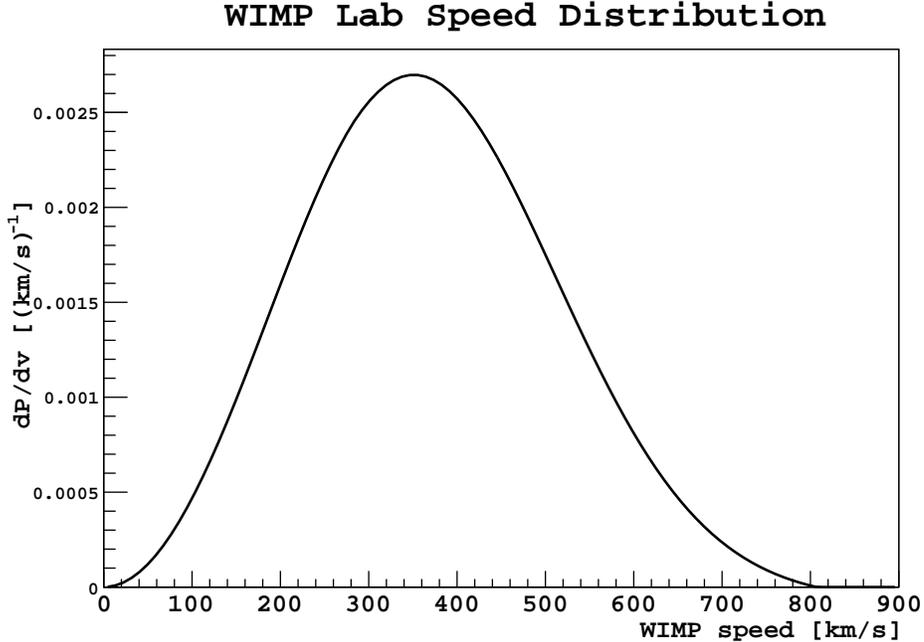


Figure 1.6: Distribution function of dark matter velocities in the Earth frame, for the parameters given in table 1.1.

target is $\sigma_0 \Phi_D$, with σ_0 given by equation (1.29). The term $(\mu/\mu_n)^2$ goes roughly as αA^2 , where α varies slowly with A and is of order 0.1-1. Since the form factor F^2 is also of order 1 for elements with moderate mass, if $f_p = f_n$, then the cross section per atomic nucleus goes roughly as

$$\sigma_0 \approx \sigma_n A^4. \quad (1.39)$$

For argon ($A = 40$) and a WIMP-nucleon cross section $\sigma_n = 10^{-45} \text{ cm}^2$, the total nuclear cross section is then $\sim 10^{-39} \text{ cm}^2$. The expected rate per argon atom should be on the order of 10^{-34} s^{-1} , and the rate per kilogram on the order of $10^{-4} \text{ kg}^{-1} \text{ day}^{-1}$.

Now we will compare these “back of the envelope” calculations to those obtained by evaluating Equation (1.30) using the velocity distribution function in Equation (1.37). First,

we need to know the quantity

$$\int_{v_{min}}^{\infty} \frac{f(v_D)}{v_D} dv_D = \frac{\pi^{3/2} v_0^3 k}{2v_E} \times \begin{cases} \operatorname{erf}\left(\frac{v_E - v_{min}}{v_0}\right) + \operatorname{erf}\left(\frac{v_E + v_{min}}{v_0}\right) - \frac{4v_E}{\sqrt{\pi}v_0} \exp\left(-\frac{v_{esc}^2}{v_0^2}\right) & \text{if } v_{min} \leq v_{esc} - v_E \\ \operatorname{erf}\left(\frac{v_E - v_{min}}{v_0}\right) + \operatorname{erf}\left(\frac{v_{esc}}{v_0}\right) - \frac{2(v_E + v_{esc} - v_{min})}{\sqrt{\pi}v_0} \exp\left(-\frac{v_{esc}^2}{v_0^2}\right) & \text{if } v_{esc} - v_E < v_{min} \leq \\ & v_{esc} + v_E \\ 0 & \text{if } v_{min} > v_{esc} + v_E \end{cases} \quad (1.40)$$

To actually estimate the rate and spectrum of WIMP-induced recoils on a given target, we need to specify a few parameters from independent measurements: ρ_D , the local density of dark matter; v_0 , the average speed of a WIMP particle (which also determines the dispersion or spread of velocities in the distribution); v_{esc} , the escape velocity of the galaxy; and v_E , the speed of the earth with respect to the rest frame of the WIMP gas. If the WIMP gas has negligible net rotation, which is assumed to be the case, then v_E is the orbital speed of the earth about the galactic center, which is the sum of the rotational velocity of the local position of the galaxy, the sun's so-called *peculiar* motion with respect to nearby stars, and the orbit of the earth about the sun. The WIMP dispersion velocity v_0 is also determined by the galaxy's rotation curve [40].

The local dark matter density, ρ_D , is determined primarily by the galaxy rotation curve, and can be further constrained by numerical simulation. Estimates for this parameter range from $0.2 \leq \rho_D \leq 0.4 \text{ GeV/cm}^3$, though there are arguments that it should be higher [40]. This number is very dependent on assumptions about the shape of the dark matter halo, and local variations (so-called "clumps" and "streams") in the dark matter halo profile can change this number by factors of 2 or more [41]. The community has largely adopted a value of $\rho_D = 0.3 \text{ GeV/cm}^3$ as canonical, and I will do likewise. The galactic escape speed v_{esc} is estimated to be 544 km/s, with uncertainties on the order of 50 km/s (90% confidence) [42].

Parameter	Definition	Value	Uncertainty	Ref.
ρ_D	local DM density	0.3 GeV/cm ³	\sim factor 2	[40]
v_{esc}	galactic escape speed	544 km/s	\sim 9%	[42]
v_0	dispersion velocity of DM temperature distribution	250 km/s	\sim 10%	[43]
\bar{v}_E	average relative speed of earth to DM halo	263 km/s	\sim 10% (from v_0)	[43]
Δv_E	amplitude of annual modulation to v_E	15 km/s	small	[43]

Table 1.1: Galactic halo and velocity parameters needed to determine the expected WIMP-induced nuclear recoil spectrum. See text for definitions of parameters.

The local rotation velocity is measured recently to be approximately 250 km/s [43], although many earlier measurements indicated a lower velocity of \sim 220 km/s. The small (relative to the galactic rotation velocity) velocity of the solar system radially and out of the galactic plane can be neglected, leaving 13 km/s along the direction of rotation [43]. We can likewise disregard the orthogonal part of the Earth’s rotation about the sun; if we also ignore the small eccentricity of the Earth’s orbit, it’s velocity along the direction of rotation can be approximated by $15 \cos(2\pi(t - t_0))$ km/s, where t is measured in years and t_0 , the maximal velocity, is approximately June 2 [43]. The parameter v_E is then given by

$$v_E = 263 + 15 \cos(2\pi(t - t_0)) \text{ km/s.} \quad (1.41)$$

The \sim 6% annual modulation in the velocity will have important implications for the expected signal rate, as will be shown later in this section. Table 1.1 summarizes these parameters.

Combining equations (1.30), (1.28), (1.40) and (1.41), we can calculate the expected WIMP-induced nuclear recoil rate for any spin-0 target, ignoring effects of detector resolution. Figure 1.7 shows the expected WIMP-induced nuclear recoil spectrum for the elements sodium, argon, germanium, and xenon (used in important dark matter searches) given the parameters in table 1.1, for a 100 GeV WIMP with a WIMP-nucleon cross section of 10^{-45} cm².

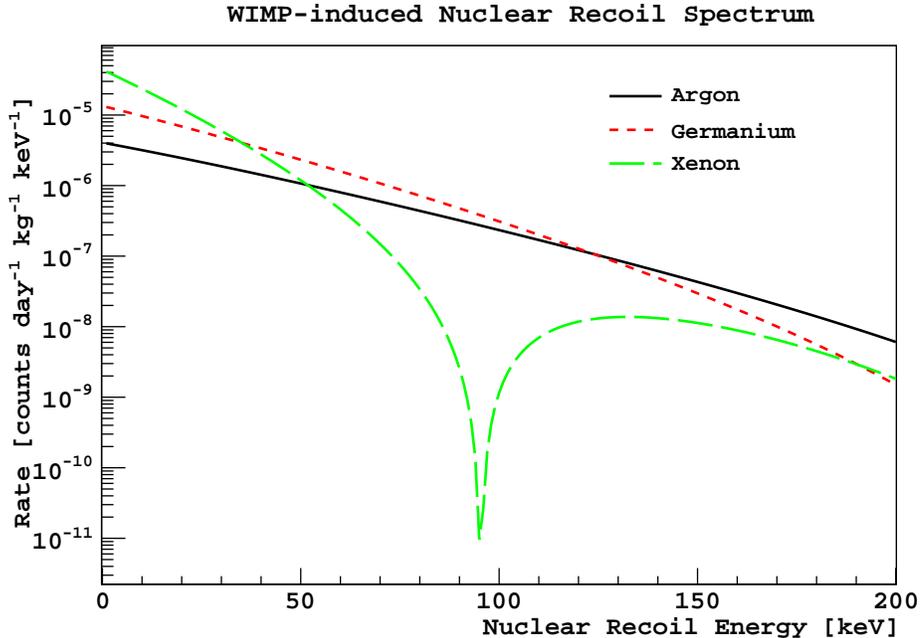


Figure 1.7: Predicted WIMP-induced nuclear recoil spectrum, for xenon, argon, germanium, and sodium targets, for a 100 GeV wimp with a scattering cross section of 10^{-45} cm², with galactic halo parameters from table 1.1.

The zeros in the spectrum arise from corresponding zeros in the form factor model, where the scattering amplitudes from the various nuclei interfere completely destructively.

Because of the very low rates involved (especially compared to typical backgrounds at low energies), direct dark matter searches at this stage don't hope to measure a spectrum from dark matter, but instead attempt to identify and count WIMP-induced recoils above some threshold (set by detector efficiency and backgrounds). Therefore, the most useful comparison between different targets is the integrated recoil rate above threshold, which is shown in Figure 1.8. Here again we see that the A^2 advantage of the heavier elements such as xenon and germanium is only realized if the detection threshold is low enough to take advantage of the full coherent scattering region.

As discussed previously, the speed of the WIMP wind with respect to a detector on earth is not constant, but varies throughout the year due to the earth's rotational motion adding constructively or destructively with the sun's motion about the galaxy, given approximately

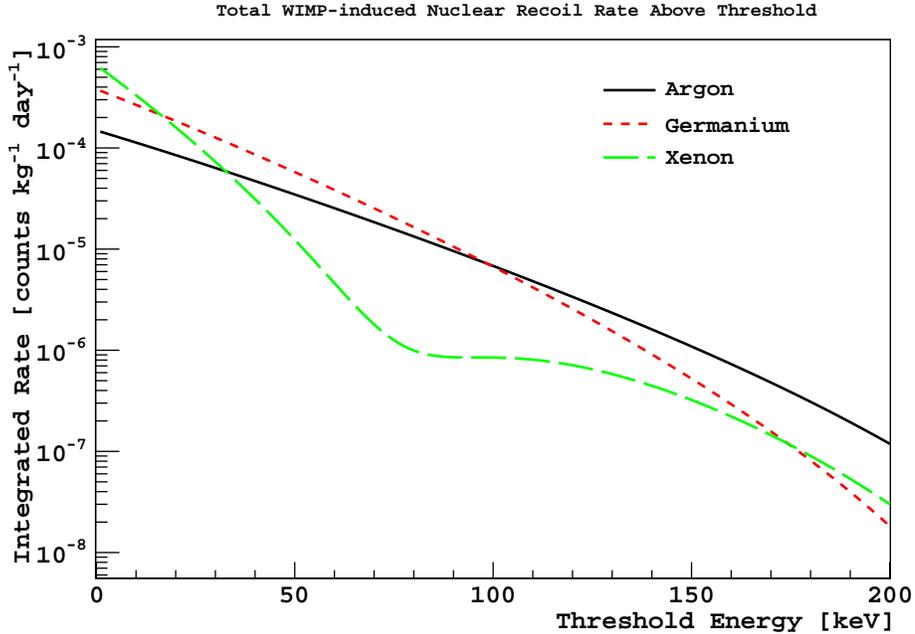
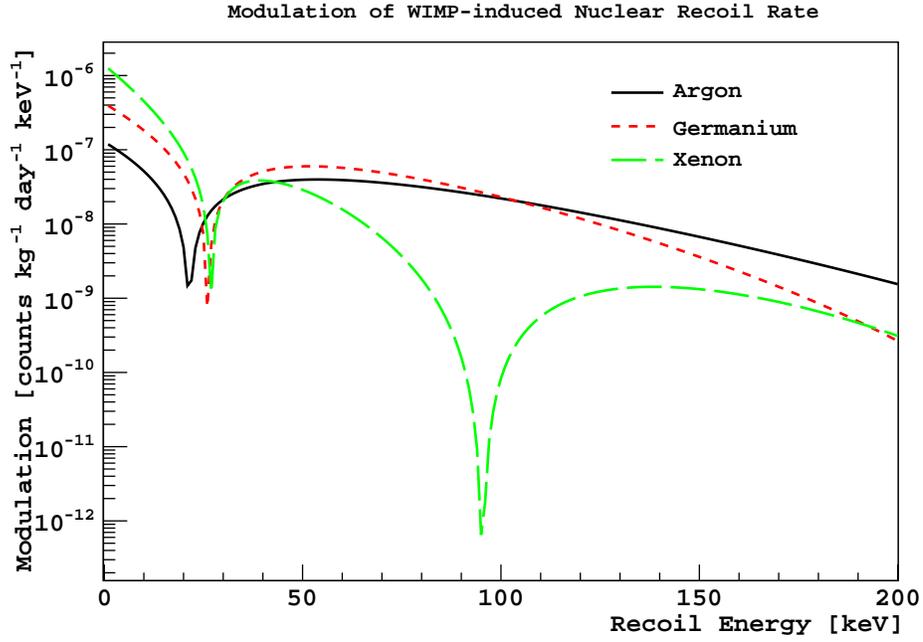


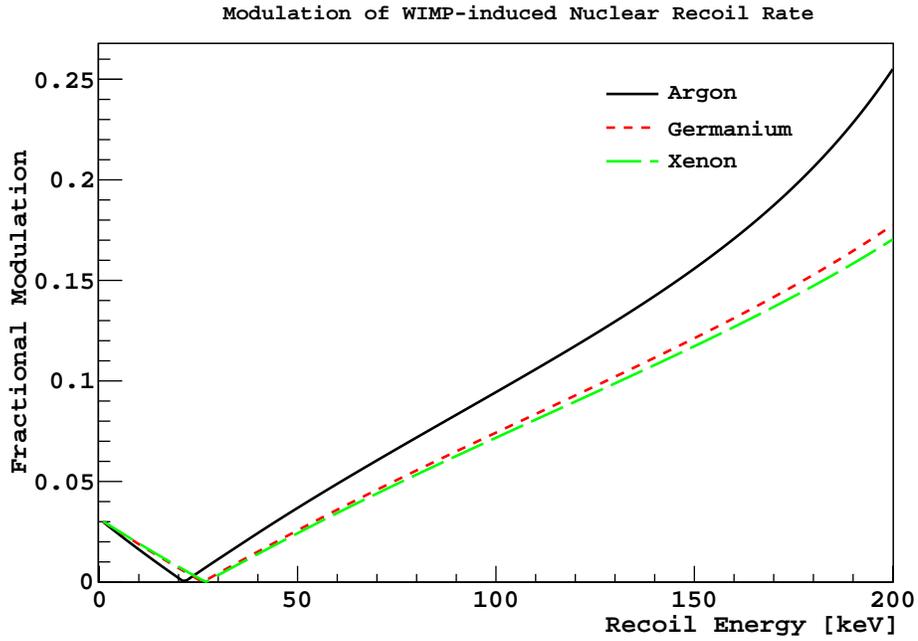
Figure 1.8: Predicted total integrated WIMP-induced nuclear recoil rate above threshold, for xenon, argon, germanium, and sodium targets, for a 100 GeV wimp with a scattering cross section of 10^{-45} cm², with galactic halo parameters from table 1.1.

by equation (1.41). This modulation in velocity will lead to a modulation in the rate and spectrum of WIMP-induced nuclear recoils observed by any terrestrial detector and provide important confirmation of any suspected WIMP signal. (Although, as we will see in Section 1.3.2, the confirmation may not be as straightforward as it seems!) Figure 1.9 plots the amplitude of the modulation of the WIMP spectrum, as the actual magnitude and as a fraction of the average signal, for the same WIMP parameters as previously. Figure 1.10 plots the modulation of the integrated count rate above threshold.

With these considerations in hand, we now have a realistic view of the signature of a WIMP interacting with various detector target materials: very low rate, low energy recoiling nuclei. We can now begin to design a detector which may be capable of observing these signals. In Section 1.3.2 we will briefly discuss other searches which are currently underway. Section 2.2 will discuss the merits of argon as a target for WIMP searches, and Chapter 4

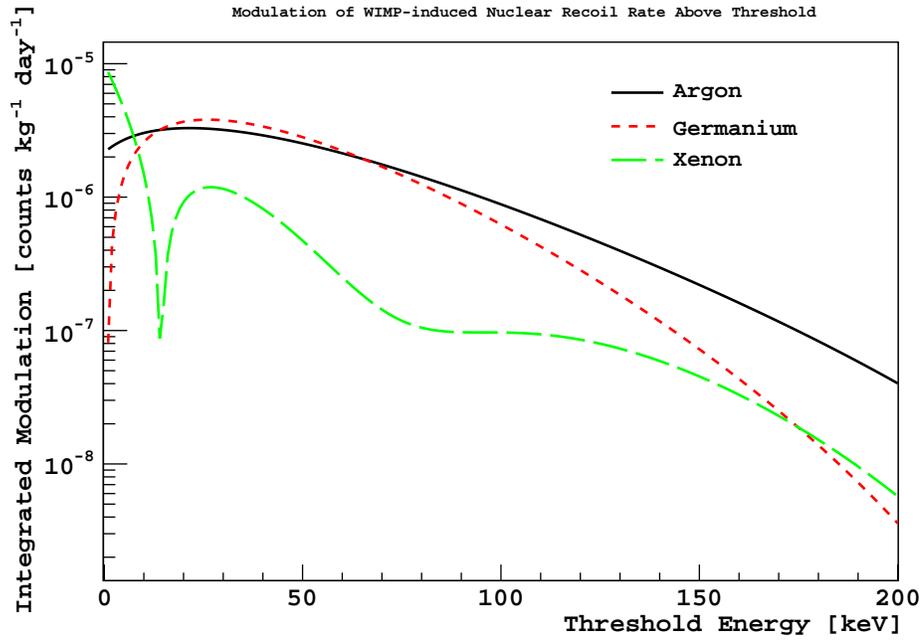


(a)

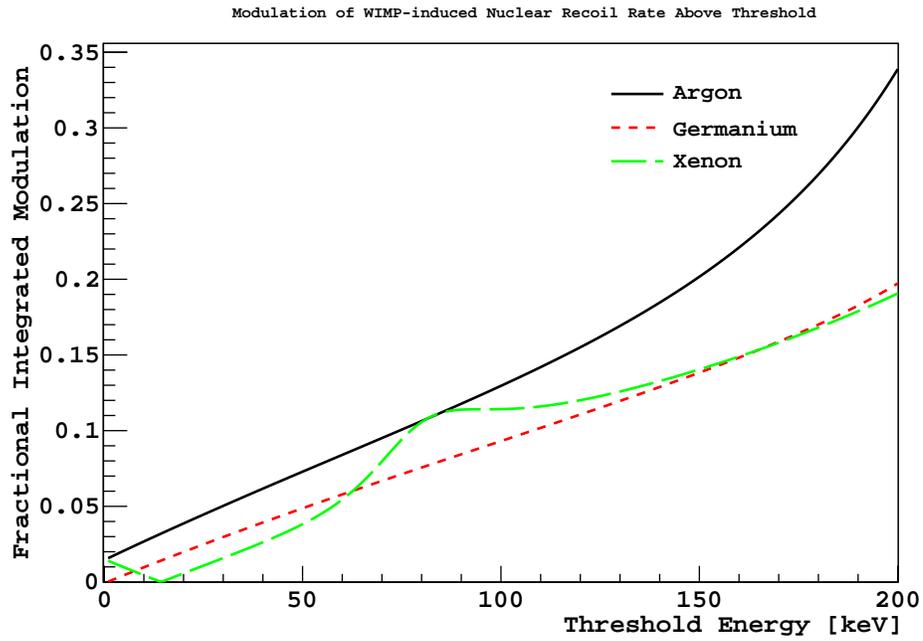


(b)

Figure 1.9: Spectrum of the amplitude of the annual modulation of WIMP-induced nuclear recoils for xenon, argon, germanium, and sodium, for a 100 GeV wimp with a scattering cross section of 10^{-45} cm², with galactic halo parameters from table 1.1. Figure (a) gives the actual amplitude of the modulation, while Figure (b) gives the amplitude as a fraction of the average signal.



(a)



(b)

Figure 1.10: Amplitude of the annual modulation of integrated WIMP-induced nuclear recoil rate above threshold for xenon, argon, germanium, and sodium, for a 100 GeV wimp with a scattering cross section of 10^{-45} cm², with galactic halo parameters from table 1.1. Figure (a) gives the actual amplitude of the modulation, while Figure (b) gives the amplitude as a fraction of the average signal.

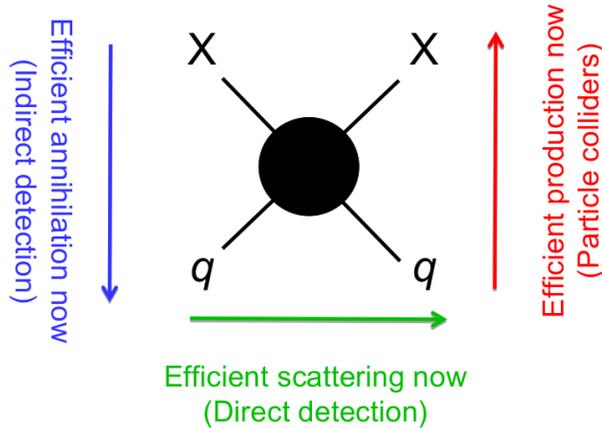


Figure 1.11: Relevant process for the three main WIMP detection techniques: indirect detection of annihilation (blue), direct detection from scattering (green), and production in colliders (red). For each process, the arrow indicates the direction of time [44].

will discuss the performance of a prototype argon detector. Finally, Chapter 5 will explore the sensitivities of a proposed series of full-scale argon-based detectors.

1.3 Current status of dark matter searches

Searches for WIMP-type dark matter present today use three separate complementary techniques: indirect detection, direct detection, and production in colliders. Figure 1.11 summarizes the process measured by these three techniques. Indirect detection experiments search for the high-energy products of annihilating WIMPs. Direct detection experiments, as described in the previous section, search nuclei recoiling from WIMP scatters. Finally, WIMPs can be created in colliders with sufficiently high energy and “observed” as a “missing mass” in the shower. This section will briefly cover the current status of direct and indirect experimental searches for dark matter.

1.3.1 Indirect detection

Indirect detection experiments can be grouped loosely according to the type(s) of particles that they detect. Neutrino telescopes, such as IceCube [45], measure the direction and

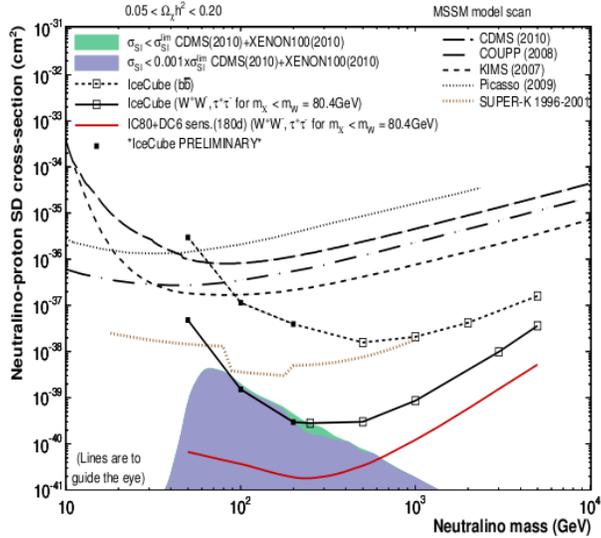


Figure 1.12: Spin-dependent cross section limits for neutralino annihilation into neutrinos set by IceCube [46].

energy of neutrinos from WIMP annihilations. To reduce the high neutrino background from atmospheric and other sources, these searches tend to focus on neutrinos emitted from the sun, but above the spectrum of neutrinos generated from solar fusion. WIMPs can scatter elastically in the sun and become gravitationally bound, over time creating a higher density in the sun than the surrounding halo. Eventually, the increased concentration of bound WIMPs will increase the annihilation rate until the capture and annihilation rates reach a stable equilibrium. Because this equilibrium concentration builds up over a very long time, the rate of annihilations is largely insensitive to variations in the density and structure of the local dark matter halo, unlike direct interactions searches. Similar arguments hold for any nearby large gravitational well, such as the earth or the center of the galaxy, although neutrinos from the direction of the sun are the primary focus. Figure 1.12 shows the current limits on the spin-dependent cross section from IceCube, under the two assumptions that WIMPs (specifically neutralinos) annihilate into neutrinos through W bosons or b quarks.

Other indirect detection experiments search for gamma rays or particle cosmic rays, usually from satellite or balloon-borne platforms. The most prominent of these experiments are ATIC [47], Fermi [48], HESS [49], PAMELA [50], and the recently launched AMS [51].

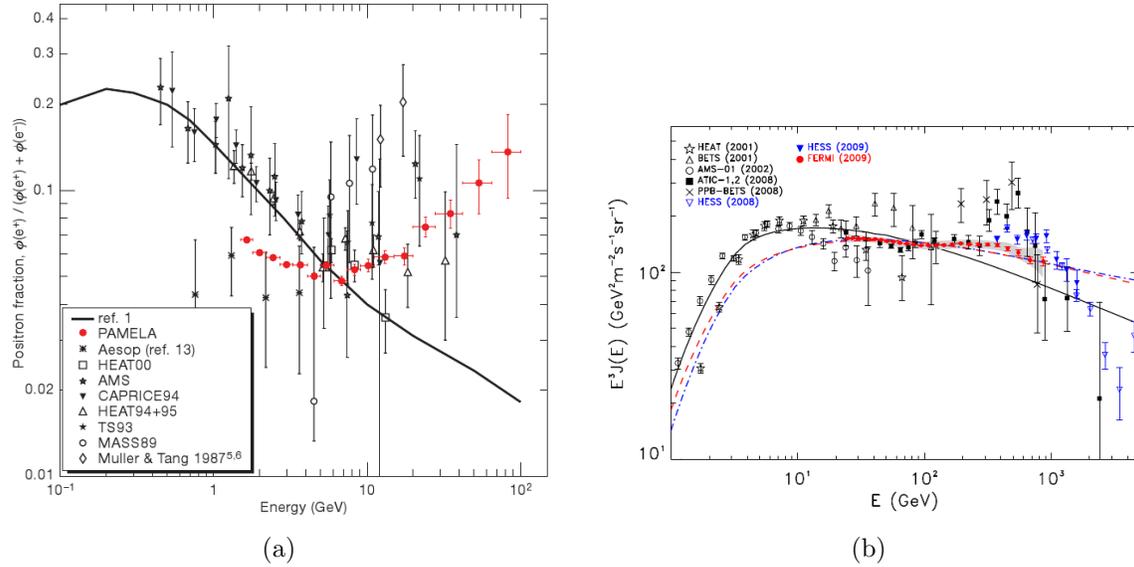


Figure 1.13: (a) Excess of high energy positron flux measured by PAMELA [53]. (b) Excess of high energy cosmic rays measured by ATIC, Fermi, and HESS [55]. The three curves (solid black and dashed colored) are attempts to fit the observed data with continuous astrophysical cosmic rays sources.

Unlike the neutrino observatories, which have not observed any excess neutrino flux above the expected background, the gamma and cosmic ray telescopes have multiple unexplained signals. In 2008, ATIC released the discovery of an excess of cosmic ray electrons with energies in the range 300–800 GeV, with a sharp peak at high energies, which could be explained by annihilating dark matter with an annihilation cross section of $\sim 10^{-33}$ cm² [52]. This result was followed shortly by measurements from PAMELA that indicated an excess positron fraction in the energies 1.5–100 GeV [53] without a corresponding excess in the anti-proton fraction [54]. Finally, Fermi and HESS released separate measurements of the cosmic ray spectrum that confirmed an excess above theoretical predictions in the same region as that seen by ATIC, but the Fermi and HESS excesses were smaller and less peaked than the ATIC measurement. The PAMELA measurement is shown in Figure 1.13(a), and the combined Fermi, ATIC, and HESS results are shown in Figure 1.13(b).

Taken together, these measurements strongly indicate some unpredicted source of high energy cosmic ray electrons and positrons with a higher positron fraction than predicted for

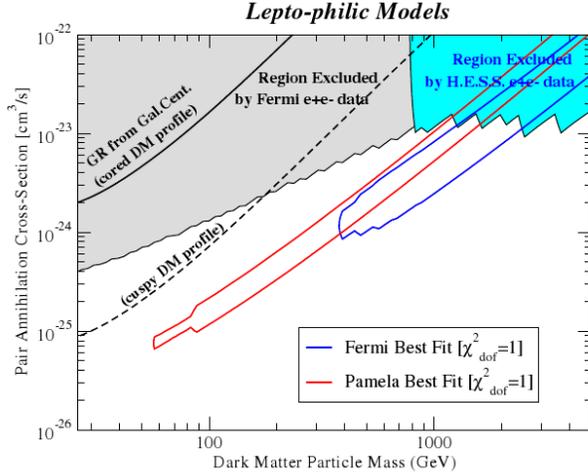


Figure 1.14: Best fit to the dark matter annihilation cross section from PAMELA and Fermi data, assuming that the DM particles annihilate equally into the three lepton species [55].

these energies. The source could be dark matter annihilations, but there are some difficulties with this idea. First, one must provide a mechanism for the dark matter to annihilate preferably into leptons (as opposed to quarks or baryons) to explain the large positron fluxes combined with predicted antiproton fluxes measured by PAMELA. If we ignore that issue for now, one calculation of annihilation cross section obtained from the measured fluxes is shown in Figure 1.14. The value obtained for the annihilation cross section is somewhat model dependent, but in most cases is around $\langle\sigma v\rangle \sim 10^{-24} \text{ cm}^3/\text{s}$ ($\sigma \sim 10^{-32} \text{ cm}^2$ since $v \sim 10^{-3}c$), a value 2–3 orders of magnitude higher than that required to obtain the correct relic abundance (See Section 1.2.1). Some have proposed that the present day annihilation cross section may effectively be larger today than it was at the time of freeze-out due to effects such as Sommerfeld enhancement (which increases the effective cross section at low velocities due to multiple exchange of some mediating particle) [56–58]. Others point out that the observed excesses could have astrophysical origins, such as nearby pulsars [55, 59]. Future high energy measurements from Fermi and AMS are expected to resolve many of these theories, in particular the dark matter versus pulsars debate, as the dark matter spectrum would have a sharp cutoff set by the WIMP mass.

1.3.2 Direct detection

The current limits on WIMP-nucleon interaction cross sections (assuming the interaction preserves isospin symmetry, see Section 1.2.2) are shown in Figure 1.15, for the case of spin-dependent (1.15(a)) and spin-independent (1.15(b)) interactions. Most direct detection searches are essentially “counting” experiments, employing low radioactivity materials and sophisticated background-rejection techniques, then hoping to observe an excess of nuclear recoils over the predicted remaining background. For spin-independent searches, the two leading experiments are Xenon100, which measures scintillation and ionization in liquid xenon using a time projection chamber similar to that described in Section 2.3, and CDMS, which measures ionization and phonon signals in superconducting germanium crystals. The two leading experiments for spin-dependent searches, PICASSO and COUPP, both employ bubble chambers with the pressure and temperature tuned so as to be insensitive to minimum ionizing events.

No statistically significant signal indicating WIMP interactions has been observed by any of these counting experiments that employ selection of nuclear recoil events. COUPP observed 3 candidate events in 28.1 kg-days of livetime with no specified estimation of background [60]; PICASSO observed 0 candidate events in 13.75 kg-days of livetime [61]; Xenon100 observed 3 candidate events with an expected background of 1.8 ± 0.6 events over 4850 kg-days of livetime [62]; and CDMS observed 2 candidate events with an expected background of 0.9 ± 0.2 events over 612 kg-days of livetime [63].

The null results of Xenon100 and CDMS are in tension with measurements obtained by the CoGeNT and DAMA/LIBRA collaborations. CoGeNT is a single very low background germanium crystal, similar to a single detector element of CDMS, but CoGeNT does not employ any nuclear recoil selection techniques. DAMA is constructed of several ultra-pure sodium-iodide (NaI) crystals triggered in anti-coincidence. DAMA also does not select for nuclear recoil events. Unlike all of the other experiments so far discussed, DAMA is not a counting experiment, but instead looks for a signal with an annual modulation consistent

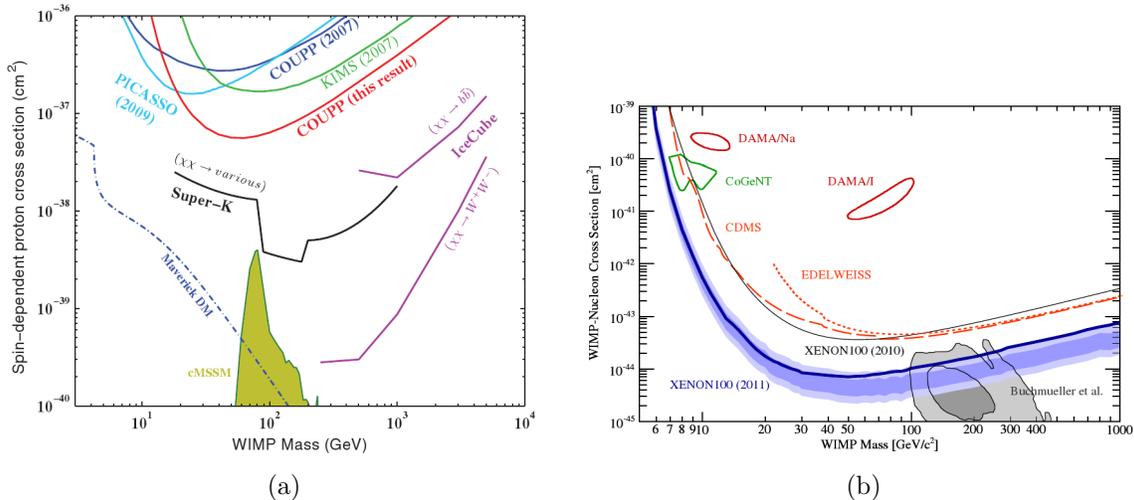


Figure 1.15: Current experimental limits on direct detection of WIMP-induced nuclear recoils, assuming that the interaction preserves isospin symmetry, for spin-dependent (a) and spin-independent (b) interactions [60, 62]. The shaded regions at the bottom of both figures are predictions of the lightest neutralino interaction cross section in the minimum supersymmetric model (MSSM).

with what one expects from WIMPs as the earth moves about the sun (see Section 1.2.2). Over measurement campaigns spanning 10 years, the DAMA collaboration has measured such a modulation with the appropriate phase and 1-year period to very high confidence level [64]. The amplitude of this modulation is consistent with interactions with WIMPs with masses of approximately 7 GeV and cross section $\sim 10^{-40}$ cm² [43, 65]. The CoGeNT collaboration has also measured a low-energy spectrum consistent with that due to a WIMP with the DAMA-favored parameters [66], and have very recently released a measurement of an annual modulation of their rate that is consistent at 1- σ with the DAMA modulation [67].

Both the DAMA and CoGeNT favored parameter spaces are excluded by the Xenon100 measurement. Many *ad hoc* theories modifying the “vanilla” WIMP interaction have been proposed in attempts to reconcile the measurements, but at this time they remain in conflict. Testing this controversy is a major goal of future WIMP search programs.

Chapter 2

Detecting Dark Matter: The DarkSide Program

2.1 Backgrounds to direct dark matter searches

As we showed in Section 1.2.2, the expected signature for a WIMP interaction is a recoiling nucleus with typical energy <100 keV, and with rates on the order of a few interactions per year per ton of target material or lower. With this kind of signature, nearly any low-energy particle interaction in the detector will present a background to dark matter searches. The difficulty of detecting dark matter is not so much finding a dark matter signal, but sufficiently reducing the overwhelming background. For the purposes of this discussion, this background can be divided into four rough categories:

1. β/γ radiation
2. surface α -particle radiation
3. radiogenic neutrons: neutrons from radioactive decays
4. cosmogenics: cosmic-ray muons and induced neutrons

For this discussion, the backgrounds are differentiated mainly by the techniques used to reduce them. Many of these removal techniques are similar across most dark matter ex-

periments. Intrinsic radiation, due to radioactive contaminants such as uranium, thorium, cobalt, and potassium, is reduced by constructing the detector of very radiopure materials. This reduces both the β/γ background and the radiogenic neutron background. Because the intrinsic rate of gamma decays in even the cleanest materials is well above the expected dark matter interaction rate, dark matter searches must also perform some active discrimination to reject these events in the detector. Since WIMPs produce recoiling nuclei, while gammas produce electron recoils, this discrimination is usually accomplished by exploiting the fact that recoiling nuclei are highly ionizing, while recoiling electrons are minimally ionizing. Therefore any observable that depends on the ionization density can be used as a handle to reject gamma-induced events. Neutrons, although generally creating a much smaller intrinsic background rate than gammas, are potentially a much more serious background, since they will also produce nuclear recoils indistinguishable from WIMP-induced signals.

Assembly of the detector in a radon-free clean room environment reduces surface alpha activity. Since, as we will show in Section 2.4.1, only alphas present on detector surfaces present a significant background, detectors with a monolithic target and some form of per-event position reconstruction can employ fiducialization to remove these events. That is, all events within some range of any surface are rejected, and only events within the accepted fiducial volume are considered as possible dark matter candidates.

To reduce cosmogenic activity, experiments are placed in deep underground sites, so that the muon flux from cosmic rays is attenuated by traveling through the earth overburden. The cosmogenic neutron flux is generated primarily by the muon flux. Very roughly, the muon flux with a flat overburden is attenuated by about an order of magnitude whenever the mass of overburden through which a cosmogenic muon must travel increases by the equivalent of a kilometer of water (km.w.e.) [5]. (The reduction is much more pronounced at lower depths, due to the low penetration of the low-energy portion of the cosmic muon spectrum). Working at depth has the added benefit (for dark matter searches, anyway) of stiffening the muon spectrum (i.e., weighting the spectrum toward higher energies), so that

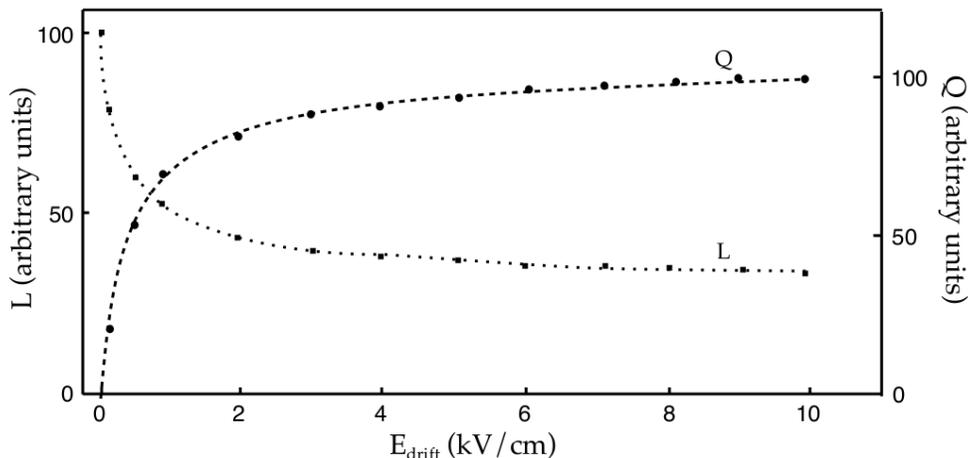


Figure 2.1: Relative amounts of scintillation (squares, L) and ionization (circles, Q) produced per minimum-ionizing energy deposit in liquid argon as a function of applied electric field [68].

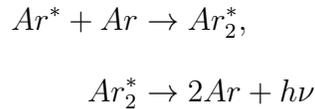
muons that do pass through the detector tend to deposit enough energy that they are outside the range of dark matter searches. Finally, experiments are surrounded by shielding, which can be either active or passive, to further reduce both cosmogenic fluxes and radioactive emissions from the lab environment.

Although the general approach to background reduction varies little across experiments, the specific implementation and efficiency, especially for gamma rejection, depend on the type of detector and the properties of the target. Therefore, in order to facilitate a more detailed discussion of backgrounds, I will first discuss the merits of liquid argon as a target, and describe in general one way to instrument argon for a dark matter search. Specific implementations and projected sensitivities of this type of detector, a two-phase argon time projection chamber (TPC), will be discussed in the remaining chapters. With this detector framework in mind, we can then return to the list of different backgrounds presented above and discuss the specific methods for background reduction.

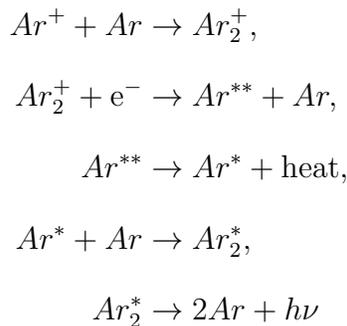
2.2 Argon as a dark matter target

Argon is an extremely efficient scintillator, comparable to NaI. Minimum ionizing particles produce about 40,000 photons per MeV of energy deposited [69], while more heavily ionizing nuclear recoils are quenched by a factor of ~ 0.25 [70, 71]. Because of this quenching factor, experimental results for argon will often use units of keV_{ee} for electron-equivalent energy or keV_{nr} for nuclear recoil equivalent energy, where $\text{keV}_{\text{nr}}=4\text{keV}_{\text{ee}}$ in the absence of an applied electric field. In the presence of an electric field, the light output for nuclear recoils is approximately constant, but the scintillation yield for minimum ionizing particles is reduced, following Figure 2.1. An efficient energy response is important for dark matter detectors because it allows for a lower threshold energy, which, since the WIMP-induced recoil spectrum decreases exponentially, greatly increases the search sensitivity.

The basic process of scintillation in noble liquids is well understood. Energy deposited in liquid argon will excite some atoms and ionize others, at an ionization/excitation ratio of ~ 0.21 [69]. The excited atoms produce scintillation by forming a weakly bound dimer (self-trapped exciton [72]):



and the ionized atoms first capture an electron and go through non-radiative deexcitation:



where the $*$ indicates an excited atomic or molecular state. Notice that the last two steps in the deexcitation path for ionized atoms are identical to initially excited atoms. The exciton trapping occurs on a scale of a few picoseconds. It has been proposed [72, 73] that excited atoms can also undergo a non-radiative biexcitonic quenching:



Since the stopping power or dE/dx for nuclear recoils is higher than that for electrons, the density of excited atoms (either directly or indirectly via ionization) will be higher, thus leading to a higher probability of biexcitonic quenching before dimerization. This mechanism is proposed largely to account for the observed light yield quenching factor between electron and nuclear recoils.

The excited dimers can form into either singlet or triplet states of approximately the same excitation energy, which decay by emitting a single VUV photon, sharply peaked around 128 nm. The singlet state has a lifetime of ~ 7 ns, while the triplet state has a lifetime of ~ 1.6 μ s. The lifetime of the individual states does not seem to depend on the ionization density, but the ratio of dimers formed into the singlet or triplet states does: the singlet to triplet ratio is about 0.3 for electrons, 1.3 for alphas, and 3 for heavy nuclear recoils [72]. The mechanism for generating the different populations of the two states is not well understood.

In addition to the scintillation photons, it is also possible to detect the free ionized charge from an energetic interaction in the argon. This is most commonly done by applying an electric field to the argon volume and collecting the free electrons at the anode. The probability that an initially ionized electron will be re-captured by its parent ion is primarily determined by of the average thermalization distance for the electron in the liquid and the effective Coulomb potential of the parent ion, and is largely independent of the event's energy and ionization density. Free electrons can also be captured by other ions, so the total

fraction of ionized electrons that fully escape is dependent on the density of ions created along a particle track. Therefore the proportional charge collected per unit of energy deposited in the argon will be smaller for interactions with higher ionization density.

As mentioned previously, one of the most critical features for a target for WIMP searches is a way to distinguish nuclear- and electron-recoil signatures. In noble liquids, especially argon, we can distinguish these interactions using two independent¹ variables: the ratio of singlet to triplet excitations, which is measured by the distribution of arrival times of detected scintillation photons, and the ratio of scintillation to ionization. In argon, the former method, pulse shape discrimination (PSD), is particularly powerful, because the singlet and triplet lifetimes differ by almost 3 full orders of magnitude. Compare this to, e.g., xenon, which has singlet and triplet lifetimes of ~ 4 and ~ 20 - 30 ns[72]. Overall, the PSD and scintillation-to-ionization discrimination result in a rejection of electron-recoil events by a factor of better than 10^8 [74–76]. This background rejection capability is the primary motivation for employing argon as a dark matter target. Of the noble liquids, argon is also promising for a number of other reasons:

Moderate mass (40 amu) Recall the approximate A^2 dependence on total WIMP scattering cross section. Argon is less sensitive than xenon ($A=131.3$), but more so than neon or helium.

Moderate cryogenic requirements Argon at 1 atm liquifies at ~ 87 K, which is obtainable with liquid nitrogen. Cryogenic containment systems for this temperature range are relatively simple compared to the 27 K necessary for neon, and commercial photodetectors are available.

Abundance Argon accounts for $\sim 1\%$ of the atmosphere, and so is easily obtainable at costs comparable to liquid nitrogen. This makes feasible the use of large amounts of liquid argon for cooling purposes, which can simplify many detector designs. The other noble gases are present only in trace amounts, and therefore have much higher cost.²

¹The independence of these variables will be discussed in Chapter 4.

²Unfortunately, this argument is largely negated by the need for depleted argon, discussed next.

Argon does have a few disadvantages as a detector medium however:

Scintillation wavelength The 128 nm scintillation photons emitted by argon are deep in the ultraviolet, and are outside the sensitive range of most commercial photodetectors and reflectors. This necessitates the use of waveshifters: fluorescent compounds that absorb the UV photons and re-emit them at longer wavelengths that are easier to detect.

³⁹Ar ³⁹Ar, a cosmogenically produced isotope found naturally in atmospheric argon, undergoes beta decay with endpoint energy 565 keV and a half-life of 269 years. The ³⁹Ar fraction in atmospheric argon leads to an event rate of 1 Bq/kg ^{nat}Ar [77].

³⁹Ar will be the primary limiting factor in using argon for dark matter searches. In a ton of liquid argon, which is the scale of the next generation of dark matter searches, ³⁹Ar represents a total background of around 3×10^{10} decays per year (20-30% of which are within the WIMP search window), compared to just a few WIMP signals. Although the discrimination power of argon to remove beta events is on better than 10^8 [74–76], this may not be enough to overcome the high rate of ³⁹Ar in natural argon. Moreover, even with perfect discrimination, the high rate of decays in a large detector will eventually lead to a large pileup fraction and dead time.

Argon-based dark matter searches to date have been small enough that the ³⁹Ar rate has not posed a significant problem, but some way to reduce the ³⁹Ar background will be necessary for the next generation of experiments. The DarkSide collaboration was formed with the intent to develop a series of detectors utilizing a naturally occurring source of depleted argon (i.e., with low ³⁹Ar content) found in underground gas wells. The development and characterization of underground argon will be discussed in detail in Chapter 3. For now, the relevant point of the underground argon is that it has at least 25 times less ³⁹Ar content than natural argon, which brings ton-scale detectors into the realm of possibility. To stay on track with our discussion of background reduction techniques, the next section will describe the liquid argon time projection chamber detector type used by DarkSide.

2.3 Liquid Argon TPC

Figure 2.2 shows a rough cartoon of the configuration and operation of a dual-phase argon time projection chamber (TPC). Figure 2.2(a) describes the geometry. The active target is a single monolithic volume of liquid argon. Above the liquid region, a volume of gaseous argon (light green in Figure 2.2) is maintained by the equilibrium vapor pressure. Photomultiplier tubes (PMTs), shown as the light blue cylinders in Figure 2.2(a) are used to detect the argon scintillation light. In the DarkSide TPC design, the active volume is segregated from the rest of the cryostat by a semi-sealed inner vessel; the top and bottom of the vessel are transparent to allow the scintillation light to reach the PMTs, which are in a non-active region of “buffer” argon. In other designs, PMTs may not necessarily be placed at both top and bottom, and the active volume may not be segregated (so the top PMTs may be in gas instead of liquid).

The segregation of the active volume from the rest of the cryostat is mainly for the purpose of cleanliness. Studies have shown [78, 79] that the scintillation and ionization signals in liquid argon are strongly quenched by the presence of impurities, particularly electronegative ones. Outgassing from surfaces, in particular the long lengths of cabling for the PMTs, is one of the major sources of these impurities. The physical segregation of the inner volume decreases the amounts of these impurities that can diffuse into the active region and degrade signals.

As mentioned previously, the scintillation light of argon is peaked sharply around 128 nm, which is outside the detection range of PMTs. Transmission and reflection rates for all but some very few materials are also quite low at these wavelengths. Consequently, the inner surface of the active volume is coated with a fluorescent material, so that it is always the first surface encountered by a scintillation photon. For argon, tetraphenyl-butadiene (TPB) and p-terphenyl (PTP) are commonly used. The fluorescent waveshifter absorbs the scintillation light and re-emits in the visible range. To maximize the amount of light collected, the part of the detector surface not instrumented with PMTs is made to be highly reflective.

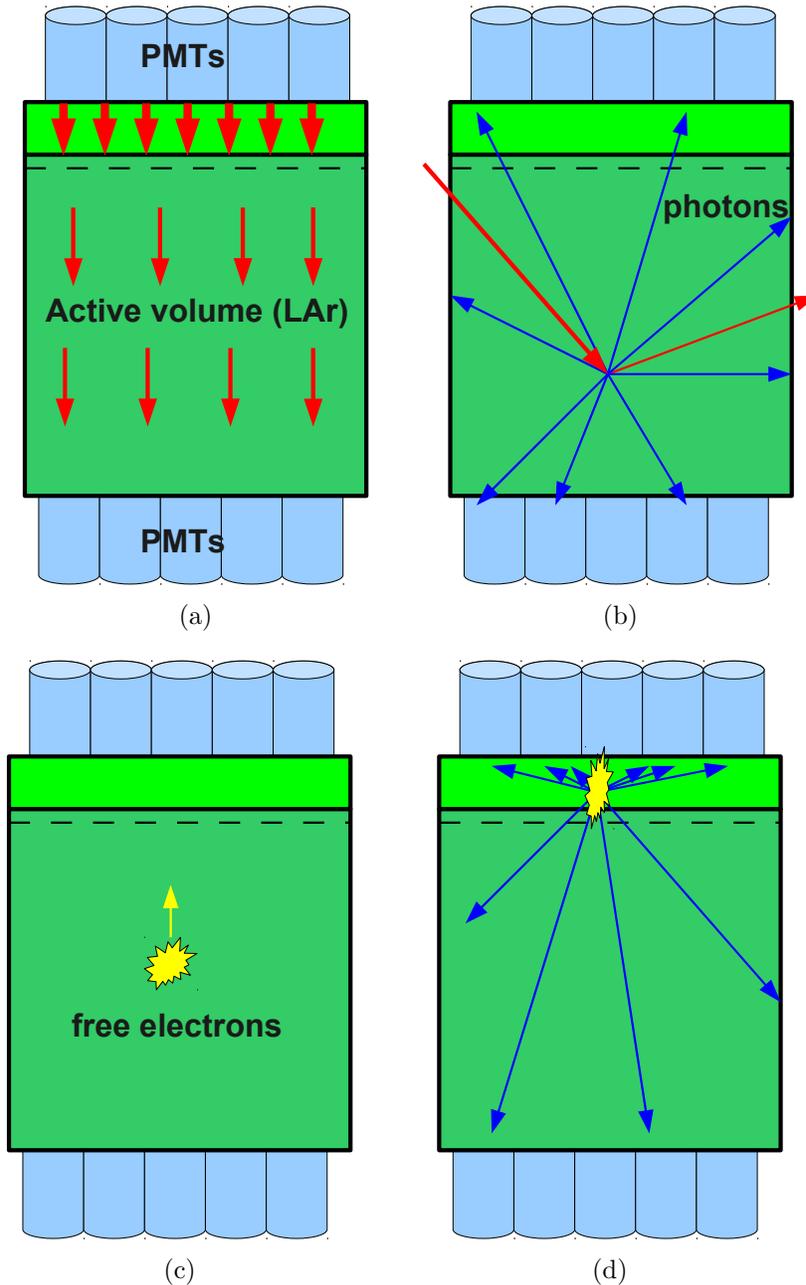


Figure 2.2: Cartoon of dual phase argon time projection chamber (TPC) operation. A thin layer of gaseous argon is maintained above the active liquid argon volume. An electric "drift" field on the order of 1 kV/cm is maintained in the liquid volume, and a higher electric "multiplication" field ($\sim 3\text{--}10$ kV/cm) is maintained in the gas region (a). An energetic particle interacts in the liquid argon volume, releasing prompt scintillation photons (b) ionized electrons, which are drifted upward (c) by the electric field, and are extracted into the gas volume, where they are accelerated and generate secondary scintillation (d).

A uniform electric field, around 1 kV/cm, referred to as the drift field, is maintained in the liquid volume, and a stronger field, around 5 kV/cm, the extraction and multiplication field, is maintained in the gas region (represented by the red arrows in Figure 2.2(a)). A wire grid, shown as a dashed line in Figure 2.2(a) allows for the discontinuity at the join of the two field regions.

When an energetic particle, e.g., a gamma or WIMP, interacts with the argon in the active volume, two signals are measured in the TPC. The particle deposits some energy in a small volume (directly in the case of alpha or beta decays, or indirectly, by recoiling of argon nuclei or electrons, in the case of gammas, neutrons, or WIMPs), which excites or ionizes some of the argon atoms in that region. Some of the ions recombine, and they and the initially excited atoms produce scintillation photons as described in Section 2.2 (Figure 2.2(b)). These prompt photons are first converted to the visible by the fluorescent waveshifter and then detected by the PMTs; we will call this signal S1. The number of photons detected in S1 is proportional to the total amount of energy deposited by the interaction (modulo quenching effects discussed in Section 2.2), and the singlet excitation fraction, which we will call $F_{singlet}$ can be determined from the photon arrival time distribution. $F_{singlet}$ is the variable used to perform background reduction via pulse-shape discrimination (PSD). Since the first photons are emitted within a few nanoseconds of the interaction, S1 also gives the time of the interaction.

Some of the electrons ionized by the initial interaction avoid recombination and are freed from the interaction site. The electric field maintained in the active volume causes those electrons to drift upwards, as shown in Figure 2.2(c). The drift speed is generally on the order of a few mm/ μ s and varies as a function of temperature and applied field as shown in Figure 2.3. Once the free electrons reach the top of the liquid region, the larger electric field there is enough to extract them from the liquid surface into the gas region. Here the field is strong enough that the electrons can gain enough energy between collisions to excite the argon gas atoms, producing a secondary scintillation pulse, which we will call S2, as shown in

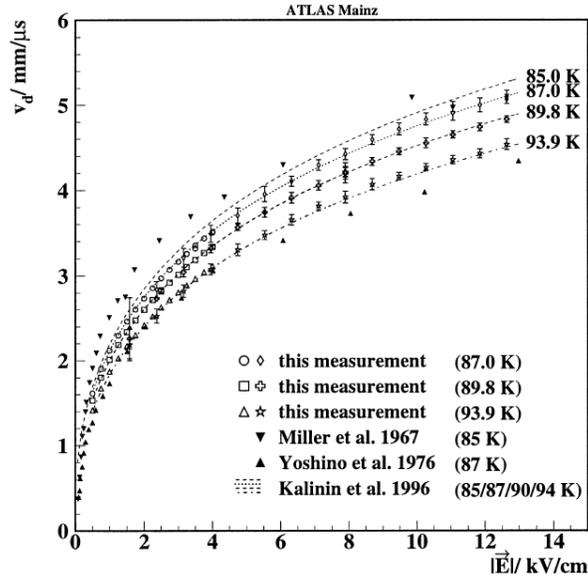


Figure 2.3: Free electron drift speed in liquid argon as a function of temperature and applied electric field [80].

Figure 2.2(d). The total number of photons collected for S2 is proportional to the number of ionized electrons that escaped recombination. Since the total amount of ionization is roughly proportional to the total energy deposited, the variable $S2/S1$ is particularly interesting. This gives the scintillation to (free) ionization ratio, which, broadly speaking, is independent of the energy deposited and depends only on the ionization density of the interaction (for given values of the fields, etc.).

The S2 signal also allows for a very accurate position reconstruction of the initial event site in 3D space. The z position relative to the liquid-gas interface can be determined by the drift time between S1 and S2. Since drift speeds are on the order of a few $\text{mm}/\mu\text{s}$, the z position can usually be determined with sub-millimeter accuracy. Since the top array of PMTs is very close to the S2 production area, the pattern of the fraction of S2 photons detected by each of the top PMTs can be used to determine the x - y position. The x - y resolution depends mostly on the “pixel” size and density, i.e., the size and placement of the PMTs, and the total amount of light in a given event.

So, with the dual-phase TPC, we can measure most of the parameters of interest for an interaction: the position and time of the interaction, the total amount of energy deposited, and two handles to evaluate the relative ionization density: $F_{singlet}$ determined from the S1 time profile, and S2/S1. With these variables defined, we can now proceed to explore the specifics of reducing different backgrounds to WIMP searches.

2.4 Background Rejection in an argon TPC

2.4.1 Surface alpha decays

Alpha decays are generally the easiest background to remove from WIMP searches. Intrinsic alpha radiation can come from two sources. The first is the presence of alpha-emitting isotopes, primarily in the uranium and thorium chains, contaminating the active volume. These in general do not present a significant source of background for argon, because (a) the alpha decay energies, even with quenching effects taken into account, are generally above the region of interest for nuclear recoils, and (b) it is generally possible to purify the active volume of these elements to the extent that the decay rate becomes negligible. In addition, some degree of separation between alphas and recoil events is obtained through pulse shape and S2/S1 discrimination. Alpha decays that take place in the detector material surrounding the active volume will not be detected, due to the very short (on the order of microns) travel range of alphas.

Alpha-emitting isotopes adsorbed on the sides of the container of the active volume can potentially be a much more serious background. Because of the extreme difficulty in excluding radon from even clean-room operations, higher concentrations of radon daughters, particularly polonium, tend to develop on all surfaces. Borexino, among the cleanest detectors in the world, measures an alpha-decay rate on the order of 10 cpd/m² of surface area. In approximately half of these decays, the alpha will be directed into the surface, causing

the heavier daughter nucleus to recoil into the active volume, mimicking perfectly the signal of a WIMP-induced recoil.

These surface recoils can be rejected in various ways. Since the recoiling nucleus and alpha particle are emitted back-to-back, if the nucleus enters the active volume, the alpha is necessarily pointing out. Since alpha decays have energies around 4-7 MeV, even a small amount of scintillation in the fluorescent waveshifter (always the innermost surface) will both alter $F_{singlet}$ and drive the apparent total energy from the S1 pulse above the region of interest for nuclear recoils. TPB has been measured to scintillate at approximately 900 photons/MeV of incident alpha energy [81], so a 4 MeV alpha will produce ~ 3600 photons, which is equivalent to a 360 keV nuclear recoil (using 10 p.e./keV_{nr} in argon). PTP has a similar or even better scintillation response.

A more active technique is to reject events whose reconstructed position is within some distance of any surface. The size of this cut depends on the position reconstruction resolution. Although fiducialization is a standard technique in low-background detectors, it is not ideal for an argon TPC. The x-y position reconstruction resolution is generally worst towards the sides of the detector, since the PMTs are all on one side, i.e., there are no PMTs at a radius greater than the wall. This necessitates a large cut in x-y to guarantee removal of surface events, which decreases the live mass and therefore sensitivity of the detector. Moreover, it is difficult to precisely calibrate the position reconstruction right at the surface, which leads to larger systematics in estimating the surface contribution to the background.

Nevertheless, with both the energy and an appropriately conservative fiducial volume cut, we can generally assume that the surface contribution to the total background will be negligible. The final size of the necessary fiducial cut must be determined individually for a given experiment based on position reconstruction calibration.

2.4.2 Radiogenic Neutrons

Neutron radiation is the most difficult background to exclude from WIMP searches. In essence, a neutron acts as a strongly interacting WIMP, producing a nuclear recoil that can not be rejected by pulse shape or S2/S1 discrimination. Radiogenic neutrons generally come from spontaneous fission of heavy elements, and (α, n) reactions on light elements. Therefore, the neutron emission rate for a given material is largely proportional to the contamination of uranium and thorium isotopes. Usually the most active source of neutrons in a liquid argon detector are the PMTs, which typically emit on the order of 20+ neutrons per year per channel for 3 inch PMTs [82]. Special low background PMTs can bring this background down to a few neutrons per year per PMT, but this is still the largest source in most detectors. Other large sources of neutrons are the cryostat holding the liquid argon, and the support structure necessary to hold the PMTs, electric field grids, etc. These are typically made of steel, which typically emits on the order of 1 neutron per year per kilogram, and copper, which typically emits on the order of 10^{-3} neutrons per year per kilogram [82]. Finally, naturally occurring radioactivity in the walls of the experimental hall will produce a large radiogenic neutron flux. This is rarely an issue, however, as even a modest amount of shielding is generally sufficient to bring the flux reaching the active volume to rates below that of internal components.

One method to reduce neutron background is to remove events with multiple recoils. The probability of a WIMP scattering multiple times in the detector is negligible, so any time two or more separate localized energy deposition sites are observed in the same event, it is almost certainly due to a neutron. The efficiency of this cut is limited primarily by the detector's position resolution and the total size of the active volume. For a typical liquid argon TPC employing tight-packed 3 inch PMTs at the top for x-y imaging, one should be able to distinguish two recoils separated by at least ~ 1 cm vertically or 15 cm laterally. (The large separation needed laterally is due to the coarseness of the imaging elements, even though one can typically achieve lateral resolution for single scatters on the order of a

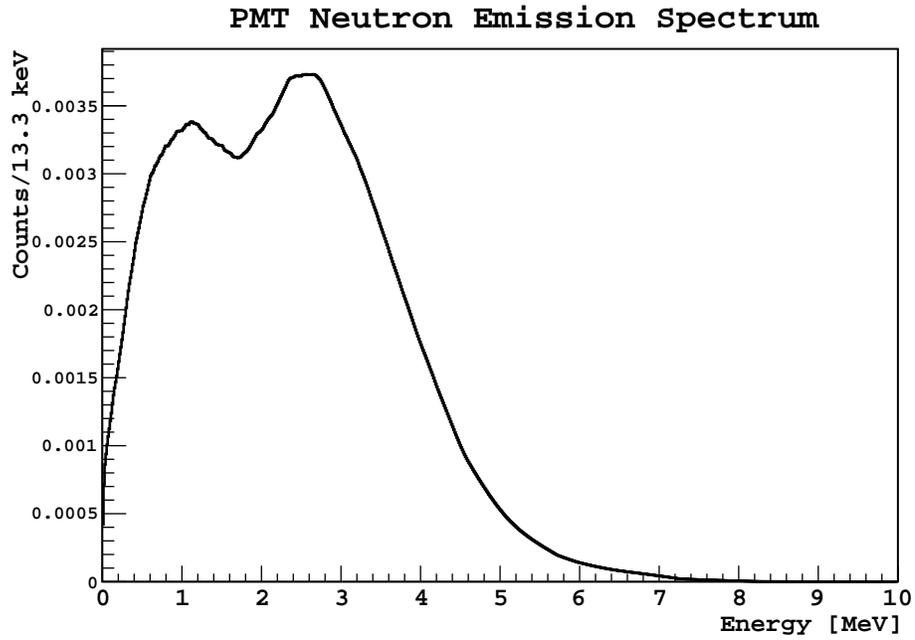


Figure 2.4: Spectrum of neutrons emitted from Hamamatsu R3998 PMT via fission and (α, n) processes, generated using Souces4A [82].

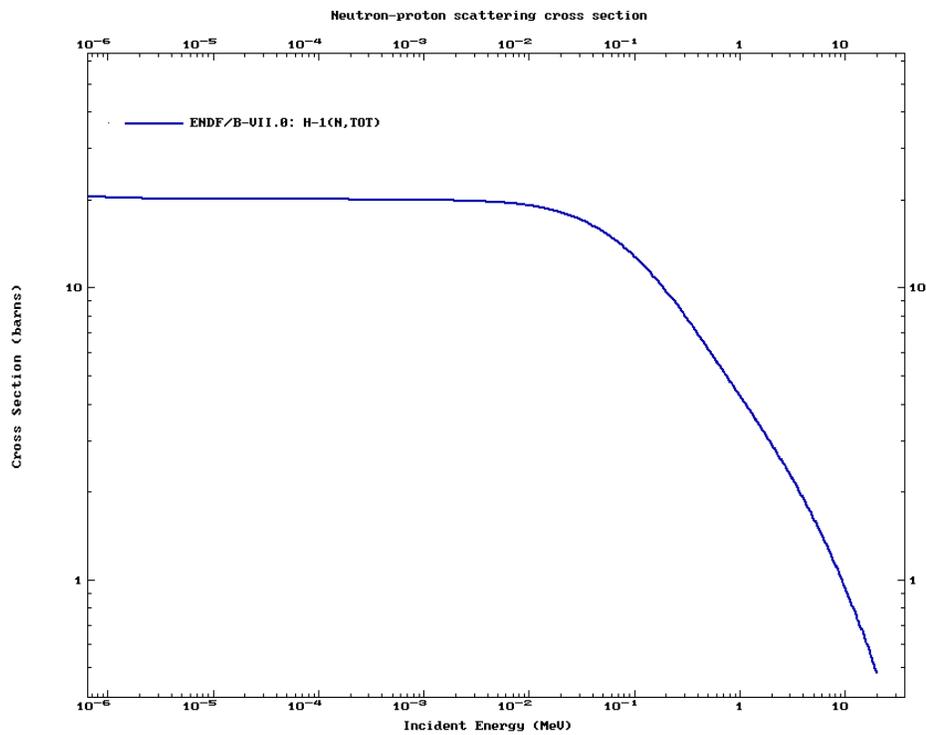


Figure 2.5: Total cross section for neutron scattering off hydrogen [7].

centimeter.) `GEANT4` simulations indicate that multiple recoil cuts with this resolution can typically lower the background by a factor of ~ 4 (see Chapter 5).

The background rate due to radiogenic neutrons from components very near the active volume can be further reduced by two somewhat mutually exclusive methods: passive shielding or active vetoing. Passive shielding relies on the fact that fission and (α, n) spectra are relatively low energy, typically below 10 MeV. The spectrum of the Hamamatsu R3998 PMT as calculated by `Sources4A` is shown in Figure 2.4 for example. The scattering cross section (off hydrogen) for neutrons of this energy is shown in Figure 2.5; scattering cross sections off other nuclei in this energy range are similar. Therefore, these neutrons are relatively easy to shield using materials with a high number of nuclei, especially hydrogen, per unit volume, such as polyethylene and acrylic.

Simulations with `GEANT4`, discussed in further detail in Chapter 5, have shown that acrylic very effectively attenuates neutrons of this energy, with an attenuation length of approximately 6 cm. Moreover, it is transparent to most optical wavelengths, allowing it to be placed in front of the PMTs. The total thickness that can be used is limited by the loss of light by attenuation in the acrylic and the loss in x-y position resolution due to the PMTs being farther from the S2 emission point. High density polyethylene is an even better neutron absorber, and can be used for shielding wherever optical transmission is not required, such as to line the inside of the cryostat. Passive shielding does have an issue with diminishing gains, however. For instance, if polyethylene is used to shield the active volume from the cryostat, the cryostat must be enlarged somewhat to accommodate the thickness of the shielding. But this increases the mass of and background from the cryostat, requiring more shielding, etc.

The alternative to passive shielding is active vetoing. In this case, little or no passive shielding is added between the active volume and neutron-emitting elements; rather, the entire detector is surrounded by a separate detector that serves as an active veto. Then neutron-induced events are tagged and rejected by detecting their thermalization and/or

capture in the veto (or in the active argon region). The efficiency of this method depends on the likelihood of detecting every neutron that enters the active volume. This requires that the amount of non-active material surrounding the active volume be minimized, since any neutron that captures on non-active material can't be vetoed. This is the reason that passive shielding and active vetoing are mutually exclusive: more passive shielding will reduce the effectiveness of the veto.

The optimal design of the active veto is not straightforward, but the most important feature is that the active element of the veto have a high cross section for both neutron scattering and capture, so that as many neutrons as possible are thermalized and captured inside the active region. The veto must also be instrumented to detect as often as possible the energy released by the capture, which is usually on the scale of a few MeV. Most neutron captures release gammas, which require tens of centimeters on average to deposit their full energy. Boron, in particular ^{10}B , with natural abundance $\sim 20\%$ [83], is a particularly effective veto target, as it has both a large neutron cross section and it produces an alpha particle on capture [7], which deposits all its energy in a few microns, reducing the necessary size of the veto. A major disadvantage of active vetoing with respect to passive shielding is that vetoing is minimally effective for gammas, which are not likely to deposit energy in both the active volume and the veto region. On the other hand, the intrinsic gamma rejection properties of argon make gamma backgrounds less serious than neutrons.

Active neutron vetoing offers a substantial bonus for background-free WIMP searches in the form of signal confirmation. That is, if a detector conducts a WIMP search and finds no events in the signal region, it can be difficult to demonstrate conclusively that the detector *could* have seen a signal if present. On the other hand, with an active veto, there will likely be a population of neutron-induced nuclear recoils that are removed only by the veto, both demonstrating the detector's sensitivity and providing a convenient real-time calibration source.

The first DarkSide detectors will employ an active neutron veto consisting of boron-loaded liquid scintillator.

2.4.3 Cosmogenic Muons and Neutrons

Cosmogenic muons in general do not directly present a serious background to dark matter searches. Even at very moderate depths, the typical muon energy is on the order of GeV, so, if it passes through the active area of the detector, it will deposit an amount of energy far in excess of the region of interest for nuclear recoils. They do generate neutrons via spallation, however, which can be a serious background. These neutrons are much higher in energy than radiogenic neutrons, and consequently much harder to attenuate by shielding. The most effective way to reduce the cosmogenic neutron background, then, is to reduce the flux of parent muons by placing the detector in an underground facility, where the muon flux is attenuated by the earth. A convenient parameterization of the neutron production rate and spectrum as a function of site depth is provided by [5], reproduced in Figure 2.6. The flux of cosmogenic neutrons at LNGS depth, where the DarkSide detectors will be constructed, is approximately $2.4 \text{ m}^{-2} \text{ day}^{-1}$, of which only $0.7 \text{ m}^{-2} \text{ day}^{-1}$ have energy above 10 MeV. At LNGS, the average cosmogenic neutron energy is around 90 MeV. At those energies, the total cross-section for scattering off argon is ~ 1 barn [84] ($1 \text{ b} = 10^{-24} \text{ cm}^2$), which would put the total interaction rate at a few times $10^{-33}/\text{s}$ per atom. Compare this to the rate of WIMP interactions of $\sim 10^{-34}/\text{s}$ per atom (at a cross section of 10^{-45} cm^2 , see Section 1.2.2): cosmogenic neutrons need to be attenuated by a factor of 10^2 - 10^3 to be background free at this sensitivity. Many of these interactions will be above the energy of the WIMP search window. The multiple recoil cut and active veto will also reduce the rate somewhat, but not as efficiently as for radiogenic neutrons, due to the cosmogenic neutrons' higher energy and therefore longer scattering length.

Neutrons produced in the rock outside the detector can be attenuated with sufficient shielding. Although other materials have better neutron attenuation, water is a typical

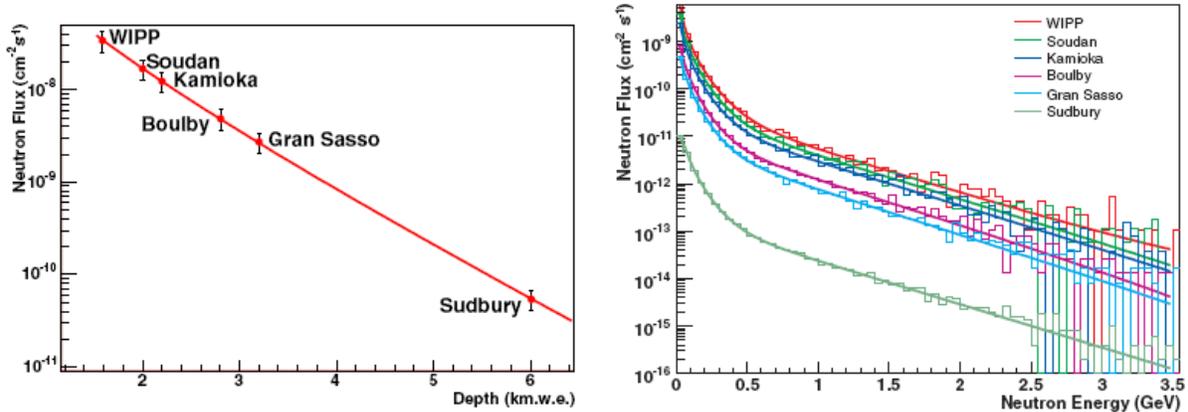


Figure 2.6: Left: Neutron flux emitted from cavern boundaries at various underground facilities. Right: Spectrum of neutrons flux. From [5]

choice for this type of shielding for several reasons, some of which are low cost and the ease with which it can be purified. The DarkSide detectors will employ a water shield. The total cross section for 95 MeV neutrons with water is 0.875 barns [84], so the mean free path for cosmogenic neutrons in water, which should be roughly comparable to the attenuation length, is around 30 cm, requiring a water shield $\sim 1 - 2$ m thick to attenuate the cosmogenic neutrons by the required factor of $10^2 - 10^3$. Neutrons passing through the outer water shield must still pass through the active neutron veto, which will further attenuate them, and also veto any that scatter in the veto volume.

Water is also an excellent shielding choice because it can serve as a muon veto by instrumenting it with photodetectors to detect the Cherenkov light generated by muons passing through the water. Muon detection efficiencies well over 99% have been demonstrated with this type of muon veto [85]. Muon veto capability is important for two reasons. First, the direction of cosmogenic neutrons is highly correlated with the direction of the parent muon at the time of production (although it is quickly randomized by scattering in rock), so many of these cosmogenic neutrons can be removed by detecting the parent muon passing through the water veto. Even if the parent muon itself doesn't pass through the veto, cosmogenic neutrons are most often produced as part of a hadronic shower and accompanied by other neutrons and charged particles, some of which may be detected by the muon veto.

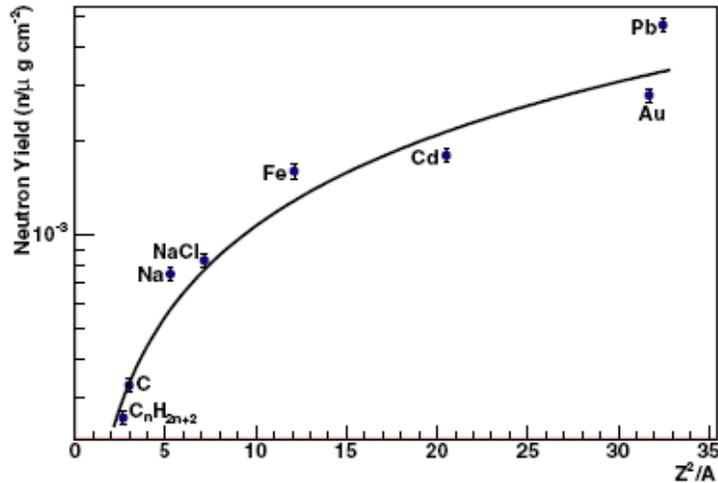


Figure 2.7: Parameterization of the rate of neutron spallation by cosmogenic muons for various nuclei [5].

Finally, an inert shield layer can only fix some of the problems of cosmogenic neutrons, since muons can also generate neutrons in the shield itself (or other detector components). Figure 2.7 shows a parameterization of the rate of neutron spallation (per incident unit of muon flux) for various materials. Since the spallation rate increases with atomic number Z , materials that provide better shielding from external cosmogenic neutrons will have a higher “internal” cosmogenic production rate. Since there is no way to provide shielding from the shield, it becomes almost essential to have a muon veto of some sort; having the veto double as passive shielding only makes sense.

2.4.4 β/γ Radiation

Gamma radiation is potentially the most serious background for most dark matter searches, because of the high rate and penetration of gammas generated both in detector materials and in the surrounding environment. In argon, the beta decay of ^{39}Ar can be lumped into this category, since the end product is the production of a moderate energy electron (less than a few MeV) in the active region.

In general, gammas from external sources, emanating mainly from the rock walls of the cavern, will be effectively attenuated by the neutron shielding and veto. The attenuation length in water of gammas of a few MeV, which dominate the flux, is on the order of 10 cm [86], so the external flux will be attenuated by $\sim 10^{-5}$ by a meter of water shielding. The flux of external gammas in Hall C of LNGS (where DarkSide-50 will be installed) has been measured to be $\sim 10^8 \text{ m}^{-2} \text{ day}^{-1}$ or $\sim 1 \text{ cm}^{-2} \text{ s}^{-1}$ [87]. So, the flux reaching the inner detector will be conservatively $\sim 10^{-4} \text{ cm}^{-2} \text{ s}^{-1}$. The cross section for gammas with energy of a few MeV is on the order of a few barns, so the background rate from external gammas after a moderate (1 m thick) water shield would be a few $10^{-2} \text{ kg}^{-1} \text{ s}^{-1}$, meaning that the ^{39}Ar rate in natural argon ($1 \text{ kg}^{-1} \text{ s}^{-1}$) should dominate by at least a factor of a few hundred.

Gamma backgrounds from internal components, primarily PMTs and steel, are not so easily dismissable. For now, we'll make the assumption that these backgrounds can also be brought below the rate of ^{39}Ar (this assumption will be justified in Chapter 5). Even if ^{39}Ar is the only significant source of beta/gammas, and it is underground argon depleted by a factor of 25, it creates a background rate of $\sim 10^4 \text{ kg}^{-1} \text{ day}^{-1}$, compared to the $\sim 10^{-4} \text{ kg}^{-1} \text{ day}^{-1}$ or less WIMP interaction rate.

As previously discussed, these backgrounds must be removed by some active discrimination technique. In an argon TPC, both pulse-shape discrimination ($F_{singlet}$) and the scintillation to ionization ratio (S2/S1) provide a means to differentiate between electron-like and nuclear-recoil like signatures. Using S2/S1 only, The WARP collaboration demonstrated that a reduction in the background due to minimum ionizing particles of $\sim 10^{-2}$ is possible while accepting upwards of 99% of nuclear recoils [74]. This reduction is thought to be approximately independent of the detector light yield or interaction energy.

In contrast, the effectiveness of a cut on $F_{singlet}$ will have a strong dependence on the average total number of S1 photoelectrons detected. Roughly estimating the effectiveness of a pulse shape cut on removing β/γ backgrounds from a dark matter search is a straightforward task. Recall that we can define $F_{singlet}$ for every event as $F_{singlet} = \frac{I_S}{I_S + I_T}$, where I_S and

I_T are the total number of dimers excited into the singlet and triplet states, respectively. Experimentally, this value can be determined by digitizing the signals from the PMTs detecting scintillation photons and fitting the pulse to the sum of two exponentials convolved with a gaussian representing the detector response function. Ideally, $F_{singlet}$ should follow a binomial distribution $B(x; n, p)$, where x is the fractional amplitude of the fast component, n is the number of detected photoelectrons, and p is the most probable fast fraction for a given species. At high energies, $p_\beta \approx 0.23$, and $p_{recoil} \approx 0.76$ [72]. At low energies, as dE/dx goes up for β/γ particles and down for nuclear recoils, deviation from this high energy limit is expected.

Because fitting individual scintillation shapes is not always straightforward and because of the large amount of processing time required for a large batch of events, it is convenient to define approximations to $F_{singlet}$ that are easier to obtain with real data. One common such approximation is F_{prompt} , the fraction of scintillation light detected in the first t_0 time of S1, defined as

$$F_{prompt} = \frac{\int_0^{t_0} S(t) dt}{\int_0^\infty S(t) dt},$$

where $S(t)$ is the PMT signal amplitude at time t . Figure 2.8 shows the mean value of F_{prompt} evaluated at 90 ns for electrons and nuclear recoils as a function of incident energy, as measured by [75], which highlights the deviation in $F_{singlet}$ at low energies due to changing dE/dx . Since 90 ns is such a common choice for the F_{prompt} evaluation time, we will often use the variable $F90$ to refer to F_{prompt} evaluated at 90 ns.

If a cut is made at $F_{prompt} = f_{cut}$ to remove β/γ events from the nuclear recoil search, the fraction of minimum ionizing particles that will remain as background is simply the integral of the F_{prompt} (binomial) distribution for $x \geq f_{cut}$. The cumulative distribution function of a binomial is equivalent in the continuum limit to a regularized incomplete beta function :

$$\sum_0^K B(k; n, p) = I_{1-p}(1 - K, n - 1) \quad (2.1)$$

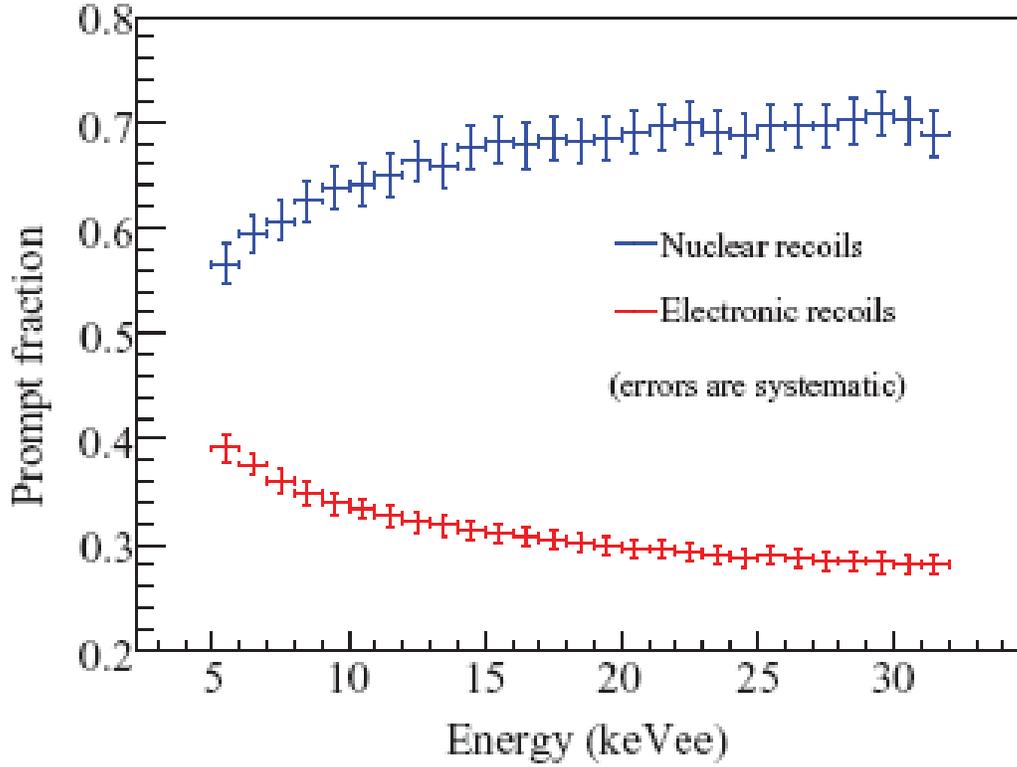


Figure 2.8: Plot of the mean value of F_{prompt} vs. energy (electron equivalent) for electrons and nuclear recoils, as measured by [75]. Here the quenching factor for nuclear recoils is 0.29.

where

$$I_x(a, b) = a \binom{a+b-1}{b-1} \int_0^x t^{a-1} (1-t)^{b-1} dt. \quad (2.2)$$

If we ignore the dependence on energy of the mean of F_{prompt} , we can write, for events with detected number of photoelectrons n :

$$A_\beta(n) = 1 - I_{1-\overline{F}_\beta} \left(n \cdot (1 - f_{cut}), n \cdot f_{cut} + 1 \right) \quad (2.3)$$

where A_β is the acceptance fraction for betas (which we would like to minimize), \overline{F}_β is the mean F_{prompt} for a beta particle, given by the “prompt fraction” in Figure 2.8. Likewise, we

can write for the nuclear recoil acceptance fraction

$$A_{recoil}(n) = 1 - I_{1-\overline{F}_{nr}} \left(n \cdot (1 - f_{cut}), n \cdot f_{cut} + 1 \right). \quad (2.4)$$

Figure 2.9 plots Equations (2.3) and (2.4) as a function of light yield for several choices of f_{cut} , taking $\overline{F}_{\beta} = 0.27$ and $\overline{F}_{nr} = 0.71$, the apparent asymptotes of Figure 2.8. As the figure shows, the effectiveness of the cut (i.e., the fraction of beta-like events removed) decreases exponentially with the number of photoelectrons detected. This means that backgrounds leaking through the cut are dominated by events near the analysis threshold (i.e., with few photoelectrons detected). Figure 2.10 shows the fraction of beta-like events passing the cut as a function of the nuclear recoil acceptance fraction, for various thresholds.

Since the performance of the F_{prompt} cut is so dependent on number of photoelectrons, the threshold in terms of photoelectrons is essentially fixed. The only practical way to lower the threshold in terms of energy, then, is to increase the detector light yield. Since the WIMP sensitivity depends exponentially on the threshold energy, it is critically important to have the light yield be as high as possible.

The ideal binomial distribution for F_{prompt} discussed above is useful to see the general dependence on parameters. In reality, however, the acceptance fractions will also contain extra factors to account for other sources of random fluctuations, such as the Poissonian statistics involved in the detection of single photoelectrons. In an effort to estimate how much the binomial distribution underestimates the width of the F_{prompt} distribution, we perform a simple Monte Carlo model of a more realistic distribution. To generate this distribution, we perform the following steps:

- Choose the number of scintillation photons emitted from a gaussian distribution with $\sigma = \sqrt{N}$.

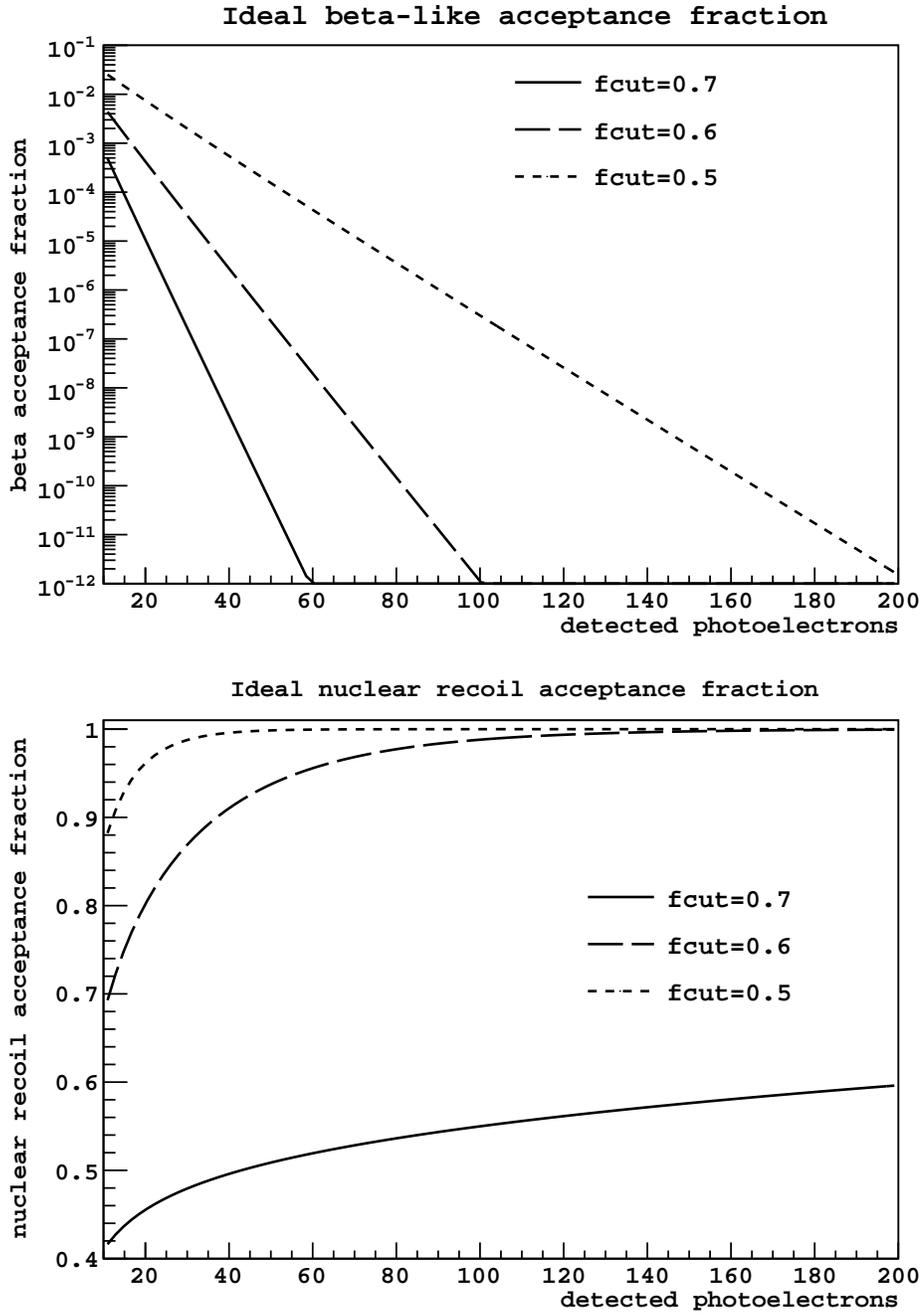


Figure 2.9: Ideal fraction of events accepted into a nuclear recoil search window based on a cut on the F_{prompt} parameter, for beta-like events (top) and nuclear recoils (bottom), assuming the mean F_{prompt} is 0.27 for betas and 0.71 for recoils.

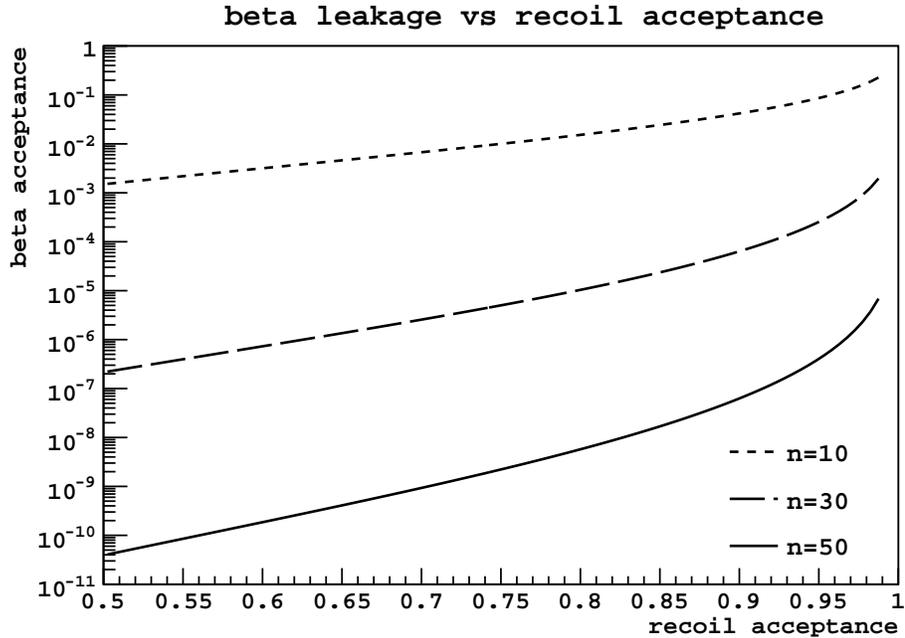


Figure 2.10: Fraction of beta-like events passing an F_{prompt} cut as a function of the nuclear acceptance fraction, for events with the indicated number of photoelectrons detected.

- Choose the number of photons emitted from singlet states from a binomial distribution with $p = 0.228$ for betas or $p = 0.693$ for recoils. (These *singlet* probabilities give the mean $F90$ for Figure 2.8.) The others are from triplet states.
- Choose the photon detection probability from a gaussian with a $\sigma = 5\%$. (I.e., assume that the light yield can vary by event position, etc).
- For each photon population (singlet and triplet), choose the number of detected photons from a binomial distribution, with p from the previous step.
- For photons from triplet states, some of them will arrive within the F_{prompt} integration window, on average $p = 1 - \exp(-t_0/\tau_3)$, with t_0 the F_{prompt} evaluation time (90 ns) and τ_3 the triplet lifetime (1.6 μs). From the triplet population, choose the number that arrive in the early portion from a binomial with p just described.
- For both populations, smear the detected number of photons by a gaussian with width $\sigma = 0.4\sqrt{N}$ to account for the width of the single photoelectron response in photomultipliers.

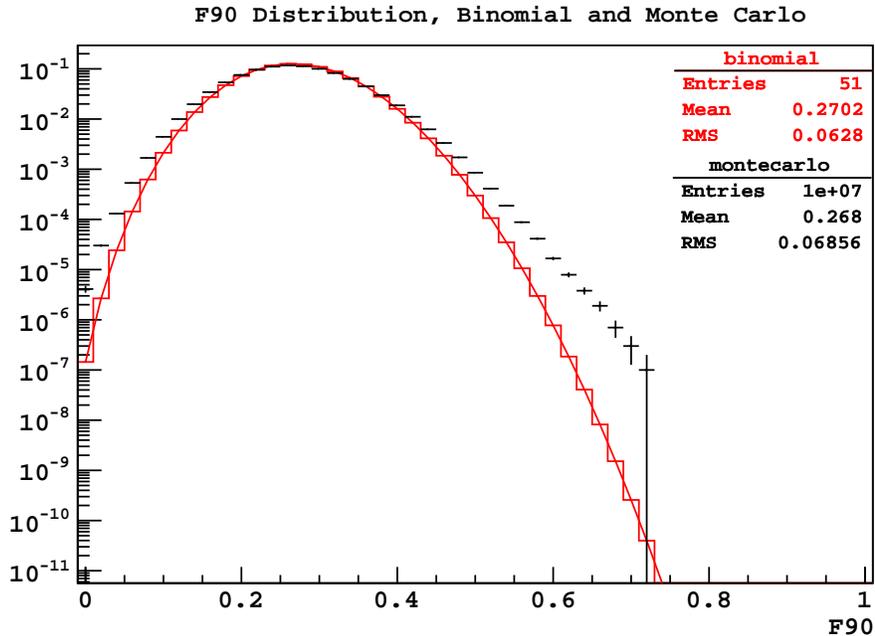


Figure 2.11: Idealized binomial distribution of F_{prompt} for electron-like events with 50 detected photoelectrons compared to a Monte Carlo distribution that takes into account multiple sources of random variation.

The results of this simulation for 10^7 events with 50 average detected photoelectrons along with the binomial distribution are shown in Figure 2.11. As the figure shows, the binomial PDF models the *shape* of the Monte Carlo distribution reasonably well, but the Monte Carlo has significantly larger tails. Since many of the parameters that went into the Monte Carlo distribution must be determined experimentally (e.g., light yield spread, width of single photoelectron, etc.), quantitative *a priori* determination of the effectiveness of an F_{prompt} cut is not practical.

2.4.5 Backgrounds Summary

In the preceding sections, we have laid out a general plan for sufficiently lowering backgrounds with a two-phase argon TPC to be sensitive to WIMP interactions, and thus the basic design of the DarkSide detectors. Minimum-ionizing beta and gamma backgrounds are reduced by factors of $> 10^8$ by pulse shape and scintillation to ionization discrimination. Alphas

and their recoiling parent nuclei, which occur primarily at detector surfaces, are identified and removed by reconstructing their positions and accepting only events within a fiducial volume. Radiogenic neutrons are tagged and rejected by a boron-loaded scintillator active veto. The rate of cosmogenic neutrons is reduced by placing the detector deep underground in the LNGS cavern. The remaining cosmogenic neutron flux emanating from the cavern walls (and environmental gamma flux) is attenuated by a large water shield. Cosmogenic neutrons that are generated inside the detector and shields are vetoed by detecting the passage of the parent muon through the water shield, which is instrumented to detect the muon's Cherenkov light. In Chapter 5, we'll revisit these general ideas with exact numbers: detector and shield dimensions, quantitative estimates of radioactive contaminations and decay rates, and Monte Carlo simulations of the effectiveness of the shielding and vetoes.

At present sensitivities, all of this is possible with standard argon derived from the atmosphere. However, the viability of future generations of argon detectors is dependent on the availability of a source of argon depleted in ^{39}Ar , which we have taken as a given up to this point. Chapter 3 describes the process we have employed for extracting and characterizing depleted argon from underground gas wells. With this in hand, all the pieces of the DarkSide program are in place, and we can realistically plan the next generations of detectors.

Chapter 3

Depleted Argon

3.1 The need for underground argon

For the reasons specified in Section 2.2, argon is potentially a very powerful target for observing WIMP interactions, because it is possible to distinguish very clearly WIMP-induced signals from minimum-ionizing interactions, the dominant backgrounds. However, commercial argon, generally obtained by concentrating (via cryogenic distillation or similar processes) the $\sim 1\%$ of argon gas that is naturally present in the atmosphere, contains the radioactive isotope ^{39}Ar . ^{39}Ar undergoes beta decay with a Q-value of 565 keV and has a half life of 269 years. Argon derived from the atmosphere contains ^{39}Ar at a level of $(8.0 \pm 0.6) \times 10^{-16}$ g/g, which corresponds to a decay rate of 1.01 ± 0.08 Bq/kg [77].

While this may be a small rate when compared to ordinary materials, in the low-radioactivity environment of a clean detector, ^{39}Ar quickly becomes the dominant source of background. Even with perfect background rejection, the dead time and pileup due to this rate quickly become unmanageable as detector size increases. As a quick example, if we require a valid WIMP signal to have only one interaction within a $200 \mu\text{s}$ window (a reasonable request for the type of detector described in section 2.3), then the effective dead time due only to ^{39}Ar would be 20% in a detector with a 1 ton liquid argon target.

Therefore, it is critical for large-scale low-energy argon detectors to obtain argon that has a reduced concentration of ^{39}Ar compared to that of atmospheric argon. One possible method to produce this depleted argon would be to remove the ^{39}Ar from atmospheric argon isotopically, via centrifugation [88] or thermal diffusion [89, 90]. However, an initial investigation of these methods by the WARP collaboration determined that obtaining sufficient quantities of depleted argon in this way would likely be prohibitively expensive and time-consuming.

If producing depleted argon is not feasible, then the only alternative is to find a naturally occurring source. ^{39}Ar is created primarily by spallation of cosmic ray particles on ^{40}Ar in the upper atmosphere, dominated by the reaction $^{40}\text{Ar}(n,2n)^{39}\text{Ar}$ [91]. ^{40}Ar , on the other hand, is primarily produced in the β -decay of ^{40}K . Because there are negligible amounts of potassium compounds in the atmosphere, the bulk of ^{40}Ar production takes place underground. The majority of this gas diffuses out and into the atmosphere, but some fraction remains trapped underground. Because the earth strongly attenuates the flux of cosmic rays, the primary creation channel (spallation) of ^{39}Ar is largely eliminated. Therefore, it is reasonable to expect that pockets of argon gas trapped underground and isolated from atmospheric mixing might have a greatly reduced ^{39}Ar concentration.

However, underground, ^{39}Ar can also be produced by stopped muon or neutron capture onto ^{39}K ($^{39}\text{K}(n,p)^{39}\text{Ar}$) [92]. Because of this extra channel, it is possible for underground argon to contain even higher concentrations of ^{39}Ar than atmospheric argon [93]. Because potassium is necessary to produce the ^{40}Ar , reduced ^{39}Ar concentrations can only be produced in rocks with a low free neutron flux. Free neutrons are produced underground both cosmogenically, due to interactions from high energy cosmic ray muons, and radiogenically, from spontaneous fission and (α,n) reactions primarily from uranium and thorium decay products. Therefore, one expects to find argon gas with the lowest fraction of ^{39}Ar in wells that are deep (to attenuate the cosmogenic neutron flux) and with a low concentration of uranium and thorium (to lower the radiogenic neutron flux). Measuring the rocks is not sufficient, however, because gas pockets can easily migrate away from their points of ori-

gin. Therefore, direct measurement of the ^{39}Ar content of any well is the only way to even estimate the depletion factor.

Rather than develop all the necessary infrastructure to mine gas wells ourselves for underground argon, the DarkSide collaboration opted to attempt to recover unused argon already present in the exhaust streams of existing commercial gas mining facilities. An initial survey of several sites found trace amounts of argon — on the order of a few hundred ppm — in many of these sites, coming from a variety of different sources with different gas compositions. Because of the very low argon concentration, we determined that it was necessary to develop a method to concentrate the argon on-site, rather than shipping large amounts of waste gas to a central facility. In order to attempt to sample the ^{39}Ar fraction at several sites, it was necessary that this preconcentration plant be portable enough to move from site to site without having to be completely rebuilt every time. After shipping, the preconcentrated gas could be further refined by other methods.

The technology chosen for the preconcentration stage was Pressure Swing Adsorption (PSA), which will be discussed in detail in Section 3.2. Section 3.3 will cover details of the development of a prototype PSA system, and section 3.7 will report on the successful deployment of this system systems to collect initial samples of depleted argon and measurements of their depletion factors. Finally, Section 3.8 will describe the full-scale two-step PSA and cryogenic distillation purification operation.

3.2 Pressure Swing Adsorption

In this section I will cover in detail the theory and operation of the pressure swing adsorption (PSA) process. For a very simple summary, PSA plants separate gases by exploiting the different adsorption strengths of different gases on certain adsorbent molecules such as zeolites and activated charcoals. Because the relative amounts of adsorption depend on the total pressure, one can regenerate the adsorption beds by altering the pressure and using some of

the purified product gas to purge the bed¹. In the standard plant, two adsorption beds (or columns) are operated 180 degrees out of phase, so that a continuous output of purified product is achieved. Because PSA plants work on pressure differences alone, the only operating costs and requirements are the power needed to run the compressor or vacuum. Moreover, the plant need not be very large (one of the primary uses for PSA plants is the production of purified oxygen for medical and home use) and can be made to be relatively portable. These characteristics make PSA technology ideal for deployment on-site at remote gas mining and processing facilities, for both limited sample collection and eventual high-volume processing.

The primary reference on PSA processes is Ruthven's Pressure Swing Adsorption [94], which contains very detailed, if arcane, explanations of the theory and applications of PSA. Although a full treatment of PSA theory requires solving explicitly the differential equations governing mass balances of gas in the adsorbent beds, many of the key details of the process can be understood by studying the much simplified model, often dubbed the Bi-Linear Isotherm (BLI) theory, presented by Knaebel and Hill [95] and experimentally verified by Kayser and Knaebel [9].

The basic premise of BLI theory is that the adsorption isotherm (the relation describing the amount of gas adsorbed as a function of pressure at fixed temperature) for a given gas species is linear and independent of the concentrations of other gases. Thus,

$$q_i = K_i P_i \tag{3.1}$$

where q is the molar concentration of gas in the adsorbed phase, P is the partial pressure, and K , which has units of [mol·volume⁻¹·pressure⁻¹], is called Henry's constant and is a function of temperature. This approximation is only valid for most gases in a fairly narrow range of pressures and temperatures, so care must be taken when applying it. Figure 3.1 is a plot of

¹When the low pressure step is below atmospheric pressure, the process is sometimes referred to as VPSA for vacuum PSA.

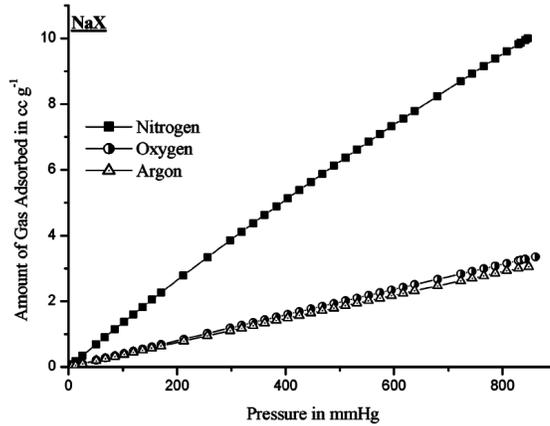


Figure 3.1: The adsorption isotherms of nitrogen, oxygen, and argon on zeolite NaX [8].

the adsorption isotherms of nitrogen, oxygen, and argon on zeolite NaX, which are in good agreement with the linear approximation in the pressure range shown.

PSA is generally performed using columns packed with beads or pellets of porous adsorbent, such as zeolite. The shape and size of the pellets can be characterized by the parameter η , which is the void volume or volume of the column not contained in the molecular pores. (The actual volume of the molecules of the cage is considered to be negligible.) In most cases 0.5 is a good estimate for η . The total amount of gas needed to fill a column of volume V at partial pressure P_i and temperature T packed with adsorbent characterized by η and K_i is therefore

$$n_i = \frac{P_i V \eta}{RT} + q_i V (1 - \eta) = \frac{P_i V}{RT} (\eta + (1 - \eta) K_i RT) = \frac{P_i V \gamma_i}{RT} \quad (3.2)$$

where $\gamma_i = \eta + (1 - \eta) K_i RT$.

Hence a column of volume V filled with adsorbent will have a different effective volume $V \gamma_i$ for each species of gas. In particular, if we introduce a mixture of two gases A and B (here I will adopt the standard convention whereby A always refers to the more strongly adsorbing or heavy species, hence $\gamma_A > \gamma_B$) at one end of a column, species B will propagate through the column more quickly than species A because it has a smaller effective volume to fill up. There are complications not addressed by this model if there are pressure variations

in the column, but it is generally valid to assume that the system comes to local equilibrium instantly.

In particular, if we have a column pre-pressurized with pure “light” (B) component and introduce our A-B mixture, species B propagates faster, and we can not only recover all of the gas initially in the column but also some extra pure component B extracted from the input stream. We can calculate exactly how much gas is extracted as follows. First, define y_F as the fraction of species A in the feed stream, so $P_A = Py_F$. Extraction of pure light component can continue until species A has filled up its effective column volume, which requires a number of moles:

$$n_A = \frac{y_F PV \gamma_A}{RT} \quad (3.3)$$

Since species B is mixed in with species A in the feed stream, the total number of moles of B introduced into the column is

$$n_B = \frac{1 - y_F}{y_F} n_A = \frac{(1 - y_F) PV \gamma_A}{RT} \quad (3.4)$$

But the number of moles of B required to fill its effective column volume is only

$$n_B^* = \frac{(1 - y_F) PV \gamma_B}{RT} = \beta n_B \quad (3.5)$$

where $\beta = \gamma_B/\gamma_A < 1$. Therefore, we can extract during the process a number of moles of B equal to

$$n_F = n_B - n_B^* = \frac{PV \gamma_A}{RT} (1 - y_F) (1 - \beta) \quad (3.6)$$

where the subscript F refers to the feed step described below. In principle, the gas extracted during this step is pure B component. This will generally be true so long as the column is initially filled with pure B and the rate of gas flow through the column is high enough such that the speed of propagation of the A “front” is higher than the diffusion rate of the A gas in the adsorbent bed, a condition that is generally held.

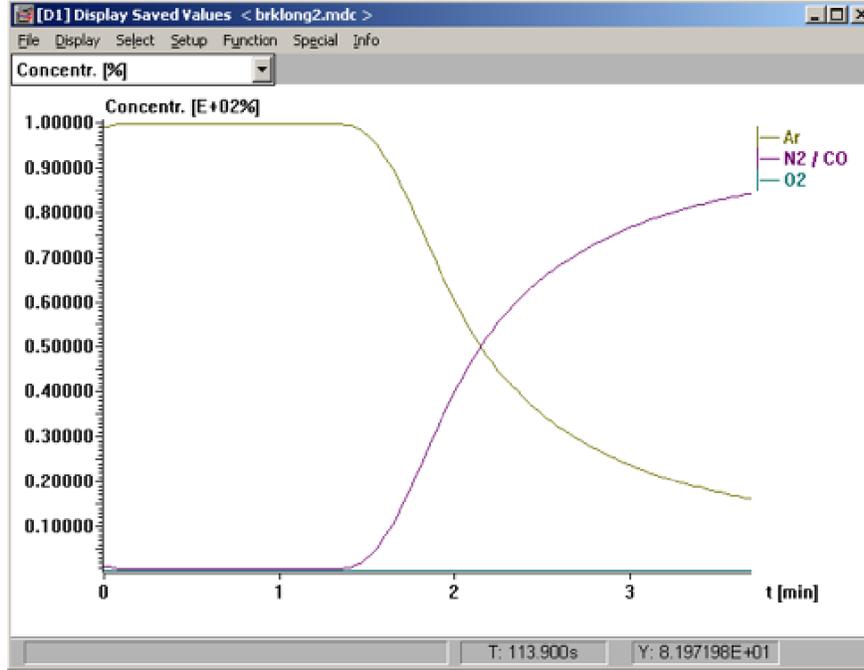


Figure 3.2: Example breakthrough plot, made with a 1% argon and 99% nitrogen cocktail fed into a bed pre-pressurized with pure argon.

Measurements of this quantity can provide a test of the accuracy of BLI theory and the validity of linear isotherms. An easier measurement to check is the breakthrough time, or time at which species A first appears in the output stream. The breakthrough time is defined to be

$$t_B = V\gamma_A/\dot{Q} \quad (3.7)$$

where \dot{Q} is the total volumetric flow rate of the A-B input mixture. Kayser and Knaebel [9] have used this technique to measure the Henry constants of various gases with good agreement with equilibrium adsorption measurements. Figure 3.2 is a plot of a breakthrough measurement we made with our system, showing the characteristic sharpness of the breakthrough curve.

With this framework, we can now begin to study a cyclic PSA process that produces a continuous output of essentially pure component B. Many different variations of PSA cycles have been studied, but we will focus on the cycle depicted in Figure 3.3, a six step cycle employing pressurization with feed.

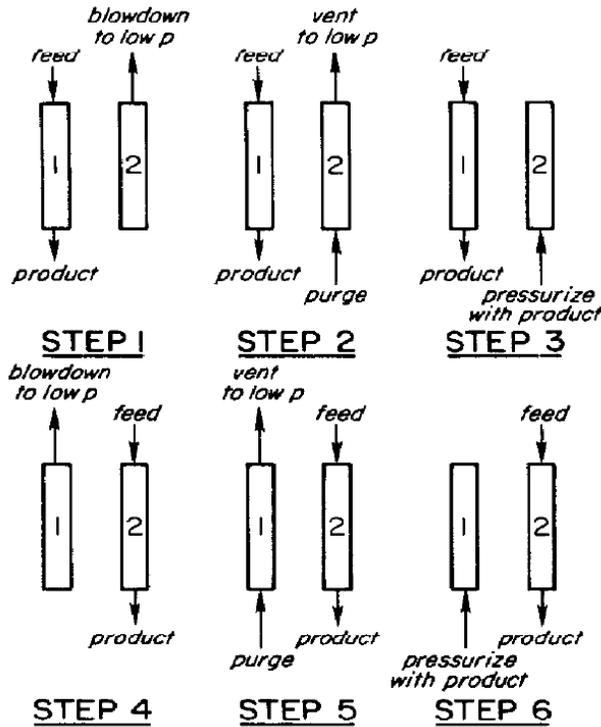


Figure 3.3: Basic six step PSA cycle employing pressurization with feed [9].

We assume that the process operates between two pressures, a high pressure P_H (which is generally the pressure of the input stream) and a low pressure P_L , which may or may not be sub-atmospheric. Another useful variable to define here is the pressure ratio $\varphi = P_H/P_L$. At the beginning of step 1 in Figure 3.3, column 1 is pressurized to P_H with pure B component and column 2 is pressurized to P_H with gas of the feed composition. During step 1, feed gas is fed into column 1 at the feed end and pure component B is extracted at the product end, while column 2 is depressurized (or undergoes blowdown) from P_H to P_L . Once column 2 reaches P_L , step 2, the purge phase, begins, where a fraction of the pure B component extracted from column 1 is redirected through the product end of column 2, which purges the A component out the feed end. Once column 2 is fully purged, the valve at the feed end of column 2 is closed, and a fraction of column 1's output stream is used to pressurize the column with pure component B. The step times are arranged so that the breakthrough time

(sometimes also referred to as feed time) is just slightly greater than the sum of the three step times. Steps 4-6 are identical to steps 1-3 with columns 1 and 2 reversed.

We can use the BLI theory to calculate the expected recovery of component B in this process. We have already calculated the number of moles extracted during the feed step n_F in equation (3.6). From this we must subtract the moles of B required to purge the column. If we assume no mixing between the gases in this step, this is equal to the amount needed to totally fill the effective volume of component A at low pressure,

$$n_{PU} = \frac{P_L V \gamma_A}{RT} \quad (3.8)$$

Unlike during the feed step, the assumption of no mixing in the purge step does not generally hold. Therefore real-world applications will require either a larger purge fraction or some A present in the column at the beginning of the feed step (and therefore in the output stream). We must also subtract the number of moles needed to pressurize the column from P_L to P_H ,

$$n_{PR} = \frac{V \gamma_B}{RT} (P_H - P_L) \quad (3.9)$$

Finally, we must account for the gas in the column at the beginning of step 1,

$$n_{HP} = \frac{PV \gamma_B}{RT} \quad (3.10)$$

Combining all these equations and defining the recovery R to be the moles of gas extracted divided by the moles of gas input, we get for steady state operating conditions

$$R_{SS} = \frac{n_F + n_{HP} - n_{PU} - n_{PR}}{n_B} = (1 - \beta) \left(1 - \frac{1}{\phi(1 - y_F)} \right) \quad (3.11)$$

Our initial talks with various gas manufacturers led us to believe that one of the most promising waste streams containing underground argon could be reasonably approximated as a binary mixture containing approximately 99.9% nitrogen (component A) and 0.1% argon

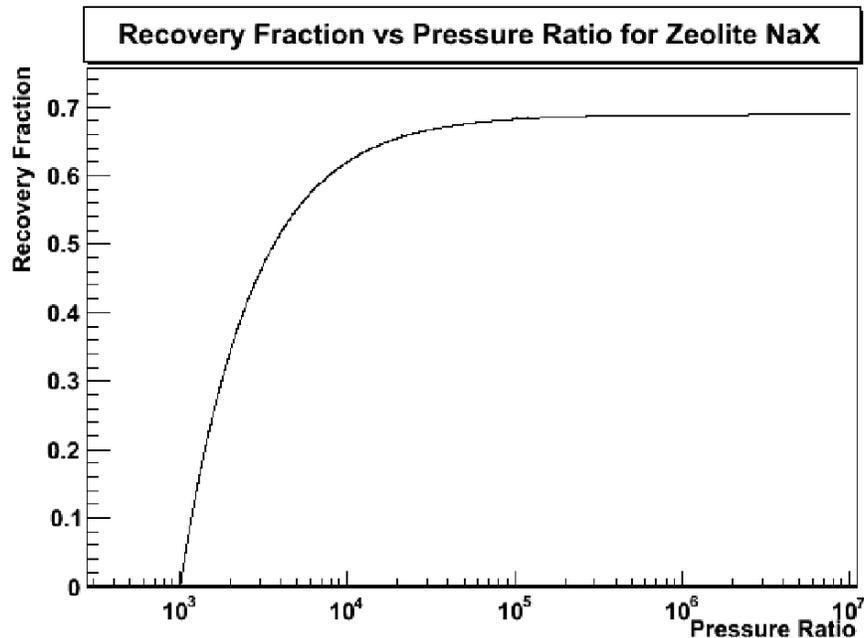


Figure 3.4: Ideal recovery of component B for $y=0.999$ and $\beta = 0.3$. Note that this can be only a very rough approximation at best, since the assumption that the adsorption fraction is linear with pressure does not generally hold when the pressure varies over a very wide range.

(B). Figure 3.4 is a plot of equation (3.11) for this mixture of gases on Zeolite 5A ($y_F = 0.999$ and $\beta \sim 0.3$ [8, 96]) . It should be emphasized here that equation (3.11) is the recovery for a process that extracts pure component B; improvements in recovery can be achieved by reducing the extent of purge, in effect sacrificing purity for added recovery [94]. However, these additional gains become minimal when the concentration of light component becomes very small.

3.3 Prototype PSA Development

Figure 3.5 is a schematic depiction of the PSA system constructed at Princeton. As Figure 3.4 shows, achieving more than minimal recovery of the argon, which is present at $< 1\%$ of the gas stream, requires operating the system with pressure ratios in excess of 10^3 . Hence it was necessary to install a vacuum oil pump on the waste line, which allows us to bring the

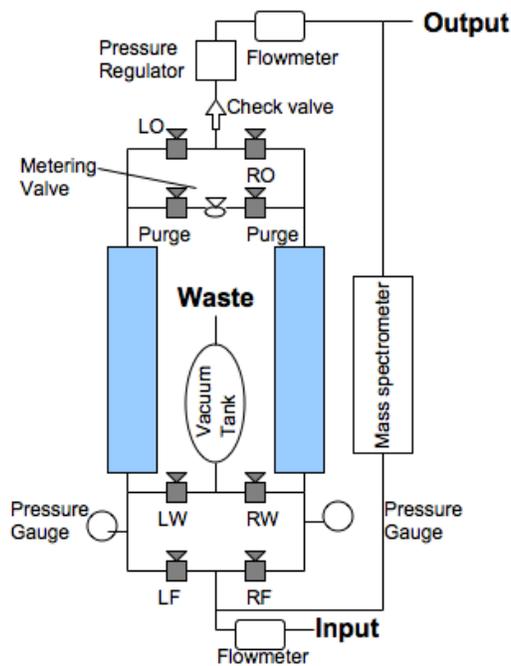


Figure 3.5: Schematic depiction of the prototype PSA apparatus.

low pressure down to ~ 2 torr. The high pressure was usually set between 5 and 25 psig, which results in a pressure ratio of ~ 500 . The cycle time is determined primarily by the time it takes the vacuum pump to fully depressurize a column to a few torr, which can take several minutes. A large vacuum “buffer” tank is installed on the waste line before the pump both in order to protect the pump from being exposed to high pressures and to decrease the blowdown time.

The rate of the feed step is controlled by throttling the output flowmeter, which, during operation, is usually set to around 40 ccm (cubic centimeter of gas per minute at standard temperature and pressure). The check valve on the output line shown in Figure 3.5 prevents the low pressure column from pulling gas backward along the output line during the beginning of the pressurization state. This decreases the number of cycles necessary to achieve steady state operation. The metering valve restricts flow during the purge step to minimize turbulence and ensure that the minimum necessary amount of purge gas is used.

The valves are solenoid valves, which are closed until energized by standard 120 V-60 Hz current. The two valves labeled “Purge” function as a single valve and are connected to the same switch; it was necessary to use two facing opposite directions in order to hold off vacuum at either end of the line, as these particular valves are only vacuum rated in one direction. Switching of the valves is handled by a National Instruments cRIO-9101 real-time configurable chassis with two cRIO-9481 4-channel switching relay modules. The cRIO chassis is controlled via a LabVIEW program. The program, a screenshot of which is shown in Figure 3.6, allows users both to control individual valves manually and to fully customize step times and valve positions of a PSA cycle with an optional separate initialization routine. For this setup, the steps are controlled entirely by timing, with no feedback from pressure or flow rate, for instance. Step times can be estimated knowing the adsorption constants of the system using the equations presented in the previous section, but precise hand-tuning is necessary to achieve optimum operation. The control system was designed with this in mind, to allow the operator to easily change step timings and valve positions even during a running cycle. From a control point of view, a step is defined by the setting (open or closed) of each of the valves and a time duration to hold that state.

The mass spectrometer is a Quadstar 427 residual gas analyzer from Pfeiffer Vacuum. The valves are configured so that the spectrometer can measure the composition of gases either at the feed or product end easily. With this unit we are able to measure breakthrough curves, or, with the data from the input and output flowmeters, determine the recovery ratio of the argon. There are two sets of adsorption columns, one with columns 10' by 1" ID and one set 4' by 3" ID. The columns are stainless steel, capped with Viton gaskets to keep everything vacuum tight. The rest of the plumbing is copper tubing with Swagelok or NPT fittings. We evaluated two different types of zeolites, a 13X brand and Molsiv Adsorbents OXYSIV-5 5A type zeolites. During the course of testing the apparatus, we switched back and forth between the long and short columns, as will be described below. The entire apparatus is shown in Figure 3.7.

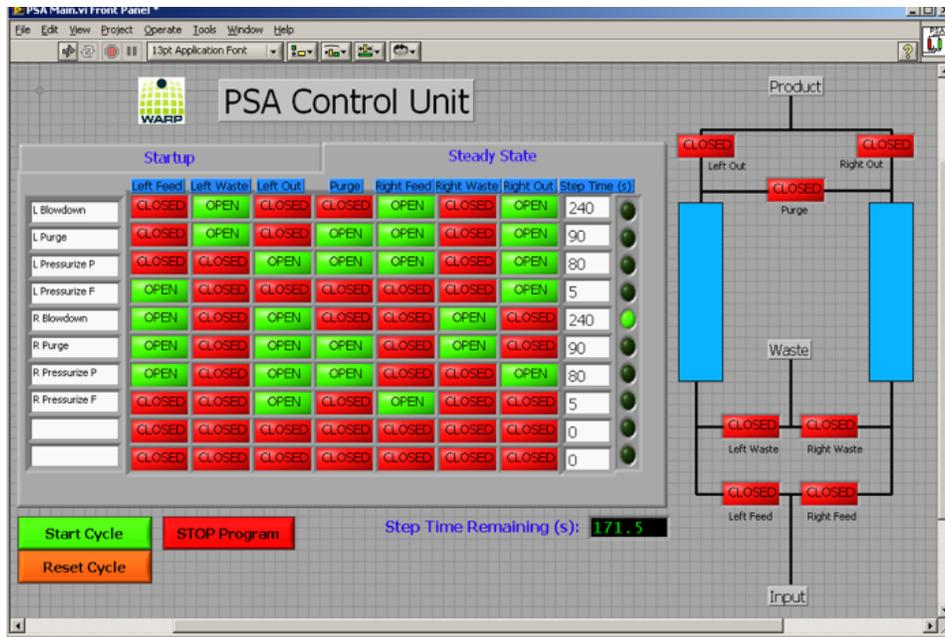


Figure 3.6: Screenshot of the LabVIEW program that controls the valves and timing of the prototype PSA system.

3.4 Prototype Testing

3.5 Oxygen Purification

As mentioned previously, portable oxygen generators are one of the main implementations of PSA technology. It turns out that oxygen and argon have very similar adsorption characteristics on most zeolites (see, for instance Figure 3.1. So the process of concentrating oxygen (and argon) out of air is almost identical to concentrating argon from a mostly nitrogen mixture, except that the oxygen in air is at a much higher concentration than the levels of argon we expected in the underground gas.

Initially, we purchased a commercial PSA unit, an AS-12 Oxygen Generator from AirSep [97], with the intention of modifying the cycle times to suit the needs of our process. Unfortunately, we discovered that the solenoid valves used in the AS-12 required voltages that our cRIO relay system could not handle, so we instead built our own unit from scratch. This had the added advantage of allowing us to include the large columns described in

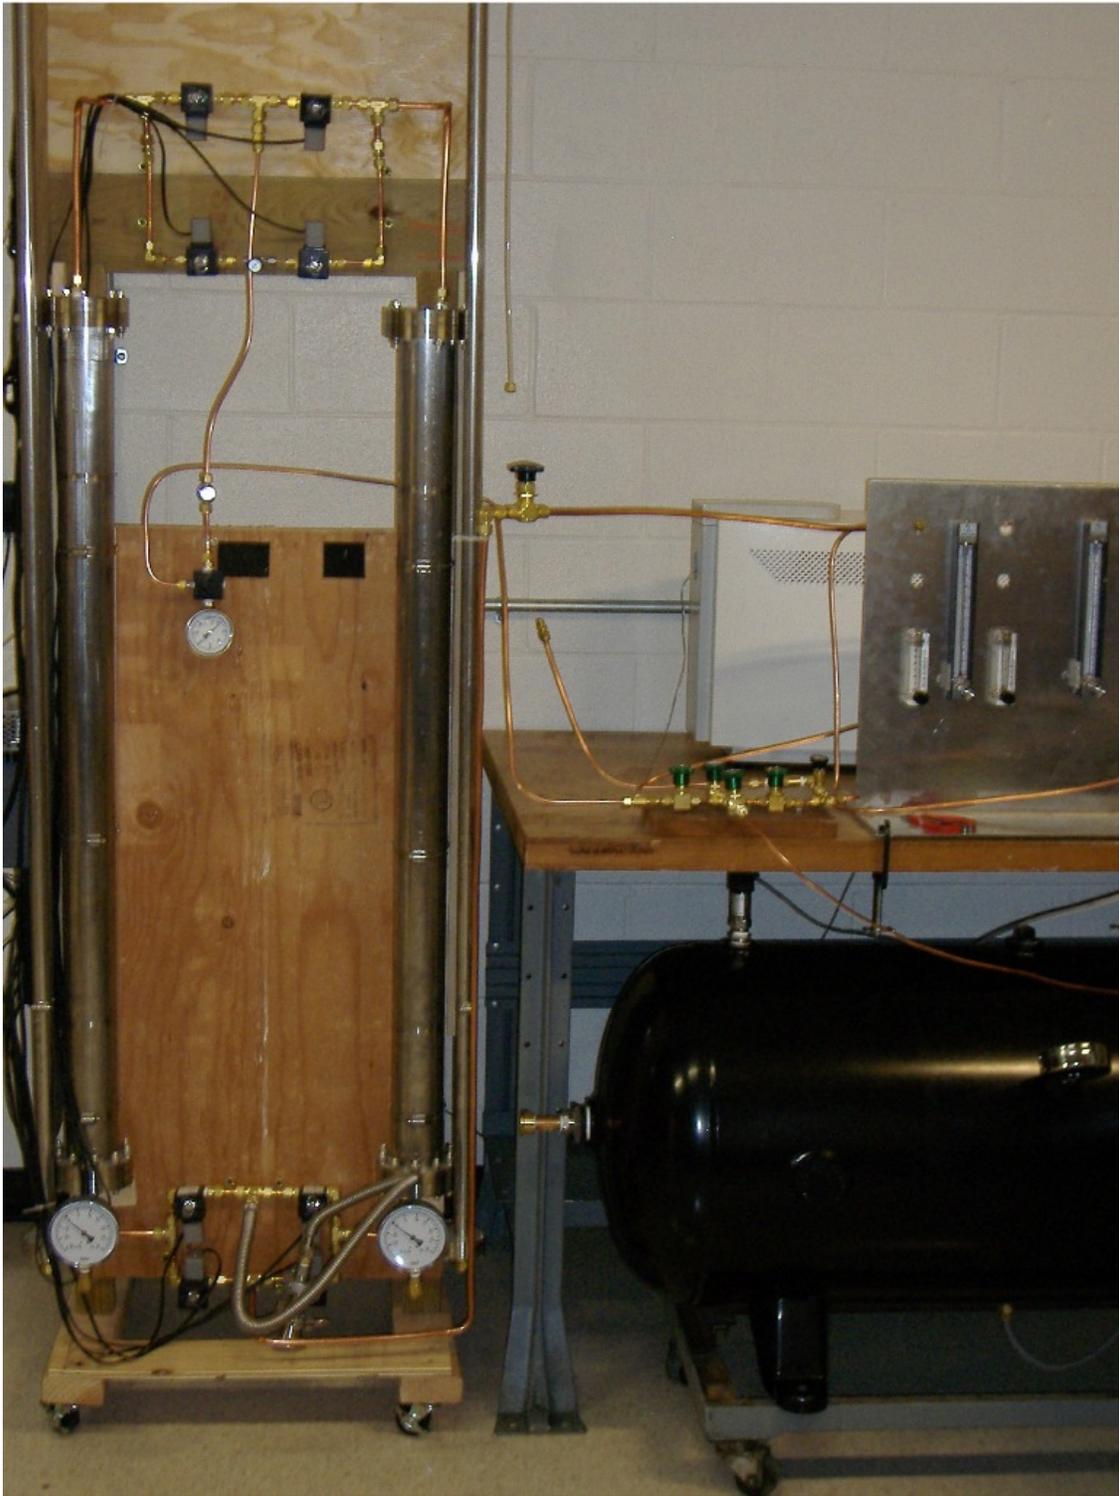


Figure 3.7: The prototype PSA system, shown here partially disconnected after maintenance.

Step Name	Open Valves	Step Time (s)
R Blowdown	LF LO RW	8
R Purge	LF LO PU RW	3
Equalize	LO PU RO	3
L Blowdown	LW RF RO	8
L Purge	LW PU RF RO	3
Equalize	LO PU RO	3

Table 3.1: Optimum PSA cycle for oxygen purification using Drierite bottles. The valve names correspond to those in Figures 3.5 and 3.6, with both valves labeled "purge" open during the PU step.

Section 3.3, which were necessary to maintain constant output with the long evacuation times. We used oxygen purification from air fed through a compressor as a first test of our new PSA system before moving on to argon separation from a primarily nitrogen stream.

While the first set of columns was being constructed, we used standard Drierite bottles, with the desiccants replaced by zeolite 5A. After identifying and fixing several small leaks in the system, we were able to define an optimum cycle given in Table 3.1, which produced ~ 5 scfh of approximately 95% oxygen and 5% argon (< 0 ppb nitrogen detected by the mass spectrometer). The Equalize step was necessary in order to bring the high pressure bottle down to a low enough pressure to expose the vacuum to (this was before we installed the large vacuum tank).

After this initial success, we moved on to attempting to use the shorter (4' long) of the two sets of stainless steel columns, also filled with zeolite 5A, and installed the large vacuum tank to facilitate the blowdown steps. We expected the process to be easier with the larger columns, since the breakthrough time (Equation (3.7)) scales linearly with the column volume, allowing more flexibility in setting the individual step times. Our expected breakthrough time, calculated from Equation (3.7) for a column of ~ 0.2 cubic feet, a value of β for nitrogen of about 10 [96], and typical input flow rates of about 20 scfh, was around 6 minutes, but we consistently saw breakthroughs while running in cyclic operation for much shorter times. One possible explanation for this is that the width of the heavy component front is dependent approximately on the ratio of the column dead volume (the space in the

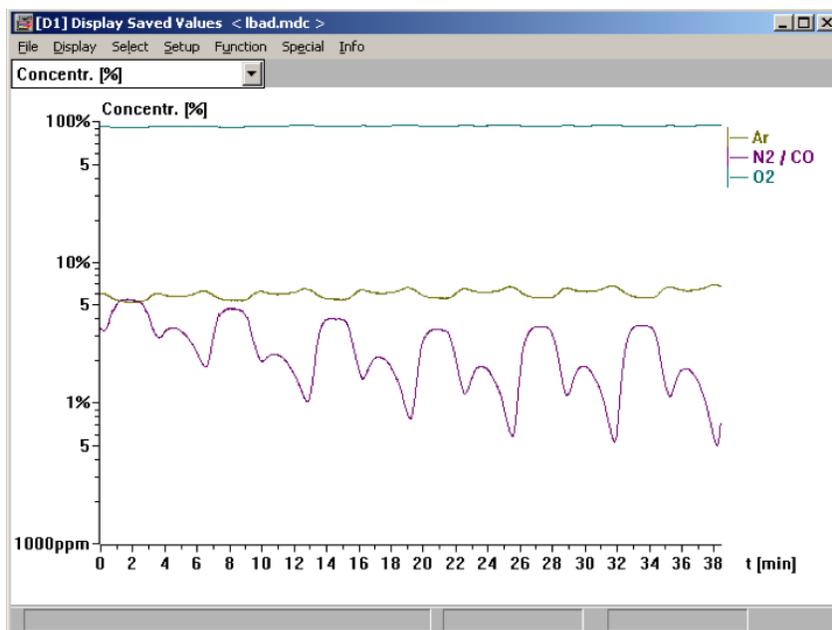


Figure 3.8: PSA output gas composition during oxygen separation showing a clear asymmetry between the two halves of each cycle.

column not filled with zeolite) and the columns length. Heavy component gas entering a column filled with pure light component mixes in the dead volume, widening the heavy component front, but the front sharpens as it propagates down the column [9]. Another possible explanation is that we were at pressures high enough that the linear isotherm approximation was no longer applicable.

Another problem we had with the short columns is shown in Figure 3.8. Despite all attempts to make the two halves of the PSA apparatus identical, a clear asymmetry can be seen in the height of the nitrogen bump during each half-cycle. We were unable to conclusively determine if this effect was due to leaks in the system, a valve that was acting asymmetrically, or settling or crushing of the zeolites in one column.

For these reasons, we switched over to the long ($10'$) columns, at first filled with zeolite 13X. Because of the high impedance of these columns, the blowdown and pressurization steps necessitated very long (15 to 20 minutes) cycles. Initialization from a dead start usually took several hours to reach steady state conditions. Also, the columns had to be opened up,

Step Name	Open Valves	Step Time (s)
R Blowdown	LF LO RW	300
R Purge	LF LO PU RW	120
R Pressurization	LF LO PU RO	300
L Blowdown	LW PU RF RO	120
L Purge	LW PU RF RO	120
L Pressurization	LO PU RF RO	180

Table 3.2: Optimum PSA cycle for oxygen purification using 10' by 1" ID columns. The valve names correspond to those in Figures 3.5 and 3.6.

packed down, and refilled several times in the first few days because the zeolites would settle and not completely fill the columns. Despite these difficulties, however, we were eventually able to determine a cycle that, with the high pressure set at 12 psig and the output flow rate limited to 0.17 scfh, resulted in 0 ppb of nitrogen at the output, but with only $\sim 27\%$ recovery of oxygen and argon. Once a stable cycle was found, however, it was a relatively simple (if slow) process to tweak the recovery by slowly increasing the output flow rate and feed time, and decreasing the purge time and flow rate, loosely following the prescription of Kayser and Knaebel [9]. Eventually we ended up with an output flow of 0.21 scfh, with the cycle settings as given in Table 3.2, which yielded a recovery of 43%. Zeolite 13X does not have quite as good a separation factor as zeolite 5A; β is approximately 0.54, which results in a theoretical recovery of about 45%, which matches well with our results.

3.6 Argon Extraction From Crude Helium

After getting a feel for the system with the tests of oxygen purification, a few attempts were made to separate $<1\%$ argon out of a mixture containing mostly nitrogen. These efforts met with little success, largely due to an inability to control precisely the argon concentration in the cocktail at such low concentrations. This, however, turned out to be irrelevant. We discovered that the helium plant from which we planned to process our first batch of waste gas mixed air into the underground gas early in the processing cycle in order catalytically burn out the methane. Since this also introduces atmospheric argon into the mix, we could

Component	Concentration
He	77%
N ₂	20%
CH ₄	2,3%
H ₂	0.3%
O ₂	0.1%
H ₂ O	700 ppm
Ar	680 ppm
CO ₂	110 ppm

Table 3.3: Composition of crude helium processed for argon by the PSA-charcoal system. The water levels may be a background in the RGA due to adsorption on the capillary.

only use gas removed from the pipeline before the mixing step. The samples we eventually took were of crude helium from the National Helium Reserve, located in the Cliffside Storage Facility outside Amarillo, TX, and had the composition given in Table 3.3.

Although the PSA system had not been designed with this gas composition in mind at all, we devised a method to extract most the argon from the crude helium with only a small addition to the system. Figure 3.9 is a pseudo breakthrough plot for the crude helium. I say “pseudo” because the system is switched to the other column almost as soon as the nitrogen breaks through. As the plot shows, the helium goes right through the column, and the hydrogen breaks through quickly. Around 3.5 minutes later, the argon breaks through, and the nitrogen does not break through for 6 or 7 minutes. (There is some evidence that the nearly constant water line is due to water adsorbed on the walls of the capillary at the input of the mass spectrometer and does not represent actual water in the output.)

Rather than try to fathom a standard PSA cycle for such a complicated gas mixture, we decided to exploit the separation in breakthrough times between the argon and nitrogen. The cycle times were set to switch columns just before the nitrogen breakthrough, allowing only helium, hydrogen, and argon to pass the PSA system. At the output of the PSA, a trap consisting of liquid nitrogen-cooled activated charcoal was installed. The charcoal preferentially adsorbs argon, while allowing virtually all of the helium and most of the hydrogen to pass [98, 99]. Once enough argon has been collected that the charcoal is nearly

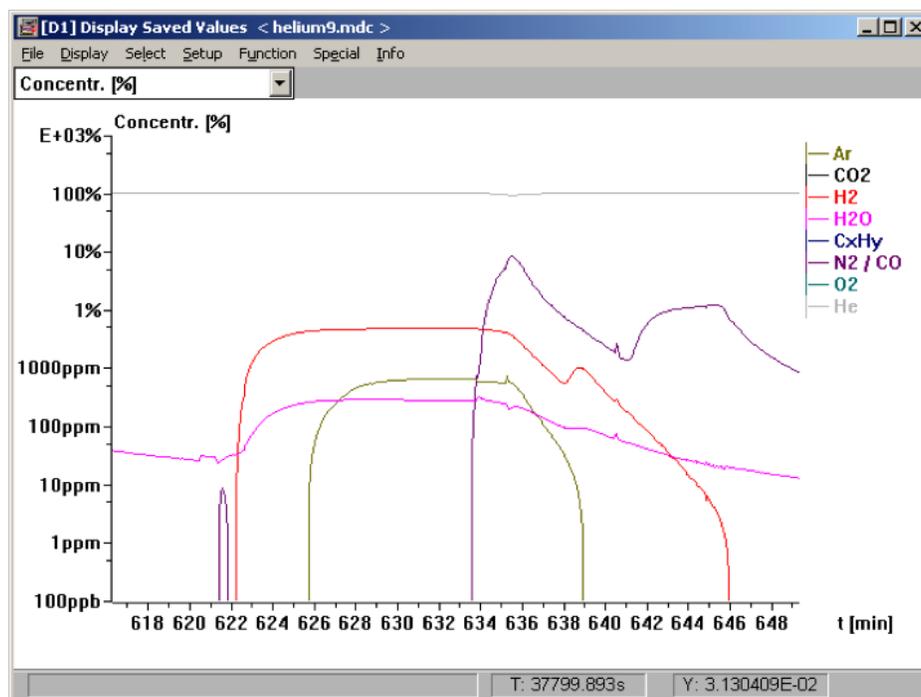


Figure 3.9: Pseudo-breakthrough curves for the crude helium. The breakthrough is stopped and switched to the opposite column just after the nitrogen breakthrough point. The plot shows a clear separation between the argon and nitrogen breakthrough times.

saturated, it is heated, allowing the outgassing argon to be collected. This method was used to capture a sample of argon with upwards of 80% purity (exact determination of the concentrations were impossible at the time due to water contamination in the mass spectrometer) and approximately 60% recovery fraction.

3.7 Measurements of ^{39}Ar fraction

In order to assess the ^{39}Ar fraction in the well gas, the processed samples were sent to the Low Level Counting Underground Laboratory at the Physics Institute at the University of Bern [1]. The laboratory measures ^{39}Ar fractions relative to atmospheric levels by counting ^{39}Ar decays using gas proportional counters, depicted in Figure 3.10. The proportional counter is 100 cm³, constructed of oxygen-free high-conductivity (OFHC) copper, and is

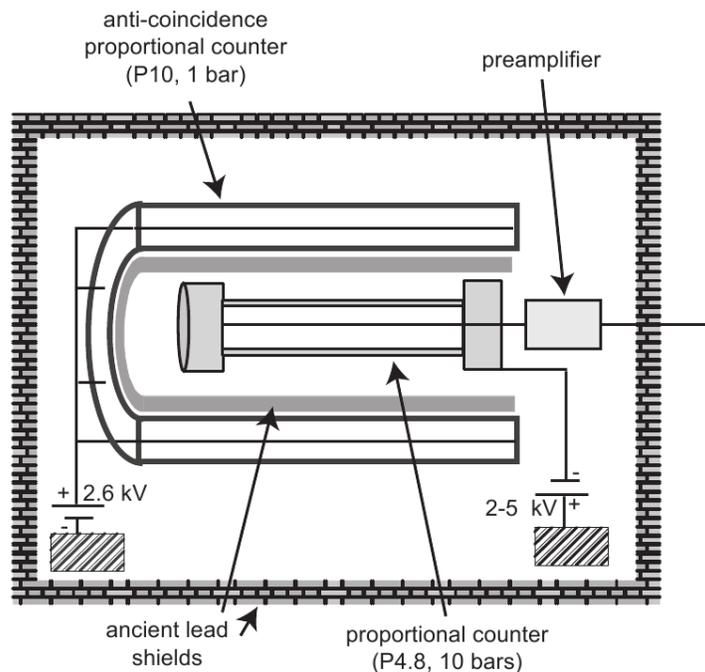


Figure 3.10: Schematic of gas proportional counters used at the Low Level Counting Underground Laboratory [1].

filled with argon at ~ 10 bar. An electric field of $\sim 2-5$ kV causes breakdown in the argon from ionization due to the beta-decay of ^{39}Ar , and the current is read out and digitized by a multi channel analyzer (MCA).

The counter is constructed entirely of materials screened to have low backgrounds, and encased in a 5 cm thick cylindrical low-radioactivity lead shield for passive shielding. A second, cylindrical proportional counter around the inner lead shield serves as an anti-coincidence to reject external radiation, and the entire assembly is enclosed in a second low-radioactivity lead shield 12 cm thick. The laboratory itself is located at a depth of 35 m (70 m water equivalent) to reduce the cosmic muon flux, and is constructed with concrete selected for low radioactivity.

Even with the care taken to reduce the radioactivity of the system, the total background is comparable to the rate of ^{39}Ar decays in atmospheric argon, which limits the sensitivity of the measurement. The sensitivity is further limited by the fact that only a small fraction of the beta decay energy is released in the gas before it reaches the wall of the proportional

chamber, and the linear energy range of the MCA only extends to 35 keV; all events above 35 keV are recorded in the last MCA bin. Since the endpoint of ^{39}Ar is 565 keV, a large fraction of events all end up in the last bin, severely limiting any analysis based on the ^{39}Ar spectrum. Finally, only a few grams of gas can be counted in a single run.

Details of the measurement process can be found in Reference [1]. The ^{39}Ar rate in the first samples from the National Helium Reserve was measured to be $\leq 5\%$ of the atmospheric rate, or depleted by a factor of at least 20 (84% CL). Subsequent measurements of other samples from the same site as well as the Bravo Dome and Doe Canyon CO_2 wells (in New Mexico and Colorado, respectively) all gave similar results, with upper limits only. The best result obtained for any sample was an ^{39}Ar rate of $\leq 4\%$ of atmospheric, or a depletion factor of at least 25. This depletion factor was deemed sufficient to begin full-scale development, and is the factor we will cite when estimating ^{39}Ar rates in DarkSide detectors.

In order to obtain a more accurate measurement of the ^{39}Ar rate, DarkSide collaborators at Princeton have constructed a low-background liquid argon counter, depicted in Figure 3.11. Scintillation signals from ~ 1 kg of liquid argon contained in a reflector-lined “cup” are viewed by a Hamamatsu R11065 3” low-background, high quantum efficiency cryogenic PMT. The argon volume is surrounded on all sides by 2” of OFHC copper and 8” of lead shielding. Plastic scintillator paddles outside the lead shielding serve as a cosmic ray muon veto, and a surrounding layer of water-filled boxes 1’ thick provides further shielding (although in practice these were found to be unnecessary). The PMT signal is digitized by a CAEN V1720 500 MS/s, 12-bit, 2Vpp waveform digitizer, which allows the full ^{39}Ar spectrum to be visible. Neutrons, which are a major contributor to the remaining background after the shielding is in place, are rejected by pulse shape discrimination.

Preliminary measurements of the ^{39}Ar rate by the low background counter give a best fit (not an upper limit) for a depletion factor of 50 with respect to atmospheric argon, but the systematics of the measurement are not well understood. For this reason and to be more conservative, we continue to use the factor of 25 depletion obtained by the gas

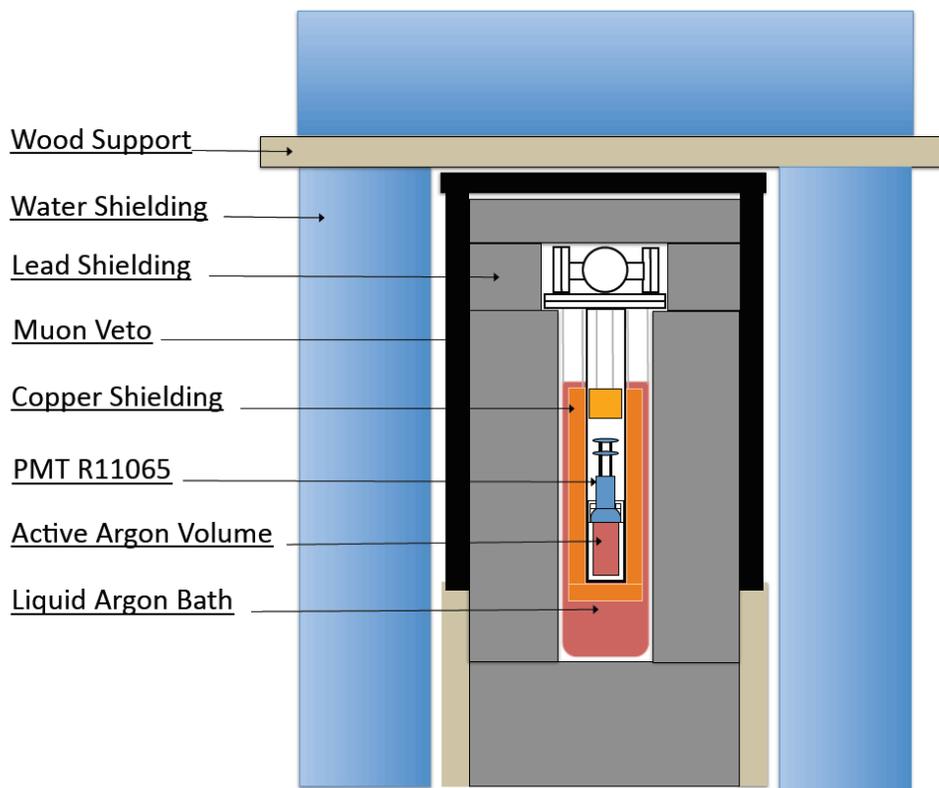


Figure 3.11: Schematic depiction of the Princeton low-background liquid argon counter.

proportional counters. Further measurements using the low background counter will take place in the Kimbalton Underground Research Facility near Blacksburg, VA over the summer of 2011. KURF is located at a depth of 1500 m.w.e., and should bring the sensitivity of the measurement to 0.5% of atmospheric argon.

3.8 Full-scale PSA and cryogenic distillation

Following the initial survey with the prototype PSA system, the collaboration selected the Doe Canyon site near Cortez, Colorado for full-scale operations. Processing of the ~ 400 ppm of argon from the well to 99.9999999% pure detector-grade argon requires three steps.

The first step is a two-stage VPSA system located at the Doe Canyon site. The basic operation of each stage is the same as for the prototype. Figure 3.12 shows the piping and

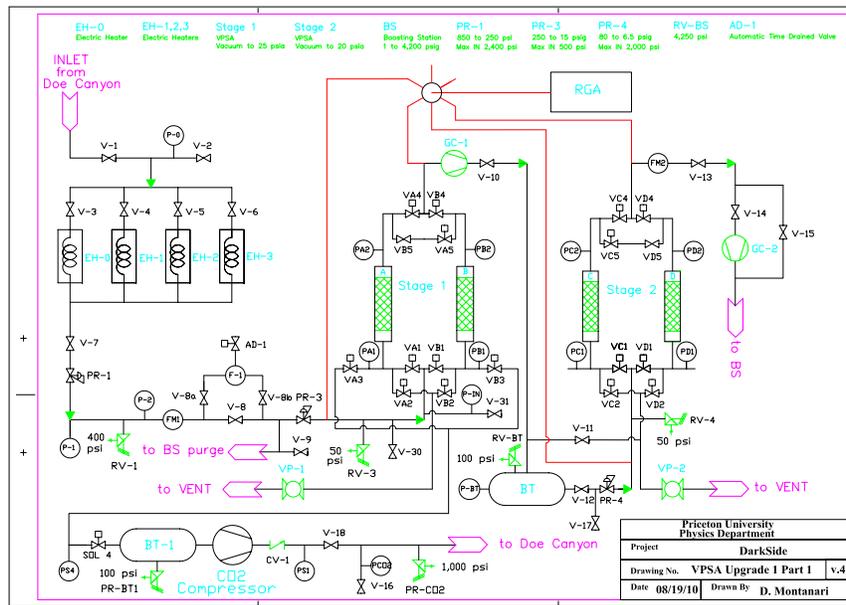


Figure 3.12: Process and instrumentation diagram of the two-stage industrial VPSA plant operating at the Doe Canyon site.

instrumentation diagram (P&ID) of the plant. The first stage columns are 14" in diameter and 56" high, containing zeolite NaX (13X). The first stage of the system removes the CO₂ and trace hydrocarbons present in the stream, producing gas that is primarily nitrogen, with some helium and argon. The pressure drop in the first stage is large enough to liquify the CO₂ in the input stream. To prevent this, the gas is heated before going into the first stage (EH-0–EH-3 in Figure 3.12). The second stage columns are 3" in diameter and 40" high and uses Li-LSX zeolite (lithium and silver-exchanged 13X zeolite). The second stage removes most of the N₂, producing as the final product a stream with ~72% N₂, 24% He, 2-4% Ar, and trace amounts of methane and hydrogen. A multi-inlet mass spectrometer allows for sampling of the gas composition at any of several key points throughout processing. The output gas pressure is boosted to ~4000 psi for storage and shipping in standard gas cylinders. The plant is currently operating, producing around 0.4 kg/day of argon. As of May 2011, approximately 25 kg of depleted argon have been produced. An upgrade planned for the near future should bring the production capability to ~10 kg/day.

The gas produced by the VPSA system is shipped to Fermi National Laboratory for second stage purification in a cryogenic distillation column. Cryogenic distillation is a standard industrial technique for gas separation that exploits the different partial pressures of various gases at cryogenic temperatures to vary the concentration of feed gases over a series of vertically stacked “stages.” The system at FNAL was designed specifically for the mixture of gases output by the VPSA step. Figure 3.13 shows the P&ID for the plant. The column is 130” tall and 1” in diameter, packed with Sulzer EX Laboratory Packing, a common industrial mesh that provides the “stages” for the distillation. The inner column is contained in a 12” OD vacuum-walled cryostat; the cryostat interior is also evacuated, forming a second insulation layer. Cooling power is provided by two 600 W cryocoolers. At the time of this writing, the distillation column is undergoing final testing and commissioning. According to the design specifications, the column will be able to produce ~ 10 kg of argon per day with < 1 ppm of impurities at 95% collection efficiency.

Finally, remaining impurities are reduced to the required ppb or better levels by hot chemical getters, e.g., calcium or zirconium traps.

Chapter 4

10 kg Prototype

In Chapter 2, we described the basic layout of the DarkSide detectors: a dual-phase depleted argon TPC inside a borated liquid scintillator neutron veto all inside a large water tank which serves both as passive shielding and a cosmic ray muon veto. The shielding and vetoes reduce most external backgrounds, and remaining beta/gamma backgrounds are removed by pulse shape and scintillation to ionization ratio (S2/S1) discrimination. A number of technical requirements must be met for the successful operation of an argon TPC:

- High chemical purity of argon. As mentioned previously, chemical impurities quench both the scintillation light and free electrons, and negatively impact both pulse shape and S2/S1 discrimination.
- High light yield (probability of detecting scintillation photons). Sensitivity to WIMP-induced nuclear recoils decreases exponentially with increased threshold; the threshold in turn depends on the light yield.
- Electric field uniformity. Non-uniform drift fields can negatively impact S2/S1 discrimination and position reconstruction.
- Environmental stability. Changes in the temperature, pressure, or level of the gas pocket can all negatively impact S2/S1 discrimination.

In addition to the above requirements, a TPC sensitive to current WIMP interaction limits must also be constructed in a clean room environment out of low-background materials. The first physics-capable DarkSide detector will be DarkSide-50, a 50 kg active mass TPC which will be installed in the Borexino CTF tank in the LNGS cavern in 2012.

DarkSide-10 is a small-scale (10 kg of active argon) prototype TPC constructed and operated at Princeton in order to demonstrate the capability to build and run a full-scale background-free detector. In addition to demonstrating operational capability, DarkSide-10 was also used to test several technologies that will be implemented in future physics-capable detectors:

- High Voltage Feedthroughs. Argon gas has a relatively low breakdown potential, so most commercial HV feedthroughs will spark in argon. All the feedthroughs in DS-10 were either modified or completely custom made.
- TPB evaporation. Efficient light collection requires that the first surface encountered by the VUV argon scintillation photons be a fluorescent wavelength shifter, otherwise a substantial amount of light is lost. Evaporation of the TPB required the construction of a large vacuum chamber and apparatus to perform and monitor the evaporation process.
- Transparent conducting windows. In other TPC designs to date, the anode and cathode for the electric fields were made from wire grids. Since the electric field cannot be completely closed by grids, multiple grid layers are necessary to completely close off the field, resulting in “dead” regions with partial field. Transparent conductors such as indium-tin-oxide (ITO) can provide a solid conducting plane to close off the electric field while allowing light to pass to the PMTs outside the field regions. ITO is fairly recently developed, and, prior to DarkSide-10, had seen little or no testing in cryogenic environments.
- Hamamatsu R11065 PMTs. These high quantum efficiency, low background, cryogenic capable PMTs, which will be used in the first iteration of DarkSide-50, were new to

the market at the time of DarkSide-10's construction; DarkSide-10 provided the first operational test of these PMTs.

With respect to DarkSide-50, DarkSide-10 is simplified in a number of ways:

- Smaller active volume, requiring less argon and allowing faster filling and emptying of the detector.
- Commercial (i.e., not depleted) argon. This allowed DarkSide-10 to be constructed before a supply of depleted argon was available, and also greatly simplified the gas handling system, since there is no need to recover commercial argon on detector warming, unlike depleted argon.
- Smaller scale. This is most important because the drift field requires a smaller maximum voltage, putting lower requirements on the high voltage feedthroughs.
- No attempt to use low-background materials or methods for construction.
- No veto or shielding.
- Fewer PMT channels.

In addition to those listed above, another major difference between DarkSide-10 and DarkSide-50 is that the inner volume container for DarkSide-10 is constructed out of acrylic, whereas in DarkSide-50 it will be composed of quartz and teflon. Since acrylic is brittle at liquid argon temperatures, this is one way in which DarkSide-10 is more complicated than DarkSide-50. Constructing the inner vessel out of acrylic was an early plan to incorporate passive shielding between the PMTs and active volume, before the decision was made to use an active neutron veto.

Two data-taking campaigns have been conducted so far with this detector. The first, from March-July 2010, was single-phase, i.e., without any electric fields. After the first campaign, all but two of the PMTs were replaced, field-generating pieces and HV feedthroughs were installed, and a number of upgrades were made to the DAQ and slow control systems. The second campaign was conducted from November 2010 through February 2011, and analysis of the data taken during this campaign will be the subject of the remainder of this chapter.

A third campaign will take place underground in the LNGS cavern beginning in the summer of 2011.

4.1 Detector description

Figure 4.1 is a schematic showing the key features of the 10 kg inner detector, and Figure 4.2 is a photograph of the fully-assembled inner detector, suspended from the top flange as it is being inserted into the cryostat. The active volume is contained within an acrylic cylinder 0.75 inches thick with a 9.5 inches nominal ID (11 inches OD) and 9.25 inches high. The cylinder is sealed at both ends by acrylic disks 13 inches in diameter and 0.5 inches thick. All of the acrylic pieces are annealed after machining but before sealing by baking at 85 °C for roughly 36 hours, with a slow transition from room temperature and back (raising the temperature at 15 °C/hour and lowering at 5 °C/hour). The annealing procedure reduced the OD of the acrylic cylinder to approximately 10.94 inches. The disks are sealed onto the cylinder using Teflon o-rings. The seal is designed to be “bubble-tight”, i.e., to allow for the presence of a pocket of argon gas at the top of the acrylic vessel maintained by boiling some of the liquid. The pieces are held fixed and compressed by a “cage” system consisting of a stainless steel compression plate on each end 14.5 inches in diameter and 0.5 inches thick, joined by stainless steel rods. The rods are tightened onto springs to maintain compression after the system is cooled to liquid argon temperature (88 K). (See figure 4.3).

The compression plates are machined with holes for the photomultipliers: seven 3” PMTs on top and a single 8” PMT on bottom (see Section 4.2.2). The PMTs are held off from the compression plates by Teflon spacers, and pressed forward by hold-down plates (one piece each for top and bottom) which are attached to the compression plates by screws tightened onto springs. The compression plates also serve to anchor the detector to the cryostat flange via four stainless steel support rods. The entire assembly is contained within a vacuum-walled cryostat, 18 inches ID, and submerged in a bath of liquid argon, up to a few inches

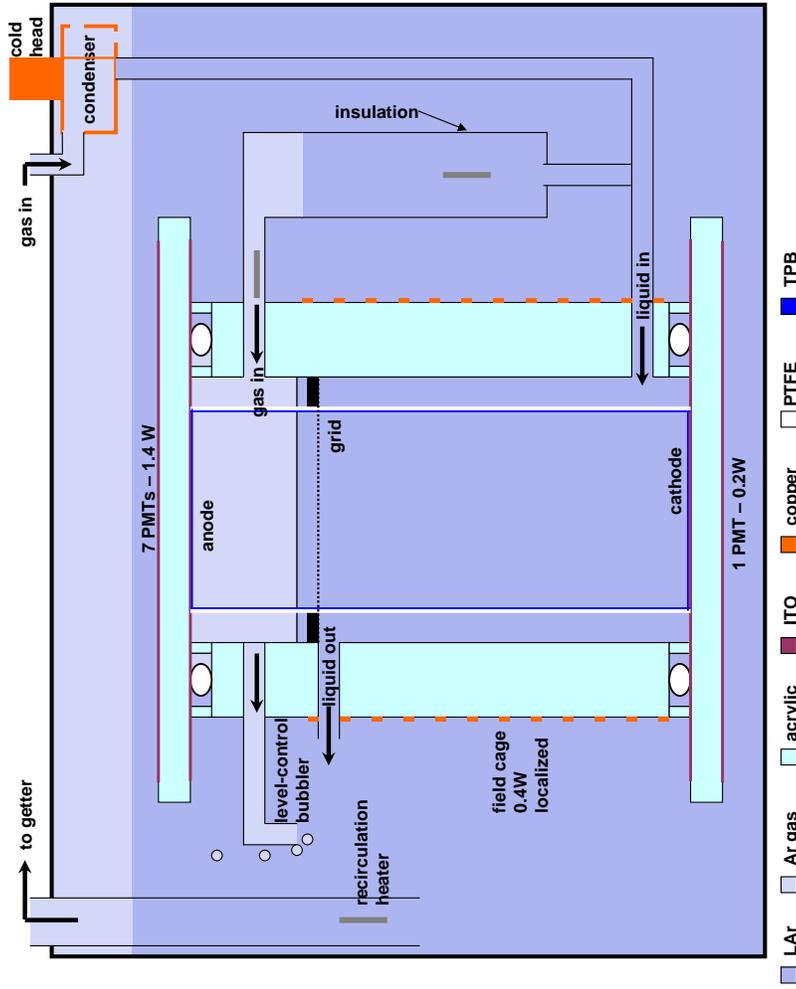


Figure 4.1: Cartoon schematic of the DarkSide 10 kg detector. Teal regions are acrylic; darker blue regions are liquid argon and lighter blue are gaseous argon. Figure is not to scale. See text for details.



Figure 4.2: Photograph of the 10kg prototype inner detector during assembly.

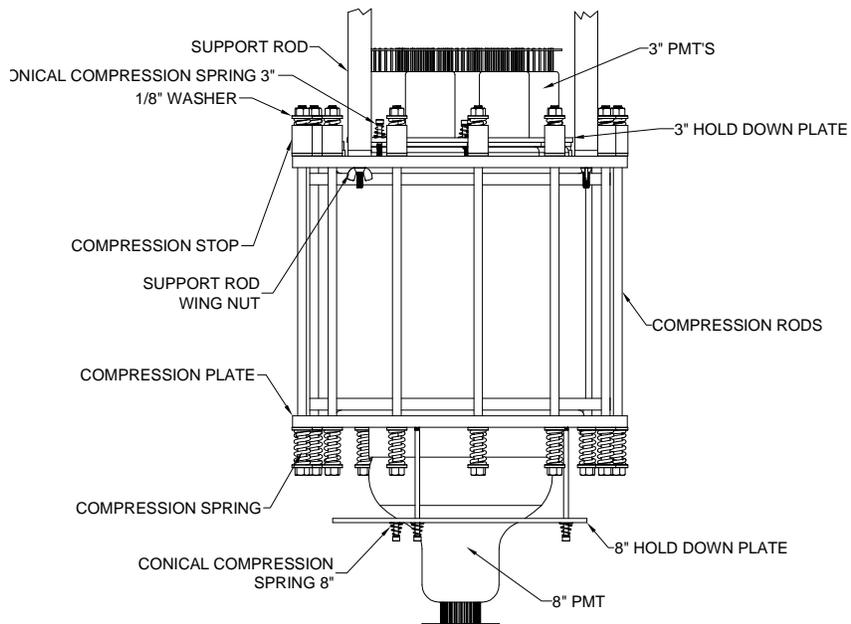


Figure 4.3: Diagram of the DarkSide-10 inner detector assembly, showing the support structure and photomultipliers. See text for details.

above the top PMT bases. This external volume of argon serves as a thermal buffer which minimizes temperature changes in the active volume.

Both sides of the acrylic disks are coated with a ~ 100 nm thick layer of indium-tin-oxide (ITO), a transparent conductor (for a total of four layers of ITO). The outer layers (nearest the PMTs) are held at photocathode voltage for optimal PMT operation. The inner bottom ITO layer serves as cathode for creating an electric field within the detector volume with the top inner layer as anode (ground, in this case). Electrical connections to the ITO layers are made by clamping beryllium-copper springs held in Teflon brackets to the ITO layer which extends beyond the acrylic cylinder. A stainless steel grid with a hexagonal mesh (100 microns thick; each “wire” segment 500 microns wide and 5 mm long) rests on a lip machined into the acrylic cylinder 1 inch from the top (anode). Electrical connection to the grid is achieved by a wire running through a small hole drilled into the acrylic cylinder just below grid height. Three separate regions of electric field strength are maintained in the detector during two-phase operation: a drift field in the liquid volume below the grid, an extraction field in the liquid region above the grid, and an electroluminescence field (sometimes referred to as multiplication field) in the gas region. Typical settings for these fields are ~ 0.7 kV/cm for the drift field and ~ 3 kV/cm for the extraction field; the electroluminescence field is not set independently, but is a factor of ~ 1.5 (the ratio of dielectric constants in argon gas and liquid) times the extraction field. Given the separations between the anode and cathode (~ 23 cm), and anode and grid (~ 2.5 cm), these field settings require a total voltage drop of ~ 25 kV. Uniformity of the fields throughout the inner volume is maintained by a field cage consisting of wide copper strips patterned onto a flexible kapton printed circuit board which is wrapped around the outside of the acrylic cylinder. The strips alternate between front and back of the board and partially overlap, providing almost complete coverage of the system. Groups of resistors are connected in parallel between adjacent strips provide the appropriate voltage drops.

Inside the acrylic cylinder is a Teflon support structure to which are mounted reflecting foils (Vikuiti reflecting film from 3M company). The foils are arranged into a cylinder 8.25 inches in diameter. A separate ring holds reflector above the grid. The inside surface of the foils and the inner ITO layers are all evaporated with $200 \mu\text{g}/\text{cm}^2$ of tetraphenylbutadiene (TPB), a fluorescent compound which converts the 128 nm argon scintillation light into the visible region (peaked around 420 nm [100]), where it can efficiently be reflected by the 3M foil and detected by the PMTs. The inner surface evaporations are arranged so that the first surface encountered by a scintillation photon should always be TPB-coated.

There are a total of four penetrations in the side wall of the acrylic cylinder (in addition to the small hole for the grid wire) to allow for recirculation of the argon inventory. Clean liquid from the recondenser (see section 4.2.1) is delivered directly to the liquid region of the active volume. In order to ensure that clean argon is also delivered to the gas phase, the liquid delivery line is also connected via a 'T' junction to one of two Teflon tubes mounted to the outside of the acrylic cylinder (generally referred to as the boiling arm). The tubes are connected to each other at top and bottom to allow both liquid and gas exchange between them. The second tube is connected to the acrylic cylinder by a single port in the gas region. The purified liquid is boiled in these tubes so that the bubbles do not perturb the liquid surface in the active volume. Inside the second tube are two parallel, vertical copper bars, the capacitance of which is used to measure the level of liquid in the boiling arm (and therefore in the inner vessel as well, assuming that any pressure drops between the two volumes are negligible). An array of five PT-1000 resistive temperature devices are attached to the capacitive level sensor (CLS) to calibrate it. The PT-1000s are used as point level sensors by driving them with a fixed voltage ($\sim 20 \text{ V}$); when in liquid the power dissipated (and hence the measured current) in the resistor increases markedly (by roughly a factor of 2). A high-power resistor was installed originally in order to provide the heat for boiling, but, in practice, we found that the bottom-most PT-1000 was sufficient to maintain the gas pocket.

To maintain the level of the liquid-gas interface (at ~ 2 cm below the top acrylic plate), an arm (referred to as the bubbling arm) consisting of a 0.25 inch ID acrylic tube is attached to the cylinder opposite the boiling arm, forcing all the boiled gas to exit at a fixed height. In order to reduce variations in the level due to the formation and release of macroscopic bubbles, the aperture of the arm is actually a flattened pipe edge with many small holes drilled in it, rather than an elbow as shown in Figure 4.1. The final penetration of the acrylic cylinder is a simple hole for liquid outlet, situated below gas outlet (near the top of the liquid volume) to maximize mixing of the liquid.

4.2 Data taking conditions

4.2.1 Recirculation

To maintain the purity of argon in the detector, liquid argon in the buffer volume surrounding the inner detector is constantly boiled off. The boil-off is pumped to a commercial SAES hot zirconium getter (part number PS4MT3R1). A mass flow controller maintains a flow rate of 15 std. liters of argon gas per minute, which is the maximum rate the getter can purify to its rated cleanliness. The purified gas, which comes out of the getter near room temperature, is directed to a cold head which consists of two copper cans cooled by a Cryomech PT-90-UL cryocooler. The cryocooler delivers a constant ~ 90 W of cooling to a cold finger directly attached to the condensing cans; the temperature is maintained at 88.3 K by a resistive heater controlled by a Lakeshore 340 PID temperature controller.

The two condenser cans are separate; one is connected to the inner volume (bubbling arm) inlet, and one to the outer volume. External valves allow one to direct the purified gas stream to either or both volumes. During the early part of this campaign, it was standard practice to open both valves, allowing the purified gas to flow into both the inner and outer volumes. Because the pressures into the two condensing cans are equal, and the inner volume is at higher pressure than the outer, only liquid should be delivered to the inner volume in

this mode. If we assume that all of the surfaces of the heat exchanger condense argon at the same rate, we can estimate that roughly 1/4 of the purified argon goes to the inner volume. This mode was generally referred to as “gentle recirculation.”

Beginning on December 11, 2010, we instituted new procedures for purifying the argon inventory of the inner detector. During actual data acquisition, or whenever the HHV fields were on, all of the purified gas flow was directed to the outer volume. Overnight, the HHV is turned off, and all of the flow is forced into the inner volume, providing so-called vigorous recirculation. Due to inefficiencies in the heat exchange in the copper condensing cans, when in this mode most of the purified argon is delivered as gas, quite possibly warm, which produces violent bubbling in the side arm and blows warm gas at high rate and possibly elevated pressure through the active volume. This regimen was generally followed for the remainder of the campaign. The switch to the vigorous recirculation mode occurred largely due to not observing a substantial increase in purity (as measured by the electron drift lifetime, see section 4.5.3) from the gentle recirculation mode.

4.2.2 Photomultipliers, DAQ, and trigger

Table 4.1 lists some characteristics of the photomultipliers used in this campaign. The top PMTs are Hamamatsu R11065 high quantum efficiency 3” PMTs with flat photocathode, and the bottom is a model R5912-02 with an 8” hemispherical photocathode. The signal and trigger path is diagrammed in Figure 4.5. The PMT signals are amplified 10x by a CAEN V974 fast amplifier before being digitized by a CAEN V1720 250 MHz, 12-bit digitizer. The digitizer can be controlled directly through a fiber-optic interface to the DAQ computer or via a VME crate controller, with its own fiber-optic interface. The firmware of the digitizers is the standard provided by CAEN, but the software running on the DAQ computer was entirely custom-written. The PMT bases have 50 Ω back-termination to prevent reflections. The HV supply of each PMT is set so that the mean of the integrated single photoelectron response falls around 80 ADC counts \times samples, which corresponds to

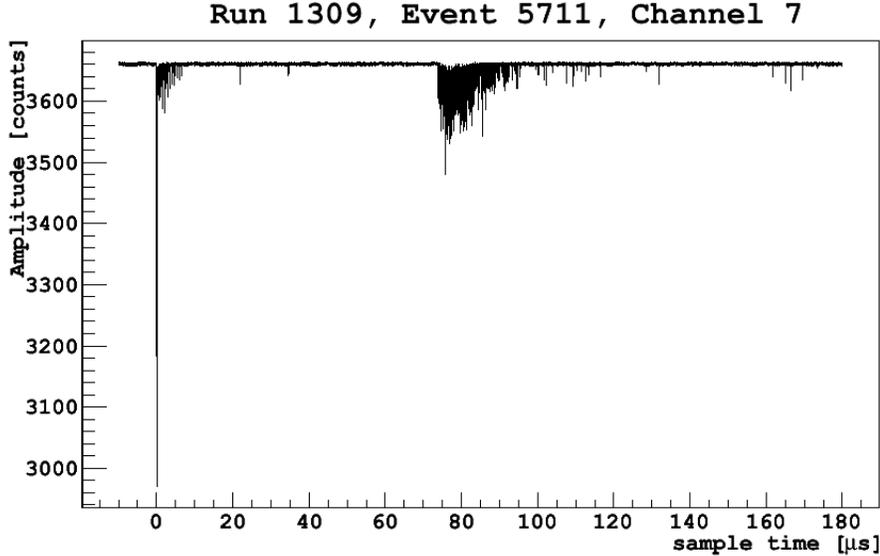


Figure 4.4: Example event, showing the digitized signal from channel 7 for a gamma event that deposited ~ 65 keV in the active volume.

a gain of around 4×10^6 . The 8" PMT has a base with a custom voltage divider which is intended to produce a lower gain (to avoid saturation) without losing collection efficiency.

A copy of the amplified 8" PMT signal is sent to a discriminator with a 60 mV threshold, which forms the primary internal trigger. When the trigger is received, the digitizer records for each channel a set number of pre-trigger and post-trigger samples with a 4 ns sample interval. For the runs considered in this analysis, we used 2500 pre-trigger samples ($10 \mu\text{s}$) 45000 post-trigger samples ($180 \mu\text{s}$). To save disk space, some runs without the drift and extraction fields (i.e., no S2) were acquired with only 5000 post-trigger samples ($20 \mu\text{s}$). Figure 4.4 is the recorded trace from channel 7 for an example event, a gamma that deposited ~ 65 keV in the active volume.

An external coincidence trigger can be generated by discriminating the output of a commercial NaI detector. Beginning at run 1292 on December 15, a $60 \mu\text{s}$ veto was added after every internal trigger after a delay of 20 ns. This was intended to allow the internal trigger to fire only once at the beginning of each S1 and S2. Prior to this change, the internal trigger tended to be live for tens of microseconds at a time during large S2 pulses, which caused a

high rate of accidental coincidences. The final trigger (either just the internal trigger or the coincidence of the NaI plus internal) opens a veto which is set (manually) to approximately the length of the acquisition window. This veto prevents retriggering within the acquisition window; in this way we can estimate the dead-time of the system by counting all of the generated triggers without being thrown off by pileup. After each run, basic statistics for the run such as start and end time, number of triggers, and the estimated live time are automatically saved to a database.

In order to calibrate the response of the PMTs, we use a pulsed diode laser with a 425 nm wavelength (near the emission peak of TPB) with a pulse width of tens of picoseconds. The laser output is focused onto an optical fiber by which it is carried into the detector. The other end of the fiber is mounted at the bottom acrylic plate so that it shines onto the side wall of the inner volume. The evaporated TPB is a fairly good diffuse reflector, so the light is then scattered roughly evenly about the inner volume. The intensity of the light is controlled by a series of seven filters between the laser and fiber mounted so that they can be rotated in and out of the beam. Figure 4.6 shows a photograph of the laser, filters, and fiber setup. The filters themselves are pieces cut from a single neutral density Kodak gel filter so that each additional filter should attenuate the light by the same amount. The filters have a nominal optical density of 1 (i.e., 10x attenuation), although in practice at 425 nm we find the attenuation to be a factor of ~ 14 , which has been confirmed by a spectrophotometer.

PMT calibration data is acquired in dedicated runs, for which the trigger is generated by the laser controller box. The trigger rate for these runs is approximately 4 kHz, which is near the maximum rate at which events of this size can be written to disk without any dead time. The laser events are automatically analyzed to find the single photoelectron response for each PMT (see Section 4.4.1) and, after visual inspection, the analysis results are also written to the database. The single photoelectron fit machinery performs most reliably when the average number of photoelectrons per laser trigger on a given channel is around 0.1. Because the 8" PMT has approximately seven times the surface area of each of

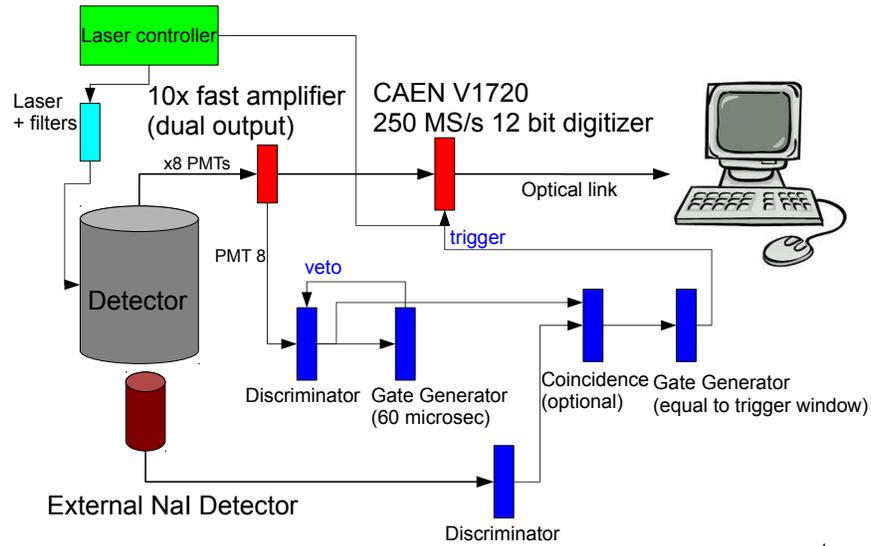
PMT / cable	DAQ channel	Model	Size	PMT S/N	Base S/N	% QE @420 nm	HV Supply
1	0	R11065	3"	ZK4997	5	32.6	1580
2	1	R11065	3"	BA0041	3	35.5	1650
3	2	R11065	3"	ZK4995	1	32.9	1410
4	3	R11065	3"	ZK4976	4	31.1	1520
5	4	R11065	3"	BA0042	2	35.6	1580
6	5	R11065	3"	ZK5174	9	33.4	1400
7	6	R11065	3"	ZK5174	6	31.5	1460*
8	7	R5912-02	8"	FA0023	n/a	17.5	1200

Table 4.1: Information on PMTs used in this data campaign. All are manufactured by Hamamatsu. Channels 0, 3, and 6 (marked in bold) were omitted from most of the analysis presented in this memo due to poor behavior. (*)The HV on channel 6 was raised to 1560 V on Dec. 30 and 1700 V on Jan. 7.

the 3" PMTs, it receives much more light than they do, and so we take separate runs with a different number of filters to give each the best amount of light for the analysis. During analysis, the the calibration information for the valid laser run with the appropriate number of filters closest in time to the run being analyzed is read from the database and used to scale the PMT signals. Channel 7 is calibrated on laser runs with 3 filters, while the others use two-filter laser runs.

Because the detector is unshielded and at the surface, the rate of background events in the detector is extremely high, on the order of 500 Hz. On January 24, 2011, we installed some moderate shielding, consisting of boxes 1 foot on each side filled with water, stacked 6 high around 3 sides of the detector, plus a layer of copper plates (1-2 inches thick in total) stacked below the detector. This shielding brought the background rate down to ~ 300 Hz, which is consistent with expectations based on the solid angle coverage of the shielding. (This shielding was installed after the period considered in this analysis.)

Because of the high interaction rate, the PMTs are exposed to a very high amount of light when the HHV fields are present due to the S2/S1 multiplication. This light level is in fact comparable to the level at which the photocathode is known to be non-linear. (Hamamatsu shows a 5% non-linearity around 10-100 nA of anode current. For gains of $\sim 10^6$, this would be around 10^5 photoelectrons per second, which is easily reachable with trigger rates of a



4

Figure 4.5: Diagram of the data acquisition and trigger setup for the second DarkSide-10 campaign.



Figure 4.6: Photograph of the diode laser, filters, and optical fiber used to calibrate the PMT response.

few hundred Hz and order of 100 p.e. in an S2 pulse.) Two of the channels (channels 0 and 3/PMTs 1 and 4) showed severe degradation in performance when the S2 signals were present. This degradation consisted of a broadening and delay of the current pulse of the tubes, as well as an apparent decrease in the overall efficiency of photon detection. The effects were confirmed both by pulsing with a laser and in scintillation responses [101].

At some point beginning around December 23, PMT 7 (channel 6) began exhibiting a different sort of poor behavior, appearing to be a significant, and gradually increasing, drop in gain. For some time, the problem seemed correctable by raising the supply HV to the PMT, but eventually it stopped producing measurable signals, even with a supply of 1700 V. Later measurements of the base showed a failed resistor, so it is assumed that this was a mechanical failure of the base, quite possibly due to shorting of the PMT can (which is at negative high voltage) to the nearby clamping plate, which is grounded.

In order to avoid systematics due to changing numbers of channels, all three of these problematic channels (0, 3, and 6) were removed from all analyses, including runs where they were presumably performing well, unless specifically mentioned otherwise. This has an obvious impact on the total amount of light collected as well as the observed resolution of scintillation signals, particularly S2, which occurs very close to the top PMTs. Where appropriate, we have performed separate analyses with all channels included in order to estimate the scale of this effect.

4.2.3 Data sample

Again because of the very high rate of background events in the detector, most of the data collected during this campaign was taken using the coincidence trigger. For gamma events, we use a ^{22}Na source, which emits a 1275 keV gamma along with two back-to-back 511 keV gammas from positron annihilation. Because the annihilation gammas are emitted back-to-back, while the 1275 keV gamma is emitted isotropically, the relative weight of triggers containing a 1275 keV gamma in the active volume is highly suppressed by the coincidence

trigger by roughly the solid angle subtended by the active volume. The height of the spectral features due to the 1275 keV gamma is further suppressed due to the fact that, if the 1275 keV gamma reaches the inner volume, the NaI trigger must have been generated by one of the annihilation gammas, and so the majority of these triggers will contain both a 1275 and a 511 keV gamma. This effect has been verified by observing the spectrum of the NaI detector during coincidence triggers, in which, due to the high resolution, three peaks are observed, corresponding to a single 511 keV, a single 1275 keV, and the sum of the two; the single 1275 keV peak was suppressed with respect to both of the other two. In the argon detector, this is not directly observable because much of the 1275 keV region and all of the region above the 1275 keV peak is outside the digitizer range. Cuts that attempt to select only events with a single interaction site (see section 4.4.2) also suppress the events in which both a 1275 and 511 keV gamma are present. To generate argon nuclear recoil events, which mimic WIMP-induced signals, we use an americium-beryllium (AmBe) source. The AmBe source emits a very high rate of 60 keV gammas, as well as neutrons with a mean energy of a few MeV which are usually coincident with a 4.4 MeV gamma[102]. Because the external NaI detector can trigger on the neutrons as well as the gammas, AmBe runs contain both electronic and nuclear recoil events. The NaI threshold is set at 125 mV, which is above the high-rate 60 keV gammas.

For this analysis, we have chosen to focus on a small subset of the runs in the second campaign which were taken under similar conditions. They generally fall into three sets: ^{22}Na runs taken with the drift field strength at 0.6 kV/cm and extraction field of 3.85 kV/cm include runs 1405, 1444, 1450, 1457, 1460, 1519, 1522, and 1523. AmBe runs taken at the same field include runs 1408 and 1447, and ^{22}Na runs taken at null-field include runs 1439 and 1456. These runs were acquired over the period from December 23, 2010 to January 12, 2011. Unless otherwise explicitly stated, results reported in this document are obtained from the 8 ^{22}Na runs with field present. Table 4.2 lists some statistics for the runs under consideration.

Run	Calibration Runs	Start time	Duration (s)	Trigger rate (Hz)	Accepted triggers	Effective livetime (s)
^{22}Na runs, drift field = 0.6 kV/cm, extraction field = 3.85 kV/cm						
1405	1403,1404	2010-12-23 13:15:08	1193	307	46787	152.322
1444	1442,1443	2011-01-05 15:47:57	1194	315	52195	165.418
1450	1448,1449	2011-01-06 09:55:18	1194	305	52326	171.078
1457	1458,1459	2011-01-07 09:38:41	1193	303	49767	164.085
1460	1461,1462	2011-01-07 16:14:28	1194	303	52244	172.175
1519	1520,1517	2011-01-12 17:03:14	292	117	12467	106.285
1522	1520,1521	2011-01-12 17:18:09	142	150	5070	33.6283
1523	1520,1521	2011-01-12 17:21:03	1193	149	48834	327.174
Total			7595		319710	1444.0
^{22}Na runs, HHV off						
1439	1436,1437	2011-01-05 13:28:12	1793	335	79059	235.908
1456	1454,1455	2011-01-07 09:00:55	1193	327	45366	138.338
Total			2986		124425	374.2
AmBe runs, drift field = 0.6 kV/cm, extraction field = 3.85 kV/cm						
1408	1406,1407	2010-12-23 13:49:45	10799	5	55179	10799.8
1447	1445,1446	2011-01-05 16:26:56	61199	5	335095	61199.9
Total			71998		390274	71999.7

Table 4.2: Statistics for the base set of runs used for analysis in this memo. The ‘‘Calibration Runs’’ column lists the laser runs from which the single photoelectron information is loaded; the first run listed is used for the top 7 PMTs and the second is for the bottom PMT. The drop in trigger rate for runs 1519, 1522, and 1523 is likely due to the ^{22}Na source unintentionally being placed in a different position than for the other ^{22}Na runs.

4.3 Low Level Analysis Description

This section contains a brief description of the steps performed in translating the digitized PMT traces to analysis variables. The structure follows that of the software modules which actually perform the analysis. In very basic terms, the analysis procedure usually consists of adding up the waveforms of each channel (scaled by gain), then, on this sum channel, identifying and subtracting the baseline, searching for physical scintillation pulses, and measuring their parameters, such as integral, amplitude, time between, etc.

4.3.1 ConvertData

This module handles the lowest level tasks of converting the raw binary information received from the digitizer into a vector waveform of samples for each channel. It also reads information such as the ID and timestamp for each trigger. Parameters:

`offset_channel` Offset one channel by the specified time (in microseconds) relative to the others. Each channel can be assigned a separate offset. Value: channel 7 (the bottom PMT) is offset by $-0.03 \mu\text{s}$ to account for observed delay. This parameter was set approximately by hand. Figure 4.7 shows a typical event, which indicates that the signal from channel 7 is still arriving later, by 10-20 ns than the 3" PMTs.

`skip_channels` Remove a channel from the analysis, usually because of poor signal. Value: Unless otherwise specified, channels 0, 3, and 6 are removed for analyses presented here.

4.3.2 SumChannels

This module takes the waveform for each of the channels, scales by the measured mean single photoelectron response, and adds them together, creating a virtual sum channel, which is assigned an ID of -2. After this point, the sum channel is treated identically to a physical channel unless otherwise specified. Section 4.4.1 describes the determination of the single

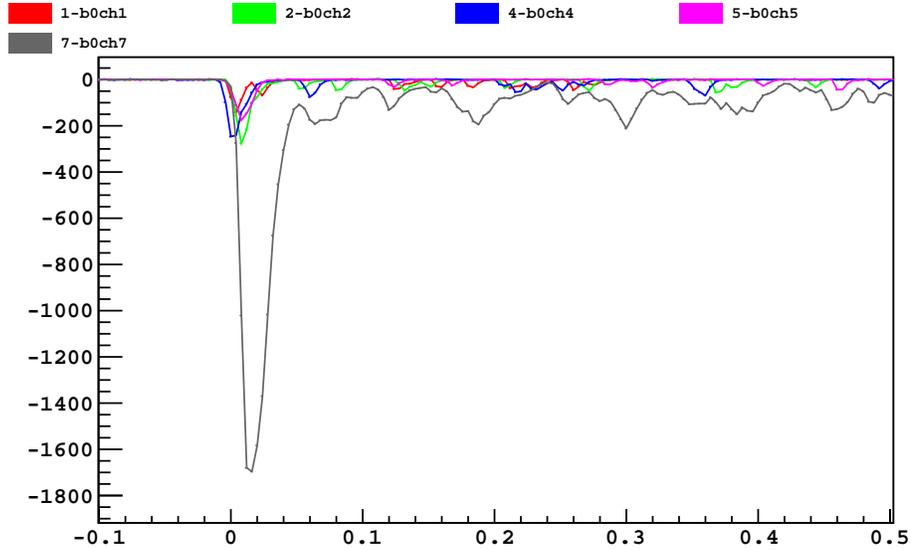


Figure 4.7: PMT waveform traces for a typical scintillation event (Run 1460, event 14). Alignment among the 3" PMTs seems good, but the 8" PMT (channel 7) is delayed, despite a 30 ns correction already applied.

photoelectron response, which is stored into and read out of a database at the beginning of run processing.

4.3.3 BaselineFinder

The digitizer is DC-coupled with a semi-arbitrary DC offset level (semi-arbitrary because, although it is controllable via software, the level does not correspond exactly to a digitizer count level). Therefore, it is necessary to determine for each channel where the baseline (or zero-voltage) level falls in terms of digitizer counts. This value is largely constant during a run, but can often fluctuate between runs when resetting the digitizer's applied DC offset.

A competent baseline finding module must be capable of identifying the baseline level of a given waveform and subsequently producing a baseline-subtracted waveform. This is normally done by analyzing several samples acquired before each trigger (the pre-trigger window), because this window is the least likely to contain any signal. Under ideal conditions, the baseline can be assumed to be a constant over an entire waveform, or over several

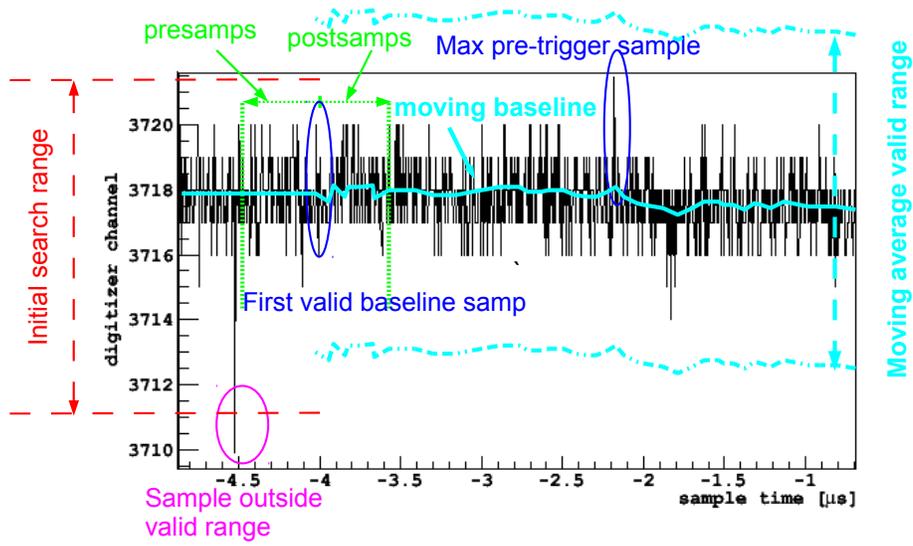


Figure 4.8: Illustration of baseline finding algorithm. See text for details.

contiguous waveforms collected within a short time. The pre-trigger window baseline, given low noise conditions, can be treated as the baseline for the entire waveform.

For the waveforms acquired in the 10 kg detector run starting Oct 2010, this assumption is unfortunately not true: one of the characteristic fluctuation time scales of the noise is on the order of several microseconds. This leads to the observation of baseline drifting on this time scale. Even more unfortunately, this noise, which recent evidence indicates is largely due to the fast amplifiers, is coherent across all channels. This means that when summing over all of the digitizer channels, the amplitude of this noise relative to the signal is not reduced. The conventional way of determining the baseline is unable to account for such features. The result of this is that the integral of the waveform often ends up in a “run away” situation, causing difficulties in downstream analysis. To account for these microsecond-timescale fluctuations, we have implemented a moving baseline, which attempts to distinguish between slow changes in the baseline level and signals coming from the PMTs.

Algorithm

The baseline is first of all assumed to be time-varying. The idea of this version of the baseline finder is to follow the drift caused by low-frequency noise. Points in a given waveform are divided into two categories: baseline points and signal points. The user defines a vertical range (in digitizer counts for individual channels, or photoelectrons for the sum channel) for baseline drift. Samples outside of this range are assumed to be real signals, as are samples in a user-defined time range before and after each signal region. Remaining regions of the waveform are considered to be baseline regions. Within baseline regions, the assumed value for the baseline at each point is calculated as a moving average, with a user-defined number of samples included before and after each point. In signal regions, the baseline is linearly interpolated between the closest surrounding moving average point. The actual algorithm is described in the following steps. Figure 4.8 illustrates the application of this algorithm to a hypothetical waveform.

1. Find the maximum amplitude data point within the pre-trigger window (the circled blue point at $\sim -2.1\mu\text{s}$ in Figure 4.8). PMT signals are assumed to be unipolar negative, while of course noise can fluctuate positive and negative. Samples in the pre-trigger window with values between this value and this value minus twice the parameter `max_amplitude` (shown by the dashed red lines in Figure 4.8 are considered to be in the baseline range.
2. Starting from the beginning of the waveform, seek forward until a sample is located such that it and all of `pre_samples` before and `post_samples` after it are within the baseline vertical range. The averaging window (greatly exaggerated for illustration) is shown by the dotted green lines in Figure 4.8. If no sample satisfying this condition is found within the pre-trigger window, this channel is flagged, and the event is removed from later analysis.
3. Calculate the average value of the `pre_samps` plus `post_samps` samples around this point; this is assumed to be the baseline value this point.

4. Once the first valid baseline sample is found, the baseline vertical range is re-defined to be the moving average \pm `max_amplitude`.
5. Step forward sample by sample, computing the baseline value at each point as the moving average within the defined window, if all samples within that window are within the baseline vertical range. The moving average is show in Figure 4.8 by the solid cyan line; the valid baseline search range by the dashed cyan lines.
6. Once a sample outside of the baseline vertical range is found, stop calculating the moving average. This is assumed to be a signal region. Skip forward until we find a new valid baseline window, searching in the same mode as for the first baseline sample. Repeat the above steps until the end of the signal.
7. For each signal region where the moving average is not computed, the baseline value at each point is linearly interpolated from the two nearest surrounding baseline points.
8. Create the baseline-subtracted waveform by subtracting the assumed baseline value from each point. This is stored in a separate waveform, so later modules can work on either the raw or baseline-subtracted waveform for each channel.

Figure 4.9 shows the sum signal for event 0 of Run 1460, zoomed around the baseline of the S2 region, showing the baseline interpolation over the S2 signal region. Though it is difficult to see on this scale, most of the individual photoelectron “spikes” are also interpolated over. This images also clearly shows the $\sim 5 \mu\text{s}$ sinusoidal background that is the main impetus for using a moving baseline algorithm.

Parameters

`max_amplitude` maximum deviation (in digitizer counts) from the previous baseline running average for a point to be considered as a baseline point for individual channels. Value: 5 counts.

`max_sum_amplitude` Same as `max_amplitude` but for the sum channel, since it is scaled to photoelectrons. Value: 0.15.

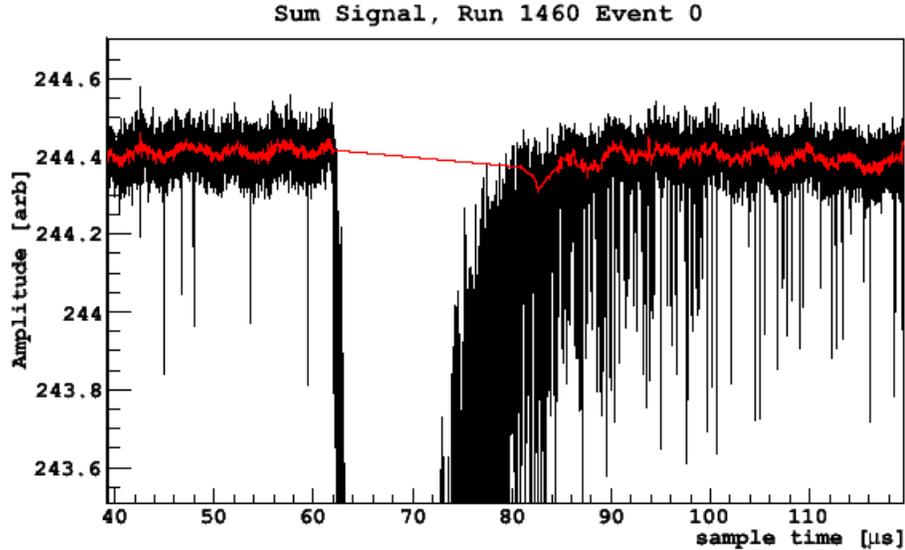


Figure 4.9: Summed signal waveform showing baseline oscillation and interpolation over S2 region.

`pre_samps` Number of samples averaged before a given point in the moving average. Value: 10.

`post_samps` Number of samples averaged after a given point in the moving average. Value: 10.

Limitations

- This algorithm makes no attempt to find the baseline under any S1 or S2 pulse. The baseline values under the pulses are simply obtained through linear interpolation. Hence, if a pulse lasts a long time, there is no compensation for the baseline drift over the pulse duration.
- It is very likely that we are operating in a regime where there is no clear cut-off between the baseline noise and single photoelectrons. If the threshold is set too high, the baseline finder could identify a significant fraction of single photoelectron peaks as part of the baseline fluctuation.

4.3.4 EvalRois

This module is used to define fixed Regions Of Interest (ROIs) in sample time for which to calculate simple statistics. A region is defined by a start and end time in microseconds, and for each region the integral is calculated (using the baseline-subtracted waveform), as well as the maximum and minimum values (raw) obtained over the interval. An “npe” variable, the integral divided by the mean spe response, is also stored. The ROI variables are used to check very basic statistics, especially for online analysis where pulse finding (which can be very slow) is not desirable.

Parameters: ‘add_roi’ can be called multiple times to add a new region. In the standard configuration, we define 4 ROIs:

1. laser window: 0.15 to 0.28 μs
2. full window -0.05 to 7 μs
3. fprompt window -0.05 0.1 μs
4. after laser window 0.3 7 μs

Because the S2 signal does not arrive at a fixed time relative to the trigger, it is not possible to define an appropriate ROI.

4.3.5 PulseFinder

Given a baseline subtracted waveform, the next step in the analysis is to identify the regions where scintillation pulses are located. The PulseFinder module attempts to identify these signal regions. This is in many ways similar to the BaselineFinder; the main difference is that PulseFinder can identify several single photoelectrons in a short time as a scintillation pulse, while BaselineFinder would identify the regions between the same single photoelectrons as baseline.

There are many possible types of triggers found in our data:

- Trigger on S1, observe a single S2 — a single scattering event

- Trigger on S1, observe multiple S2 — multiple scatterings of a single particle
- Trigger on S2. This mode is much less common when using the external coincidence trigger.
- Multiple S1 and S2
- Overlapping pulses (such as for events very near the top of the detector, where S2 starts before S1 has finished)

Given such a wide variety of events, it should be evident that no presumption about the ordering of pulses, relative amplitude or pulse duration can be made prior to processing the waveform. This version of pulse finder strictly follows this guideline.

Algorithm

The algorithm works based on the observation that each scintillation pulse (either S1 or S2) causes the integral of the waveform to rise sharply. In other terms, a pulse is identified by several photons arriving within some short time window. When the curvature (second derivative) of the integral goes above a user defined threshold, a pulse start is located. When the amplitude of the signal waveform falls below another threshold, its end is identified. Overlapping pulses are tagged if a secondary peak is present.

A down-sampling of the integral array (by a factor of 250 presently) is performed to speed up the search. This down-sampling is why we take the curvature of the integral, rather than the first derivative of the raw digitized pulse, since the down-sampled derivative would contain very little information. Pulse searching is performed only on the sum channel.

Algorithm details:

1. Traverse through the down-sampled integral waveform and compute the slope (first derivative) and curvature (second derivative) of the integral points.
2. Traverse through the curvature waveform and identify the first point that is more negative than a defined `pulse_start_curvature`. Once this is located, seek forward

- sample by sample on the original waveform until a sample is found with a value below `amplitude_start_threshold`. This is the start of a pulse.
3. Skip forward within the pulse until the first maximum is found (until the curvature becomes positive). While still within a pulse, check for a possible pileup pulse identified by the curvature falling more negative than `pile_up_curvature`.
 4. Locate the end of the pulse as the slope of the integral (i.e. the raw signal, but down-sampled) returns to within `pulse_end_slope` of zero.

Figure 4.10 demonstrates this algorithm on an example trigger. The digitized pulse is shown in black, the integral in blue (with the integral scale on the right) and the curvature in orange. The curvature line is arbitrarily scaled down by a factor of 50 and offset 5 counts from the raw waveform for visibility. Because the curvature graph is averaged to some degree (due to the down-sampling of the integral), the negative dips that mark the start of pulses appear at earlier sample time than the pulses themselves, hence the forward seeking on the original waveform after the initial curvature search. The tagged pulses are identified by the green hashed boxes.

After identifying the start and end sample of all pulses in an event, PulseFinder is also responsible for calculating several values for each pulse. Of particular note are several parameters used by other modules and offline analysis to determine the suitability of a given pulse for analysis:

`is_s1` Used to determine whether a given pulse has the time profile expected for an argon scintillation pulse (as opposed to noise, or the much broader S2 pulses). Two variables are defined: `ratio1` is the ratio of the integral of the signal from 20 ns before to 20 ns after the peak to the total integral, and `ratio2` is the ratio of the integral from the start of the pulse to 20 ns before the peak to the total integral. `is_s1` is true if `ratio1` is greater than 0.05 and `ratio2` is less than 0.02.

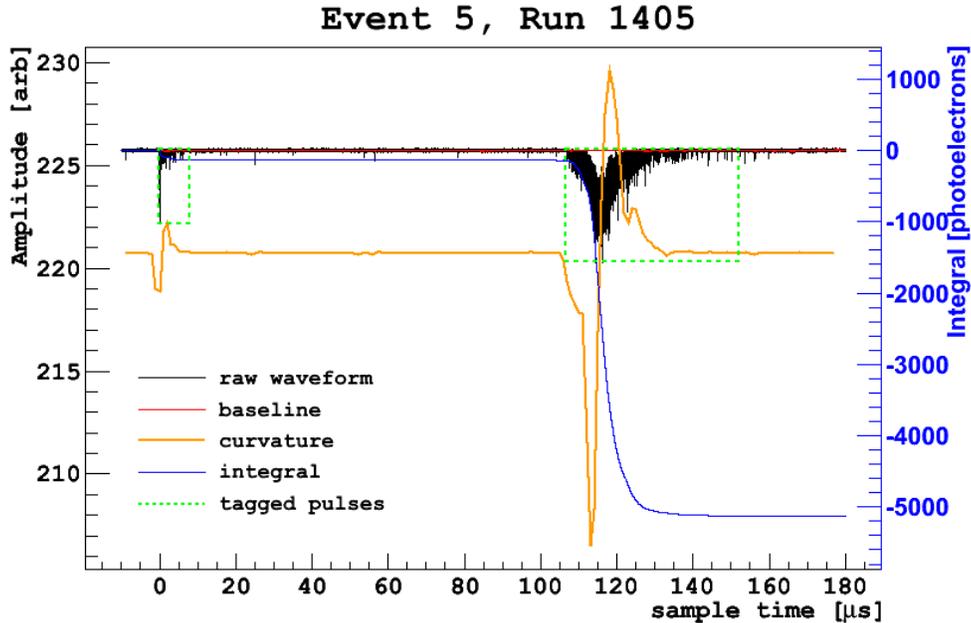


Figure 4.10: Example trigger illustrating the PulseFinder algorithm. The curvature, used to roughly identify the starts of pulses, is shown by the orange line, and is arbitrarily scaled down by a factor of 50 and offset by 5 counts from the baseline for visibility. See text for further details.

`start_clean` For the first pulse, true if `is_s1` is true. For subsequent pulses, true if it is not a pileup pulse, i.e., if the previous pulse has ended before the start of the current pulse.

`end_clean` True if the pulse has ended before the start of the next pulse and before the end of the acquisition window.

`is_clean` True if both `start_clean` and `end_clean` are true.

`fixed_int1_valid` In addition to the integral over the entire range (start to end) identified by PulseFinder, we also calculate the integral over two user-definable windows (`fixed_time1` and `fixed_time2`), where the times are appropriate for the typical lengths of S1 and S2 pulses, respectively. `fixed_int1_valid` is used to ensure that the `fixed_int1` value is considered over a valid range; i.e. it is true if `start_clean` is true and if the sample at `fixed_time1` after the start of the pulse is before the start of

the next pulse and before the end of the trigger window. Notice that this can be true even if the pulse length (as identified by the algorithm) is shorter than `fixed_time1`.
`fixed_int2_valid` See previous.

Parameters

`search_mode` Allows the user to select between different algorithms to perform the pulse search (largely historical). The method presently used (and described in this chapter) is labeled CURVATURE.

`down_sample_factor` Down sample the integral vector size by this factor. Value: 250.

`pulse_start_curvature` Identify the approximate start of a pulse when the curvature becomes more negative than this value. Value: -6 .

`amplitude_start_threshold` Once the start of a pulse is identified roughly, find the exact start when the signal falls below this value. Value: 0.3.

`pile_up_curvature` Identify a second pileup pulse when the curvature within a pulse region but after the first peak becomes more negative than this value. Value: -50 .

`pulse_end_slope` Identify the end of a pulse when the down-sampled signal (or slope of the integral) rises above this value. Value: -1 .

`fixed_time1` Fixed time starting at the beginning of the pulse over which to evaluate and store the signal integral appropriate for S1 pulses, in order to be less biased by the algorithm's identification of the pulse end. Value: $7 \mu\text{s}$.

`fixed_time2` See previous, with integration time appropriate for S2 pulses. Value: $30 \mu\text{s}$.

At this point, the parameters used for analysis have not been tuned or determined in a methodical way, rather they have all been simply estimated by hand. Since this is one of the most critical steps in the analysis chain, identifying the effect of these parameter choices is one of the primary steps necessary for improving the analysis.

4.3.6 S1S2Evaluation

This module takes the pulses identified by PulseFinder on the sum channel and evaluates them on each of the individual channels (using the identified start and end times). It then adds the statistics for the pulses of each physical channel together to determine the total values for S1, S2, F_{prompt} , etc. for the entire event. These values may differ from those evaluated on the pulses of the sum channel because the baseline of each channel is evaluated individually. These differences are discussed in more detail in section 4.4.3.

In addition, this module attempts to determine whether the event contains only a single valid S1 pulse and a single valid S2. This evaluation is done using the pulses identified on the sum channel. An S1 pulse is valid if it is the first pulse of an event and fits the 'is_clean' and 'is_s1' criteria from PulseFinder. Similarly, an S2 pulse is valid if it is the second pulse, passes 'is_clean', and *not* 'is_s1'. Finally, the event is declared 's1s2_valid' if there is both a valid S1 and valid S2 and the number of pulses found is exactly 2. In addition, for each `sX_valid` variable, there is a `sX_fixed_valid` variable, which checks the `fixed_int1_valid` value (for S1, or `fixed_int2_valid` for S2). The `fixed_valid` variables allow us to keep pulses which are well separated for the majority of the signal time, but for which the end time may not be correctly identified by PulseFinder.

4.3.7 PositionRecon

This module takes events marked as valid by S1S2Evaluation and calculates the drift time (and z-position) and x-y barycenter of the event. The barycenter is calculated as the average of the x-y position of each of the top PMTs weighted by the amount of light collected during the S2 for that event. I.e.,

$$\bar{\mathbf{r}} = \frac{\sum_{i=0}^{N_{topPMTs}} \mathbf{r}_i \cdot S2_i}{\sum S2_i} \quad (4.1)$$

where \mathbf{r}_i is the (x, y) position of PMT i and $S2_i$ is the number of photoelectrons detected during S2 by that PMT.

Parameters:

geometry Allows functioning of other detector geometries. Only currently defined value currently is TenKG.

drift_speed Electron drift speed, which is a function of the applied drift field. Generally not used, as drift time is a more reliable variable than z position. In principle, if the drift speed as a function of field were known, it would allow to compare runs with different drift fields. Value: 1 mm/ μ s.

use_full_s2 Should the value of S2 for each channel be that taken over `fixed_int2` (30 μ s) (false), or the full length of S2 as identified by PulseFinder (true)? Value: true.

use_full_valid To determine if the event is valid, should the `s1s2_valid`(true) or `s1s2_fixed_valid`(false) value be checked? Value: false

4.4 Calibration and Efficiency of Algorithms and Cuts

4.4.1 Single Photoelectron Fit

For details on the single photoelectron fit procedure, see Reference [103]. The single photoelectron spectrum is created from dedicated runs which are triggered by the laser pulser. The baseline-subtracted spectrum for each channel is integrated over the fixed window 0.15-0.3 μ s after the trigger time, which is the time window over which the laser pulses arrive. The spectrum of these fixed-window integrals is histogrammed up to 300 counts \times samples (the single photoelectron mean is set by adjusting the PMT HV supply to around 80 counts \times samples). Because the DAQ triggers every time the laser fires (and not, e.g., based on discriminating the PMT output), the probability for n photons to be detected on a given channel for a single trigger follows a Poisson distribution

$$P(n; \lambda) = \frac{\lambda^n e^{-\lambda}}{n!}, \quad (4.2)$$

where λ is the average occupancy for that channel, and is set by filtering the laser to be on the order of 0.1. The spectrum is fit to the function

$$F(x) = F_{ped}(x) * \sum_{n=0}^N P(n; \lambda) F_n(x) \quad (4.3)$$

where F_{ped} is the pedestal response function (i.e., integration over an empty window), the ‘*’ indicates convolution, and F_n is the PDF of the response for n photoelectrons. In principle, the upper limit of the summation N should run to infinity, but in practice we continue the sum to 7 photoelectrons.

The pedestal response is assumed to be gaussian, centered at 0, with width determined by the fit. If the single photoelectron response is pure gaussian as is often assumed, then the n th photoelectron response, including the pedestal convolution, will also be gaussian, with width $\sigma_n^2 = n\sigma_1^2 + \sigma_{ped}^2$. In early development, we found that the pure gaussian photoelectron response would not fit the measured spectrum unless we also allowed for an exponential noise source; even then, the fits were often very dependent on initial fit parameters. Following the work of Dossi et. al. [104], we now assume that the exponential component is actually part of the single photoelectron response (possibly due to uneven multiplication due to electrons hitting the edge of the dynodes, etc). Moreover, the single photoelectron response should not contribute negative charge, so must be zero for $x < 0$. Functionally, the single photoelectron response function is

$$F_1(x) = \begin{cases} A \cdot N(x; \mu_1, \sigma_1) + (1 - A) \cdot e^{-Bx} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \quad (4.4)$$

where $N(x; \mu, \sigma)$ is a normal distribution with mean μ and standard deviation σ . The two-photoelectron distribution is the convolution of $F_1(x)$ with itself, and the three-photon distribution is the convolution of the two-photon distribution with the one-photon distribution. In practice, there are no analytic solutions for the $n \geq 3$ distributions. For this

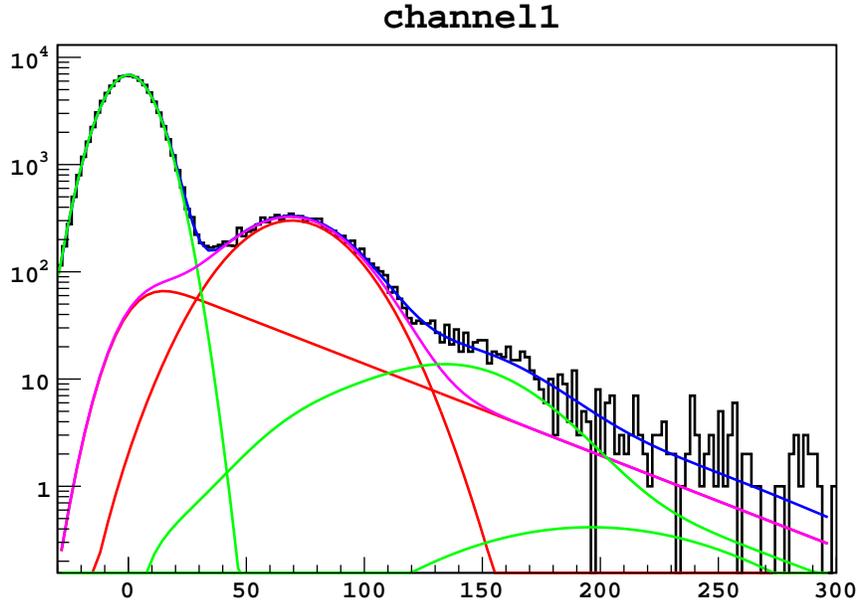


Figure 4.11: Fit to the laser response spectrum for channel 1 from Run 1461. The total response is given by the blue line. The leftmost green line is the pedestal response. The total single photoelectron response is shown by the magenta line, with the gaussian and exponential components in red. Finally, the two photoelectron response and the sum of the $3 \leq n \leq 7$ response are shown in the remaining two green lines.

reason and because the amplitude for $n \geq 3$ photoelectron distributions is quite low, we use a simple gaussian to model these, with $\sigma_n^2 = n\sigma_{1^*}^2$, where σ_{1^*} is the standard deviation of the full (gaussian plus exponential) single photoelectron response. Figure 4.11 shows an example laser response fit for channel 1 from Run 1461, with each of the components drawn separately.

We chose the gaussian-plus-exponential PDF over a pure gaussian because it seems to provide a more robust fit to our dataset over different conditions. However, the pure gaussian fit, while often being less robust (i.e., having a stronger dependence on the initial fit conditions and data cleanliness), generally had a goodness-of-fit comparable to the gaussian-plus-exponential model. Therefore, the difference in mean of the two PDFs should be treated as a systematic uncertainty when calculating scales in photoelectrons, for example the detector light yield. (Note that this does not translate into an uncertainty in the overall energy

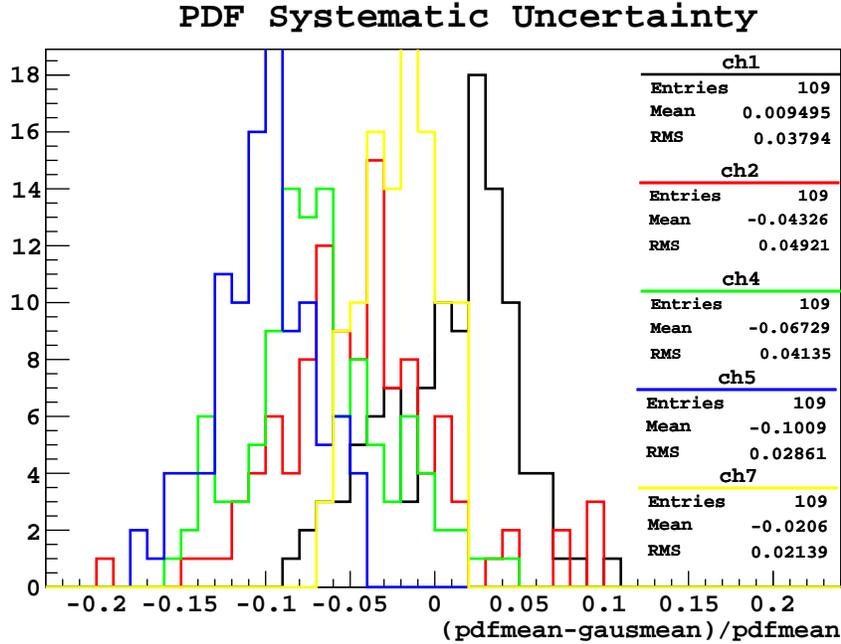


Figure 4.12: Fractional difference for the mean single photoelectron integral between a gaussian-plus-exponential PDF and a pure gaussian PDF.

scale, merely a redefinition of the ‘photoelectrons’ unit.) Figure 4.12 shows the fractional difference between the mean of the total PDF and the mean of just the gaussian portion when using the gaussian-plus-exponential photoelectron PDF for the PMTs used in the main analysis. (Technically, the fit result for the gaussian mean might be different when using the pure gaussian PDF for the fit, but this difference is completely negligible). As the figure shows, the exponential part of the PDF tends to pull the mean down from the gaussian mean by a few percent. As we’ll show in Section 4.5.1, channel 7 tends to account for some 50-60% of the measured scintillation signal, so if we average the channels with this weighting taken into account, we would get a 3-4% lower measurement for the light yield if we used the pure gaussian PDF.

4.4.2 Analysis Cuts

Table 4.3 lists the event selection cuts applied for this analysis. The general cuts are applied to all events, regardless of the type of run they come from, and are designed to eliminate

events which are either outside the range of the electronics or are not physical scintillation pulses (i.e., triggering on the tails of large (muon) signals or possibly triggering on electronics noise). Two-pulse cuts are applied to runs taken with drift and extraction fields turned on, so that an S2 signal is present, and are designed to select single-sited events, which have exactly one well-defined S1 signal and exactly one well-defined S2 signal. Here and in later discussion, when discussing signal amplitudes for interactions in the argon, we assume that the integral over some time of the signal in one channel, divided by the single photoelectron response as determined by calibration triggers, is the number of photoelectrons (p.e.) collected by that channel over that time interval. The total number of p.e. collected in that interval is the sum of the p.e. collected on each channel.

Standard Runs

Table 4.4 reports the percentage of events removed by each cut when it is applied as either the first cut (i.e., events that would be removed if that were the only cut employed) or last cut (i.e., events which would not be removed if not for that particular cut) for the standard set of ^{22}Na coincidence runs with field applied. Overall, the event acceptance for this set of runs is $14 \pm 1\%$.

Runs with Different Fields

The cut acceptances for runs with different fields were evaluated in the same way as for the standard data set, and are presented in Tables 4.5 and 4.6. All of the runs considered are uncollimated ^{22}Na coincidence runs taking place after 12/07/2010 with at least 10,000 triggers. For each field setting, the average number of events cut was taken for all runs with that field value. The errors are the standard deviation of the runs from the mean. The average acceptance for all ^{22}Na coincidence runs with field is $24 \pm 6\%$ and $76 \pm 14\%$ for runs without field.

Standard Cuts	
1. min S1	Integral over the S1 pulse window identified by PulseFinder must be greater than 50 p.e.
2. max S1	Integral over the S1 pulse window identified by PulseFinder must be less than 10000 p.e.
3. vertical saturation	Signal for each channel must remain within the 2 V input range of the digitizer over the entire acquisition window.
4. S1 length	The S1 pulse must have ended (as identified by PulseFinder) within 20 μs of the start of the pulse.
5. baseline	BaselineFinder (section 4.3.3) must have successfully found a valid baseline for each channel. Since the baseline is found independently for the sum channel, it is possible to have defined pulses, etc. for the sum but not individual channels.
6. S1 start time	The start time of S1, as determined by the PulseFinder (section 4.3.5), must be within ± 100 ns of the digitizer trigger.
Two-Pulse Cuts	
7. drift time	Drift time between S1 and S2 must be greater than 20 μs to assure good separation of the signals.
8. min S2	Integral over the S2 pulse window identified by PulseFinder must be greater than 10 p.e.
9. S2 95% time	95% of the total light collected during the S2 pulse must have arrived within 10-30 μs of the start of the pulse. This is largely to remove pileup S1 pulses which might otherwise be mistaken for S2.
10. <code>s1s2_fixed_valid</code>	This boolean combines several cuts: the first pulse must pass the <code>is_s1</code> criteria (see section 4.3.5) and be separated from the second pulse by at least 7 μs (superseded by the drift time cut). The second pulse must start after the first pulse has ended and at least 30 μs from the end of the trigger window, and must <i>fail</i> the <code>is_s1</code> criteria. Finally, there must have been exactly 2 pulses found by PulseFinder.
One-Pulse Cuts	
11. <code>s1_fixed_valid</code>	Individual cuts in <code>s1s1_fixed_valid</code> (defined above) which refer to the first pulse.
12. one pulse	PulseFinder must have identified exactly one pulse.

Table 4.3: Description of event selection cuts employed in this analysis. Other than for efficiency checks, all cuts are applied simultaneously; the numbering here is simply order enumerated in the analysis software. Note that “S1” and “S2” here apply respectively to the first and second pulse found for a trigger, not the results of `is_s1` or similar checks.

Run	Cut									
	1	2	3	4	5	6	7	8	9	10
when cut is applied first										
1405	5.7	4.0	12.6	6.0	0.1	10.4	31.0	2.3	25.5	83.5
1444	5.9	3.9	13.8	6.1	0.1	10.5	29.9	2.4	25.5	82.6
1450	6.8	3.8	15.1	6.6	0.1	12.5	31.9	2.7	27.1	83.7
1457	5.9	3.7	12.5	5.8	0.1	10.6	29.8	2.2	25.4	82.4
1460	5.9	3.9	12.4	6.0	0.1	11.1	29.8	2.2	26.0	83.2
1519	8.5	2.2	21.6	6.5	0.1	14.9	29.3	5.2	34.4	81.2
1522	8.1	2.0	22.6	6.5	0.1	14.7	28.7	4.9	34.6	81.0
1523	8.1	1.8	22.4	6.1	0.1	14.4	28.3	5.2	34.5	80.6
Average	6.9	3.2	16.6	6.2	0.1	12.4	29.8	3.4	29.1	82.3
Std. Dev.	1.1	0.9	4.4	0.3	0.0	1.9	1.1	1.3	4.2	1.1
when cut is applied last										
1405	0.4	0.0	1.3	0.2	0.0	0.0	0.4	0.0	0.0	30.4
1444	0.5	0.0	1.7	0.2	0.0	0.0	0.3	0.0	0.0	29.9
1450	0.4	0.0	1.7	0.2	0.0	0.0	0.3	0.0	0.0	28.4
1457	0.5	0.0	1.4	0.2	0.0	0.0	0.4	0.0	0.0	30.5
1460	0.5	0.0	1.4	0.2	0.0	0.0	0.3	0.0	0.0	30.8
1519	0.4	0.0	2.9	0.3	0.0	0.0	0.1	0.0	0.2	23.3
1522	0.3	0.0	3.5	0.3	0.0	0.0	0.1	0.0	0.1	23.3
1523	0.4	0.0	3.4	0.4	0.0	0.0	0.1	0.0	0.2	23.4
Average	0.4	0.0	2.2	0.3	0.0	0.0	0.2	0.0	0.1	27.5
Std. Dev.	0.1	0.0	0.9	0.1	0.0	0.0	0.1	0.0	0.1	3.3

Table 4.4: Percent of events cut by each of the standard cuts listed in table 4.3 for the set of ^{22}Na coincidence runs with drift field of 0.6 kV/cm and extraction field of 3.85 kV/cm used as the standard dataset for this analysis. The second set indicates that cuts 2, 5, 6, and 8 are completely superfluous when the other cuts are applied first.

Figures 4.13 and 4.14 show the fraction of events removed by each cut as a function of the applied drift and extraction fields, respectively. There is no clear dependence on the drift field, but in general more events tend to be rejected with increasing extraction (and multiplication) field. The cuts most responsible seem to be the vertical saturation cut and the `s1s2.fixed.valid` cut, which handles ensuring that there is only a single S1 and S2. This makes sense, as, with higher extraction and multiplication fields, S2 becomes larger, so large S2 events will more often exceed the vertical range of the digitizer, and more small S2 pulses will be identified.

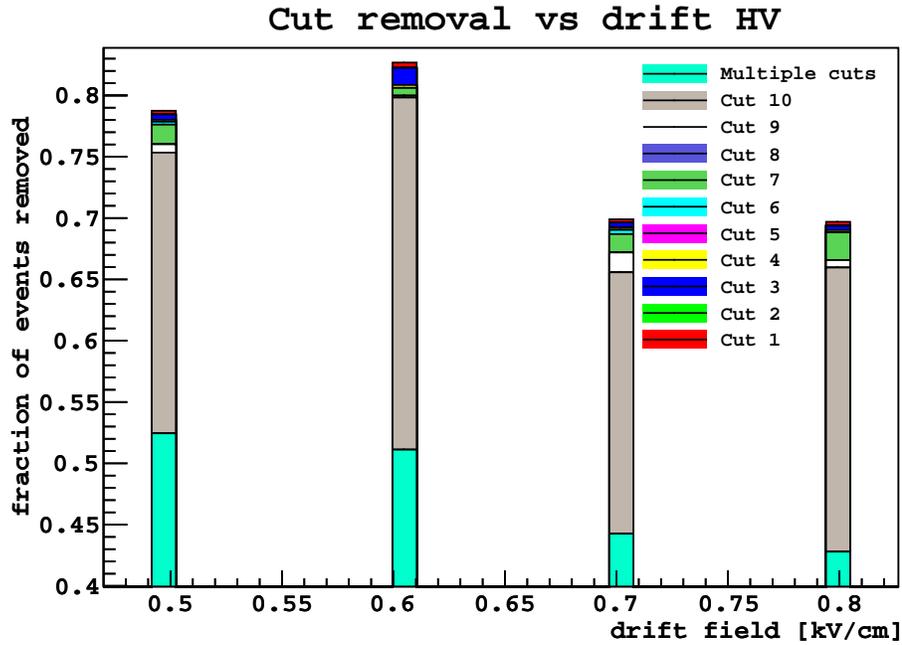


Figure 4.13: Fraction of events removed by each cut as a function of applied drift electric field. Cut numbers are those from table 4.3.

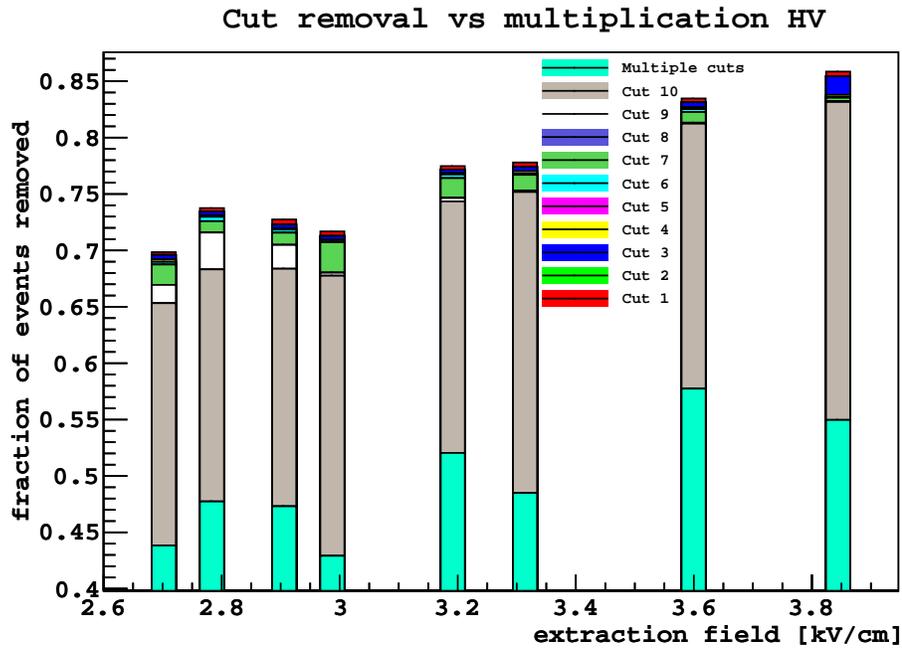


Figure 4.14: Fraction of events removed by each cut as a function of applied extraction field. Cut numbers are those from table 4.3.

Fields	Cut 1	Cut 2	Cut 3
(3.3, 0.5)	7.96/1.47±0.65/0.19	4.09/0±0.29/0	5.31/1.41±0.29/0.18
(3.6, 0.5)	11.39/2.07±0.07/0.21	5.72/0±0.26/0	7.32/2.63±0.06/0.12
(2.7, 0.5)	2.76/0.52±1.15/0.04	0.53/0±0.38/0	4.91/1.78±0.34/0.17
(2.7, 0.8)	2.68/1.38±0.67/0.32	0.29/0±0.17/0	4.39/1.27±0.45/0.34
(2.7, 0.7)	3.73/1±1.82/0.35	0.5/0±0.41/0	4.99/1.97±0.31/0.59
(0, 0)	1.22/1.12±0.05/0.07	0.07/0±0.04/0	7.94/7.93±2.66/2.73
(3.85, 0.5)	9.73/3.03±5.79/0.66	4.81/0±1.38/0	10.96/5.59±1.26/2.07
(3.0, 0.1)	2.24/1.01±0.13/0.12	0.73/0±0.29/0	12.42/1.92±1.27/0.15
(3.0, 0.4)	2.42/1.88±0/0	1.02/0±0/0	5.93/1.87±0/0
(3.0, 0.7)	2.54/1.99±0/0	1.21/0±0/0	4.97/1.34±0/0
(2.7, 0.6)	2.68/1.59±0/0	0.37/0±0/0	4.69/1.63±0/0
(3.3, 0.6)	4.25/2.11±0.82/0.39	2.33/0±0.77/0	5.84/2.1±0.54/0.31
Fields	Cut 4	Cut 5	Cut 6
(3.3, 0.5)	12.61/0.64±0.63/0.15	4.95/0±0.42/0	19.23/1.32±0.72/0.14
(3.6, 0.5)	16.07/0.89±0.12/0.05	7.63/0±0.06/0	24.07/1.43±0.23/0.24
(2.7, 0.5)	5.16/0.94±2.57/0.06	0.57/0±0.8/0	6.93/0.74±6.42/1.04
(2.7, 0.8)	2.73/0.61±0.71/0.2	0.02/0±0.01/0	2.69/0.01±1.16/0.02
(2.7, 0.7)	5.17/0.75±3.64/0.25	0.99/0±1.12/0	9.18/1.09±7.85/1.48
(0, 0)	0.01/0±0.01/0	0.01/0±0.01/0	0.07/0.01±0.01/0.01
(3.85, 0.5)	11.85/1.08±8.56/0.22	4.57/0±6.39/0	17.87/0.64±12.08/0.9
(3.0, 0.1)	3.76/1.16±1.16/0.37	0.01/0±0/0	3.26/0±0.19/0
(3.0, 0.4)	3.55/0.75±0/0	0.02/0±0/0	2.66/0.01±0/0
(3.0, 0.7)	3.48/0.61±0/0	0.02/0±0/0	2.61/0.01±0/0
(2.7, 0.6)	3.48/0.82±0/0	0.02/0±0/0	3.68/0±0/0
(3.3, 0.6)	4.3/0.95±0.78/0.31	0.06/0±0.03/0	6.44/0.01±1.62/0.03

Table 4.5: Percentage of events cut by standard cuts when cut is applied first/last. The cuts are as numbered in Table 4.3.

4.4.3 Energy and Fprompt variables

Variable Definitions

There are a number of possible variables to use to represent the energy of an event for analysis. Although different particle species will have different effective light yield (photoelectrons detected per unit of energy deposited in the active volume) due to quenching effects, it is common to measure all energies in terms of photoelectrons, which on average is directly proportional to an electron-equivalent energy deposition, except at low electron energies where dE/dx becomes large. Because photons arrive over a long period of time (i.e., the decay

Fields	Cut 7	Cut 8	Cut 9	Cut 10
(3.3, 0.5)	35.3/7.7±0.7/0.5	9.5/0±0.5/0	30.3/0.6±0.6/0.2	72.3/51.0±1.2/1.5
(3.6, 0.5)	38.0/5.4±0.1/0.2	12.1/0±0.04/0	34.0/0.6±0.8/0.2	77.4/58.5±0.8/0.5
(2.7, 0.5)	27.4/8.5±2.8/0.3	4.3/0±1.7/0	27.6/4.1±2.3/0.03	63.6/40.8±2.0/0.81
(2.7, 0.8)	24.0/7.7±5.2/3.9	3.4/0±0.8/0	27.9/2.0±4.4/0.4	63.9/41.4±4.0/4.6
(2.7, 0.7)	27.6/7.9±14.3/11.4	7.5/0±9.4/0	31.5/8.2±16.5/17.0	61.2/37.7±6.3/3.0
(3.85, 0.5)	34.0/3.4±7.6/1.1	7.5/0±7.8/0	30.2/0.3±7.1/0.1	79.7/62.2±1.2/2.4
(3.0, 0.1)	14.4/2.7±2.8/1.5	2.0/0±0.5/0	33.3/3.4±1.4/0.2	82.7/69.4±0.1/0.3
(3.0, 0.4)	23.9/7.5±0/0	2.7/0±0/0	23.7/0.7±0/0	66.5/46.6±0/0
(3.0, 0.7)	27.2/9.7±0/0	3.0/0±0/0	22.4/0.5±0/0	65.3/44.7±0/0
(2.7, 0.6)	28.2/10.6±0/0	3.7/0±0/0	29.0/9.2±0/0	65.1/45.1±0/0

Table 4.6: Percentage of events cut by two-pulse cuts when cut is applied first/last. Cuts are numbered as in Table 4.3.

times of the excited exciton states are long) and have different singlet/triplet ratios, the peak amplitude of a given scintillation pulse is only roughly correlated with the number of photoelectrons detected, and so the best estimate of the number of photoelectrons detected is to integrate the PMT response (which yields a total amount of charge collected) and divide by the average charge collected by a single photoelectron.

We commonly work in units of counts \times samples for raw integrals (both for single photoelectron and scintillation pulses), and photoelectrons (or pe) once the integral is scaled to the single photoelectron response. The CAEN V1720 digitizer has a 2V full-scale range with 12-bit resolution and 50 Ω input impedance with a 4 ns sample time, so one count \times sample is equal to 3.91×10^{-2} pC. If one wants to calculate the actual gain of a PMT, it is necessary to scale the measured charge down by a factor of 5, due to the x10 amplifiers and 50 Ω back-termination on the PMT bases.

For each of S1 and S2, there are four possible variables to choose from for defining the pulse integral, using either a fixed integration length or integrating over the exact pulse window identified by PulseFinder, and by either taking the integral of the sum of all channels (from here on referred to as IoS) or summing the integral of each channel (SoI). The fixed integration lengths are 7 μ s for S1 and 30 μ s for S2 (see section 4.3.5); note that these are generally shorter than the pulse window, but not always (see the next subsection for discussion of pulse lengths). The primary advantage of the fixed integration window is that it is not sensitive to PulseFinder correctly identifying the end of a real pulse. For this analysis, we used the full integral unless otherwise specified.

For the IoS approach, to calculate the integral, we form the sum channel as described in section 4.3.2, subtract the baseline, and calculate the integral as the simple sum of all the samples within the identified pulse window. This value is obtained in the output ROOT file via `event.GetPulse(n,-2)->npe` for the full integral and `event.GetPulse(n,-2)->fixed_int(n+1)` for the fixed integral, where $n=0$ for S1 and $n=1$ for S2.

For SoI, once a pulse has been found on the sum channel, we calculate the integral over the identified region for each other channel (subtracting the baseline), scale each for the single photoelectron response, and then add these values together. These variables are accessible in the ROOT tree by `event.sn_full` for the full integral and `event.sn_fixed` for the fixed integral, where $n=1$ for S1 and $n=2$ for S2. Unless otherwise specified, we use `event.sn_full` for the relevant energy variable.

As an estimator for the singlet-to-triplet ratio, we use the parameter $F90$, which is defined as the ratio of charge collected over the first 90 ns from the start of the pulse window to the total integrated charge. If the lifetime of the triplet state is $1.6 \mu\text{s}$, $F90$ should be related to the true singlet fraction by

$$F90 = P_s + (1 - P_s) * (1 - \exp(-90/1600)), \quad (4.5)$$

since τ_S , the singlet lifetime, is $\ll 90$ ns. The exponential term is approximately 5%. Like the total energy, there are four possible variables for $F90$, with the total charge measured over the full or fixed integral, and using either the SoI or IoS methods. Again, for this analysis we use the SoI method over the full pulse window, which is accessed by `event.f90_full`.

Checks on pulse length

Figure 4.15 shows the lengths of S1 and S2 pulses as identified by PulseFinder for events in the ^{22}Na coincidence runs which pass all cuts. The “spiky” nature of the S1 distribution is due to the fact that all S1 pulses begin within 50 ns of the trigger time (enforced by cuts, but largely true in any case), and the end of the pulse is only found with microsecond precision (i.e., 4 ns downsampled by a factor of 250). Otherwise the S1 distribution is approximately what one would expect for a pulse with a $1.6 \mu\text{s}$ lifetime (ending on average around $5-6\tau$.) S2 pulses tend to be much longer, around $40 \mu\text{s}$ or more. The second peak in the S2 time distribution near $160 \mu\text{s}$ is functionally the full acquisition time after the trigger ($180 \mu\text{s}$)

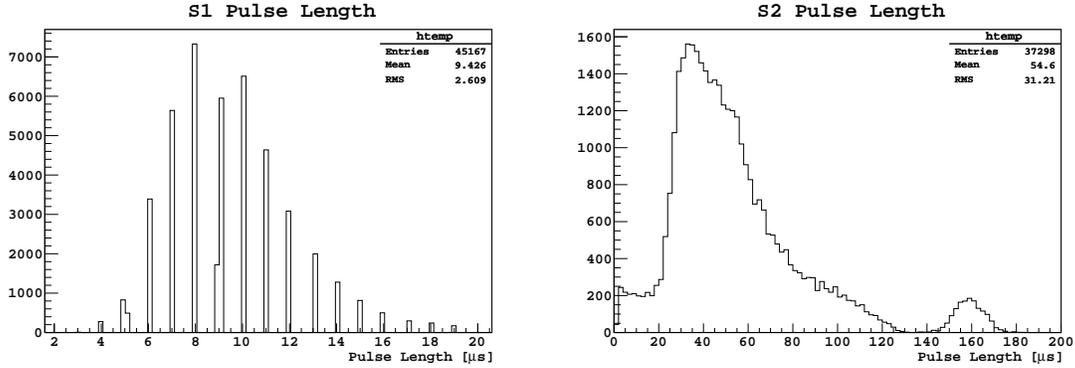


Figure 4.15: Distribution of S1(left) and S2(right) pulse lengths as identified by PulseFinder. All standard cuts are taken on the S1 length; for S2, the only cuts are that the total number of pulses is exactly 2 and that the pulse must have ended “naturally” before the end of the acquisition window.

minus the S1 pulse length. Therefore, these are the longest possible pulses that can be detected by our algorithm as two separate pulses.

In order to estimate the validity of the fixed integration windows (which are generally not used to determine signal amplitude, but to check validity of pulses), we calculate for each pulse the parameter t_{95} , which is the time at which the integral reaches 95% of the total value. For S1 pulses, (the first pulse in any event which passes `is_s1`; no other cuts applied), $97.6 \pm 0.1\%$ of pulses in the dataset have t_{95} less than $7 \mu\text{s}$. For S2 pulses (the second pulse of a trigger, which must pass `!is_s1`), $98.1 \pm 0.1\%$ have t_{95} less than $30 \mu\text{s}$. Therefore we require that a valid S1 pulse be separated from the next pulse (or the trigger window) by at least $7 \mu\text{s}$, in order to ensure that we have collected at least 95% of the total charge, and that S2 have at least $30 \mu\text{s}$ over which to integrate. In this way we ensure that, even if, for example, some small tail of S1 is still present at the beginning of S2, we have collected most of the integral (for most pulses), and so we should not reject the event.

As a check on the efficiency of PulseFinder’s ability to correctly identify the end of a scintillation pulse, figure 4.16 is a plot of the fixed integration value vs the full integration value (`event.s1_fixed` vs `event.s1_full`) for events in the dataset which pass cuts (see section 4.4.2 for a list of all cuts used). The fit line has a fixed zero-intercept and returns

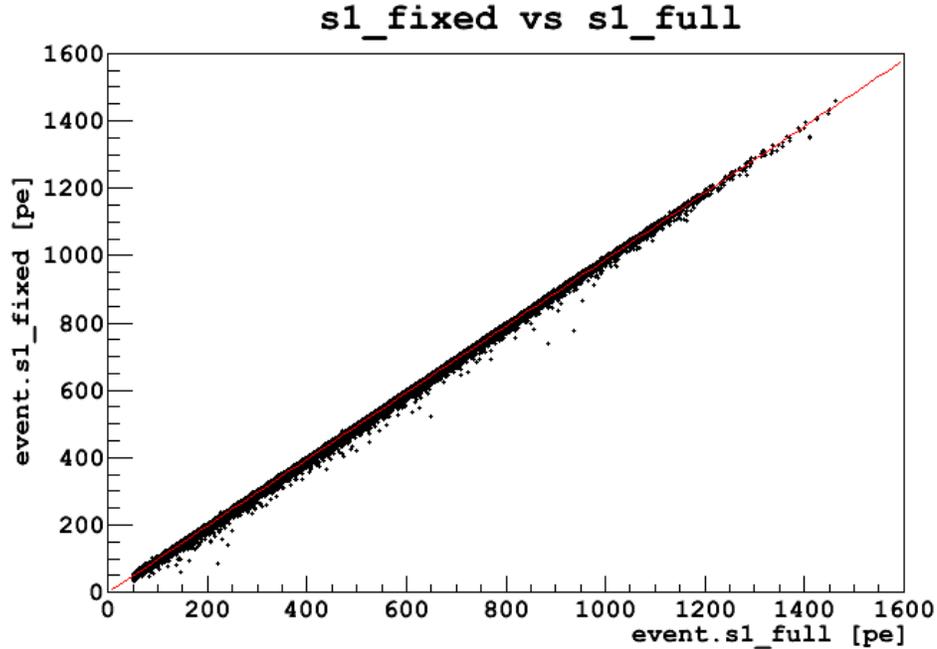


Figure 4.16: Fixed time integral vs full integral for S1 pulses in the dataset which pass all standard cuts. The linear fit has fixed zero intercept and returns a slope of approximately 0.98.

a slope of approximately 0.98 (uncertainties returned by MINUIT are meaningless without error bars). There are no gross outliers for which `s1_fixed` is much greater than `s1_full`. This indicates that PulseFinder does not greatly underestimate the length of a pulse, and so it is not necessary to place a minimum cut on pulse length. Pileup events are one exception to this conclusion, but they are removed by the cut which requires that pulses be separated by at least the fixed integration window.

Integral-of-sum versus sum-of-integral

As mentioned in the variable definitions section, slight differences can arise between the integral-of-sum (IoS) and sum-of-integrals (SoI) definition of the energy response due primarily to the BaselineFinder’s response to differently-scaled channels. Since neither approach is inherently more correct, the difference between the two approaches represents a systematic uncertainty in these variables.

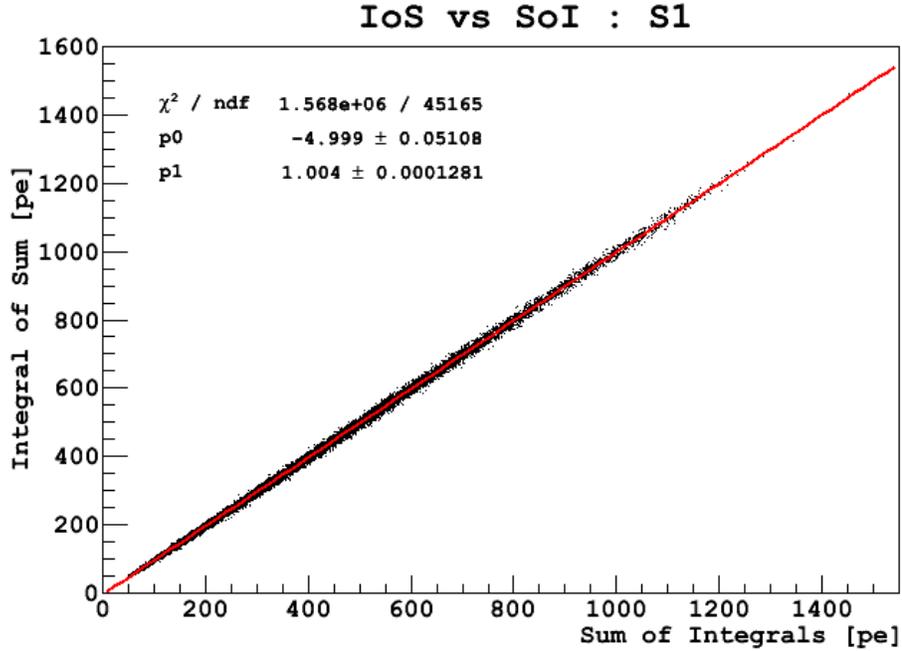


Figure 4.17: Integral of sums versus sum of integrals of the S1 pulse for events in ^{22}Na coincidence runs passing cuts.

Figure 4.17 shows the graph of IoS vs SoI for events from the ^{22}Na coincidence runs passing cuts. The scatter can be fit to a linear function with slope of ~ 1 , but with an offset (zero-intercept) of -5 ± 0.05 pe. Therefore, the difference between the two variables is not a fixed fraction.

To attempt to evaluate the magnitude of this difference, Figure 4.18 plots the mean fractional difference between the IoS and SoI approaches to calculating S1, as a function of the IoS S1 (the standard definition). As the plot shows, there is both a spread and, at low S1, a systematic offset between these two definitions for the energy. The mean offset is fairly well described by an exponential fit of $0.066 \exp(-S1/203)$. Based on this fit, the mean fractional difference in the SoI from IoS is less than (1%, 2%, 3%, 4%, 5%) for IoS S1 greater than (380, 240, 160, 100, 55) photoelectrons. If the mean systematic shift is subtracted from each point, the spread of the remaining distribution does not depend strongly on S1, and has an RMS of 2%.

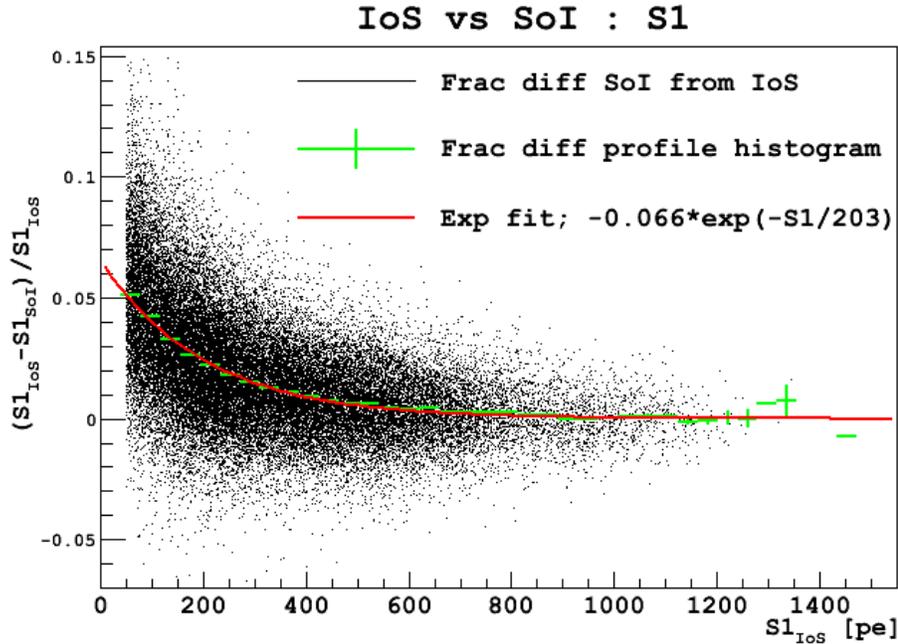


Figure 4.18: Fractional difference between S1 as calculated using the Integral-of-sum (IoS) and sum-of-integral (SoI) approach. The distribution in green is the profile histogram of the scatterplot, and is fit with an exponential function.

Figure 4.19 plots the fractional difference in $F90$ as calculated by either the IoS or SoI method as a function of S1. Similar to the S1 distribution, the mean fractional offset of $F90$ can be described by an exponential function, $-0.074 \exp(-S1/216)$. The mean fractional offset in $F90$ is less than (1%, 2%, 3%, 4%, 5%) for S1 greater than (430, 280, 200, 130, 86) photoelectrons. After subtracting the mean shift, the spread of the remaining distribution depends more strongly on energy than the S1 shift. The spread is approximately 1% at 1000 pe, and approximately 3% at 50 pe. The total distribution has an RMS of approximately 2.1%. The shift does not show any dependence on the value of $F90$ itself. In both the $F90$ and S1 offset distributions, there are very few if any obvious outliers in the distributions; visual inspection of these events did not yield any obvious causes for the offsets.

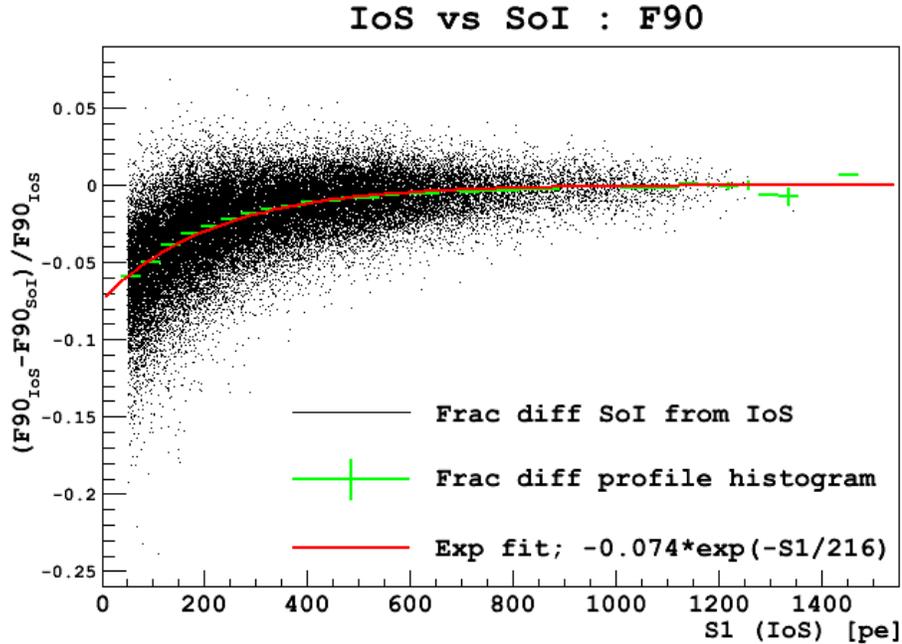


Figure 4.19: Fractional difference between $F90$ as calculated using the Integral-of-sum (IoS) and sum-of-integral (SoI) approach. The distribution in green is the profile histogram of the scatterplot, and is fit with an exponential function.

4.4.4 Moving Baseline

Averaging Single Photoelectrons

The baseline evaluation procedure allows for a moving baseline, which sets the baseline to an average of nearby points, so long as those points fall under a certain threshold. The official parameters average 10 samples before and after each time point to find the baseline for that point.

One problem with this method is that it can average away small photoelectron signals. A signal that does not trigger the interpolation will itself shift the baseline, subtracting away a noticeable part of the signal. This section attempts to quantify the degree to which this occurs using the official parameters. Potential problems come in two modes.

The first potential problem is fitting the SPE response, when small signals are in abundance. If some small SPE signals are diminished by the baseline finder averaging, this could artificially alter the SPE signal distribution and the determined PMT calibration. The base-

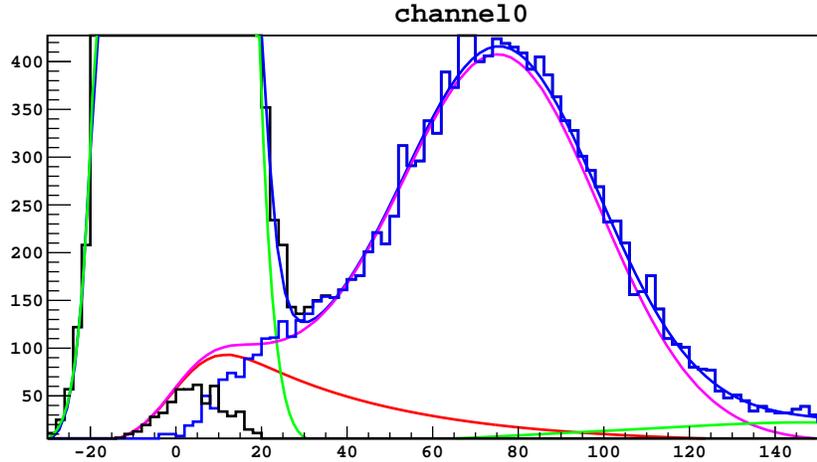


Figure 4.20: Laser Spectrum with interpolated events. The blue histogram contains events for which the BaselineFinder algorithm identified a signal region. The black histogram, which is the difference between the magenta single photoelectron fit and the blue histogram, represents the spectrum of single photoelectrons in scintillation events that could potentially be lost due to averaging over by the BaselineFinder.

line finder can be forced to not average within a specified laser window. This prevents the baseline finder from averaging away small SPE signals when they arrive when expected. This baseline freeze is turned on using the parameter `laserwindow_freeze` and the window is specified using the parameters `laserwindow_begin_time` and `laserwindow_end_time`.

The second possible problem is integrating over scintillation pulses. In particular, if the baseline finding algorithm does not identify and interpolate over small signals in the tails of scintillation pulses (such as lone single photoelectrons), the measured integral of the pulse will be diminished.

Figure 4.20 shows a standard photoelectron response fit with an additional histogram in blue. This histogram includes only events that the baseline finder identified as crossing the threshold, i.e., signal regions that are not averaged over. The small difference between the blue histogram and the magenta line near zero shows the single photoelectrons that *would* be averaged over and included in the baseline, and so their contribution to a pulse integral could potentially be diminished or lost.

In this plot, there are 472 potentially diminished SPEs (as determined by subtracting the blue histogram from the magenta fit, shown in the black histogram), with an average integral of 7.468. This is out of 13921 total SPEs. Thus, the potential for error from these diminished SPEs is at most 0.25 integrated counts per expected SPE, on a base of 68.5 integrated counts for the typical SPE, for a percent error of 0.36%. Note also that all of the difference between the SPE fit line and the BaselineFinder-flagged events falls under the integration pedestal, so the signal-to-noise ratio for these events is quite small, and therefore their contribution to a given pulse would be negligible under any circumstances.

Therefore, we can conclude that the potential systematic error due to the moving baseline averaging over small signals is negligible.

Scintillation Pulses

Without a Monte Carlo simulation of scintillation pulses or some absolute light yield reference, it is difficult to determine absolutely if the moving baseline introduces any biases in the determination of the charge for a given scintillation pulse. As a way to estimate this, we can compare the results of integration using our standard moving baseline algorithm versus a fixed baseline algorithm that simply finds the average of all samples in the pre-trigger window. Because the functioning of the PulseFinder depends on the moving baseline algorithm (in particular, to smooth out microsecond-timescale oscillations that otherwise resemble small S2 pulses), comparing flagged pulses between the two baseline modes is difficult at best. Instead, we choose to compare the integral of the waveform for each event in the fixed window from -0.05 to $7 \mu\text{s}$.

Figure 4.21 plots the integral of the sum channel in this window with a fixed baseline versus a moving baseline. Only events for which the moving baseline entry passed all-two pulse cuts are drawn. The figure also shows a profile histogram of the distribution, and a linear fit. As the figure shows, the fixed baseline appears to have an average offset of approximately 10 photoelectrons, as well as a systematic increase in the measured integral of

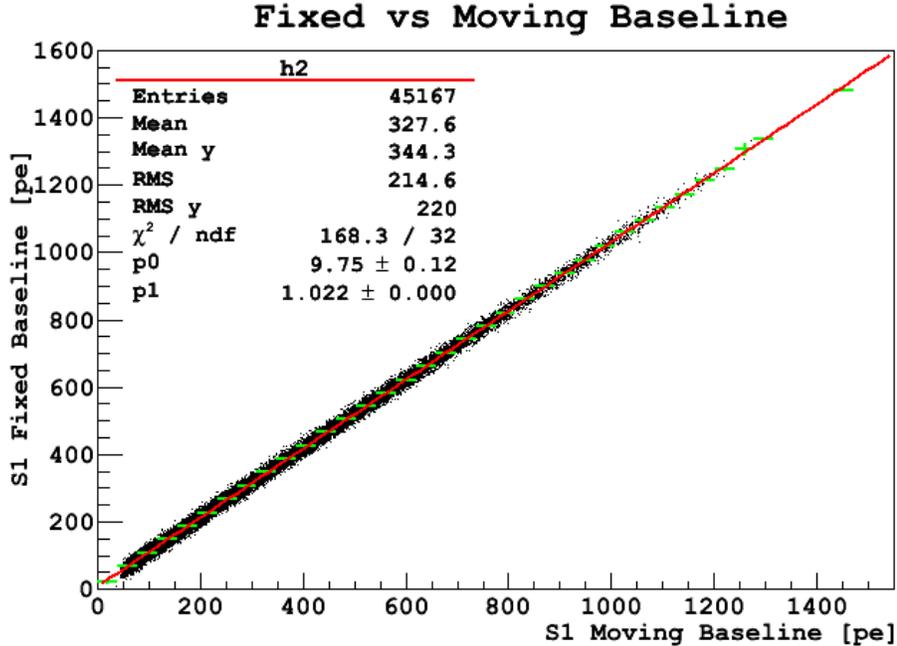


Figure 4.21: Integral with fixed baseline vs moving baseline, with linear fit.

2.2%. Therefore, the fixed baseline integral is larger by 3.2% at 1000 pe (roughly 500 keV_{ee} as shown in section 4.5.1), 7% at 200 pe, and 22% at 50 pe, near threshold. Since neither approach is inherently “correct,” we can only treat the difference as a systematic uncertainty. This effect is one of a largest contributors to energy uncertainty, especially at low energies.

4.5 High Level Analysis

4.5.1 Light Yield

Zero Electric Field

Figure 4.22 shows the S1 spectrum from the zero-field ²²Na runs combined (1439 and 1456), with successive cuts applied. Although there are no explicit cuts taken in the top (green) plot, events for which S1 is not definable (in particular, those for which no suitable baseline was found or for which no pulses were identified) are implicitly removed. As the plot shows, events begin to regularly saturate the vertical scale of the digitizer (almost always for the 8”

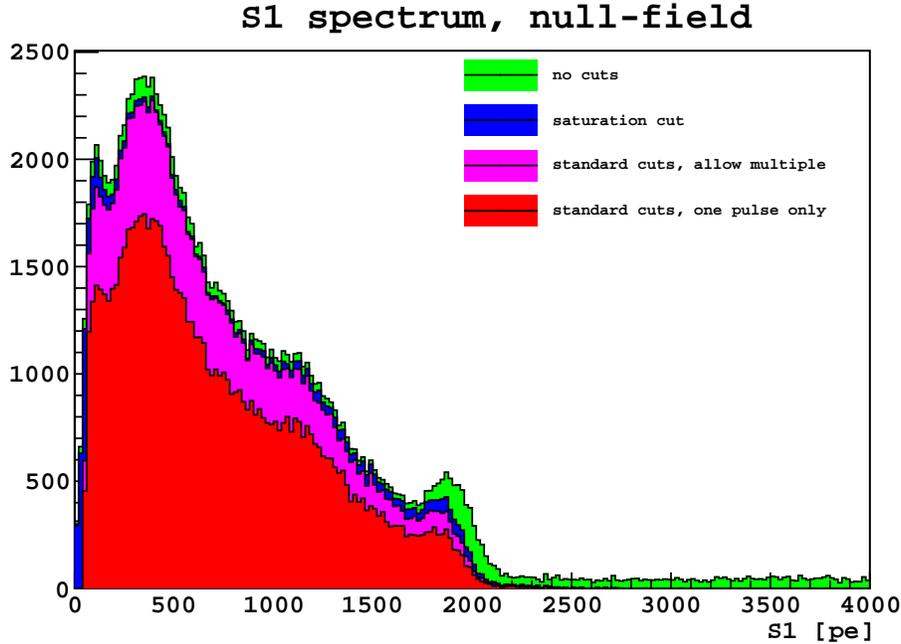


Figure 4.22: S1 spectrum of ^{22}Na events with zero field, showing the effect of additional cuts on the spectral shape. The vertical saturation cut unfortunately removes a large fraction of the full-energy 511 keV peak. Other cuts do not significantly impact the spectral shape.

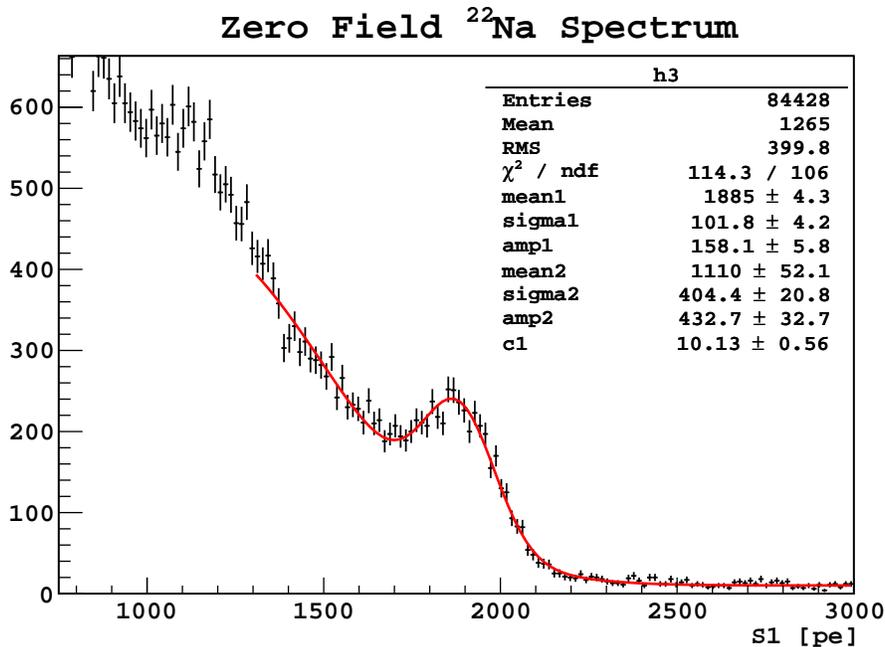


Figure 4.23: Fit of the ^{22}Na S1 spectrum for zero-field runs to the sum of two gaussians plus a constant, over the range 1300-3000 photoelectrons. Events in this spectrum pass all the usual single-pulse cuts except for vertical saturation, as that removes the high-energy end of the full-energy peak.

Fit start (p.e.)	511 keV mean	511 keV sigma	χ^2/NDF
1200	1880 ± 4	105 ± 4	128/113
1250	1881 ± 4	104 ± 4	124/110
1300	1885 ± 4	101 ± 4	114/106
1350	1889 ± 4	98 ± 4	101/103
1400	1892 ± 4	97 ± 4	91/100
1450	1893 ± 4	96 ± 4	89/96
1500	1893 ± 5	96 ± 5	89/93
1550	1893 ± 7	96 ± 6	85/90
1600	1886 ± 4	101 ± 4	82/86
Average	1888	99	
Std. Dev.	5.2	3.5	

Table 4.7: Results from fitting the sum of two gaussians plus a constant to the S1 spectrum of null-field ^{22}Na events as a function of the low-energy start point of the fit. The high end of each fit was 3000 photoelectrons.

PMT) around 1800 photoelectrons, which, unfortunately, is near the full-energy peak of the 511 keV annihilation gammas from the ^{22}Na source. This makes determination of the exact amplitude of the full-absorption response difficult. The saturated events that occur at lower energies are likely very close to the 8" PMT, causing a disproportionate amount of light to be collected by it and a correspondingly high signal. Some of these events may also contain a secondary (pileup) pulse that is saturated.

However, when events are barely saturated, in general only a few samples at the peak of each pulse are beyond the digitizer range, and so the amount of signal lost is largely negligible compared to the ~ 2000 samples over which S1 is integrated. Therefore, for the purposes of determining light yield, we fit the spectrum of events that pass all the single-pulse cuts described in section 4.4.2 except the saturation cut. As a model for the spectrum near the full energy peak, we use the sum of two gaussians (one for the full energy peak itself and one for the portion of the spectrum below the full-energy peak) plus a constant term for background plus the contribution of the 1275 keV line.

Table 4.7 reports the results of this fit as a function of the low-energy starting point of the fit range (the high end was fixed to 3000 photoelectrons.) Below about 1200 p.e., the Compton spectrum is the dominant feature in the spectrum and would need to be modeled to

be fit reliably. Figure 4.23 shows the resulting fit for a low endpoint of 1300 photoelectrons. Taking the weighted average of the results in the table puts the full-energy 511 keV peak at approximately 1887 ± 5 photoelectrons, with a width of 100 ± 5 photoelectrons. This in turn gives a zero-field light yield for 4 top plus 1 bottom PMT of approximately 3.69 ± 0.01 p.e./keV_{ee}.

In order to attempt to evaluate the maximum light yield of the detector in this campaign, that is, the light yield with seven 3" and one 8" PMTs, a sampling of nine zero-field runs was analyzed including all of the PMTs. Table 4.8 summarizes the results of fitting the function described previously to the 511 keV ²²Na peak over the range 2000-2900 p.e. for each of these runs, with no explicit cuts taken on the data, to minimize bias introduced by the saturation cut. It should be noted that the last three runs in the table (1392, 1439, and 1456) were taken during the period when channel 6 was beginning to fail. The single photoelectron calibration from laser runs largely corrects for the apparent drop in gain of the PMT, but these runs should be considered less trustworthy than the others in the table. Figure 4.24 shows a fit to the total spectrum summing over all the runs in Table 4.8, with and without any explicit cuts applied. (Even with no cuts applied explicitly, some triggers, for example those for which no baseline or no S1 pulse could be found, are removed due to the fact that the event energy is undefined.) The light yield obtained from the spectrum without cuts is 4.50 ± 0.01 (statistical only) pe/keV_{ee}, in good agreement with the average of the individual fits of 4.51 ± 0.03 pe/keV_{ee}. The fit to the spectrum with cuts returns a value within 1% of the spectrum without cuts, but the χ^2 value is much worse, likely due to distortions from the cuts (most notably the saturation cut).

To cross-check the higher-energy 511 keV ²²Na peak, Figure 4.25 shows the spectrum of ¹³³Ba, before and after background subtraction. The barium data come from run 1257, and the background from the immediately previous run, 1256, normalized to the same livetime. (Livetime is calculated by scaling the acquisition time of the run by the ratio of triggers saved to triggers pulses generated by the coincidence units.) Both runs were taken without an

Run	Entries	Light yield [pe/keV]	χ^2/NDF
1176	138413	4.46 ± 0.006	24.0/23
1224	84746	4.47 ± 0.006	19.4/23
1255	494150	4.45 ± 0.006	26.5/23
1304	99177	4.59 ± 0.006	32.8/23
1323	12386	4.41 ± 0.013	63.4/23
1335	11131	4.78 ± 0.011	15.9/23
1392	99040	4.64 ± 0.006	52.3/23
1439	78364	4.72 ± 0.006	39.5/23
1456	41340	4.57 ± 0.010	25.9/23
Average		4.51 ± 0.03	

Table 4.8: Results of fitting the sum of two gaussians plus a constant offset to the 511 keV ^{22}Na peak over the range 2000-2900 photoelectrons to zero-field runs with all channels enabled in analysis. The average value reported is weighted by the entries in each run, and the reported uncertainty is the weighted RMS/\sqrt{N} .

applied field, and all PMTs are included in the sum. All standard one-pulse cuts were applied, and the fit to the same two gaussians plus constant offset used before was performed over the range 1300-2600 photoelectrons. The fit gives a mean response of 1593 ± 3 photoelectrons for the 356 keV ^{133}Ba full-energy peak, giving a light yield of 4.47 ± 0.01 pe/keV $_{ee}$, in good agreement with the ^{22}Na results. The fact that the constant term is fit to a negative value indicates that there is some problem with our livetime normalization, but the over-subtraction is small compared to the barium spectrum ($\sim 3\%$ of the height of the 356 keV peak).

As an estimate of the systematic uncertainty due to the question of whether to use the IoS or SoI approach to calculating the total light collected (see section 4.4.3), the same spectra from figures 4.23 and 4.25 were generated using the SoI definition for S1, using the same cuts, and fitting over the same regions. For the ^{22}Na spectrum of figure 4.23, the IoS (standard) S1 definition gave a fit result of 1885 ± 4.3 p.e. for the 511 keV peak, while the SoI approach gave 1894 ± 4.3 p.e., for a difference of 9 ± 6 p.e. or $\sim 0.5\%$. For the ^{133}Ba spectrum of figure 4.25, the IoS approach gave a fit value of 1593 ± 3.3 pe for the 356 keV peak, while the SoI approach gave 1618 ± 3.5 for the fit result, for a difference of 25 ± 5 pe, or 1.6%.

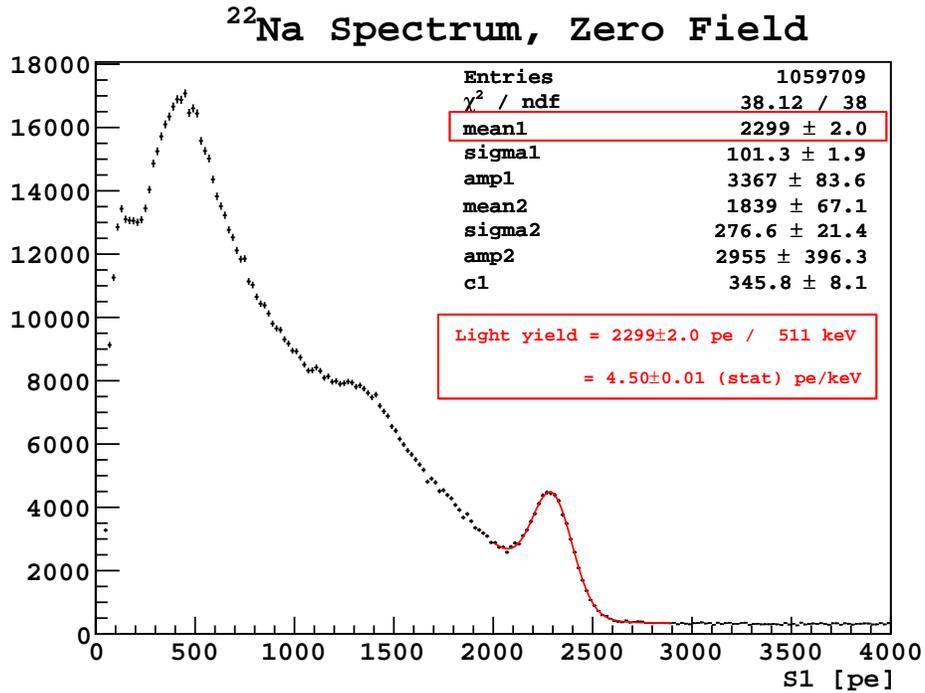
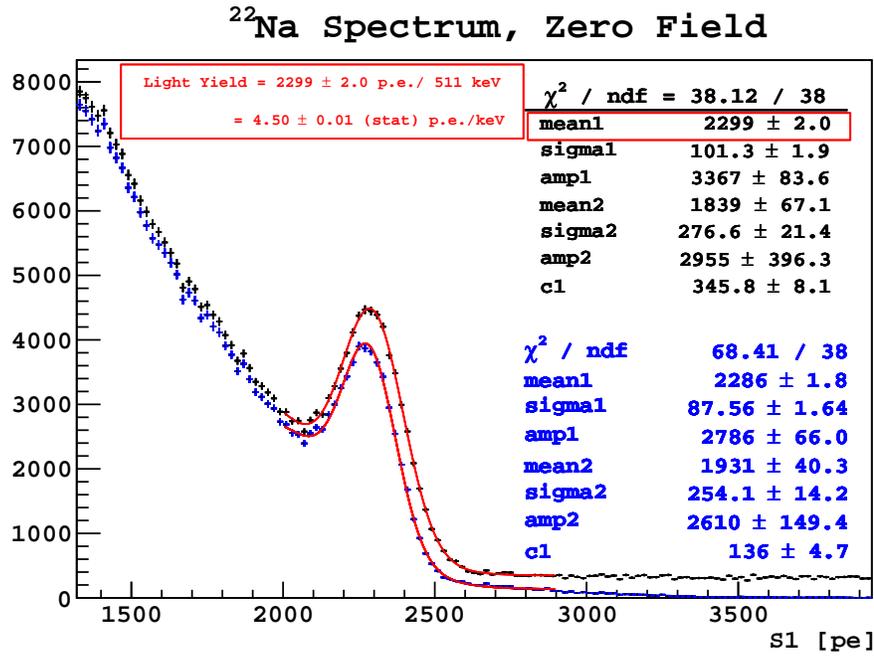


Figure 4.24: (Top): Fit to the 511 keV peak of the sum spectrum of 9 zero-field ^{22}Na runs with all tubes included, with (blue) and without (black) cuts. (Bottom): Fit showing whole spectrum, suitable for standalone reference in other documents.

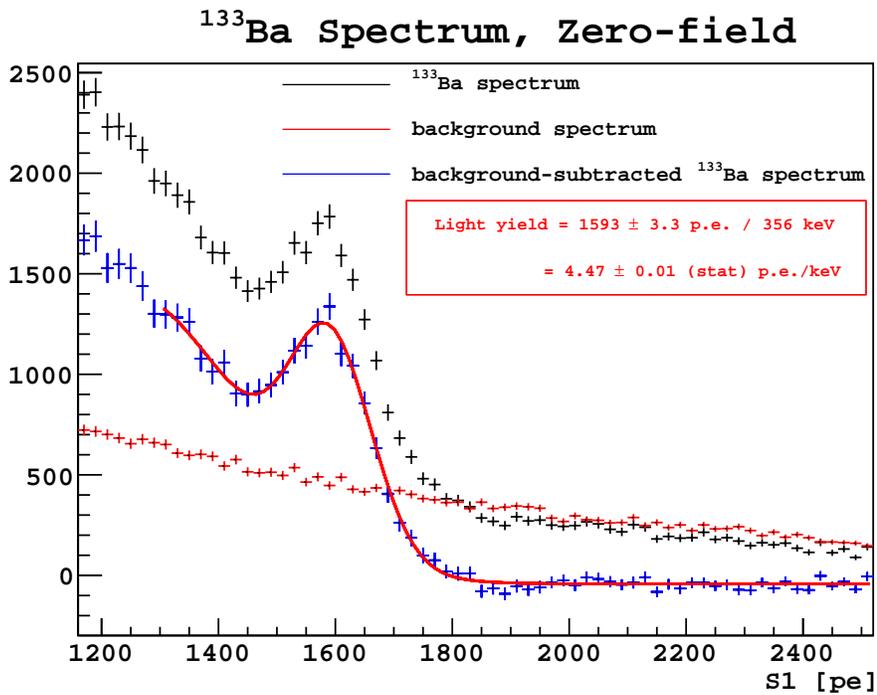
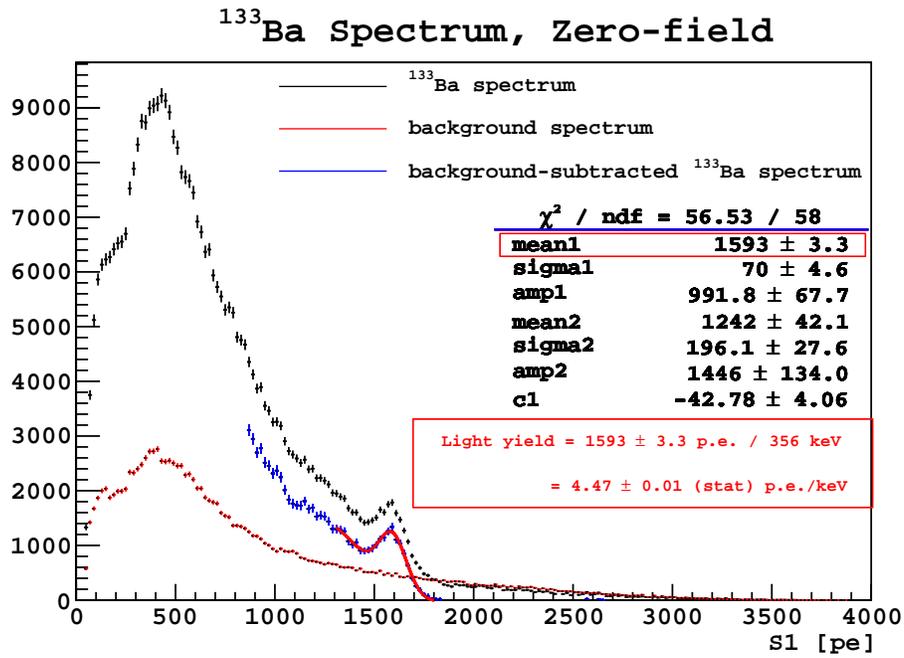


Figure 4.25: (Top): ^{133}Ba spectrum from run 1257 with all tubes included, using standard one-pulse cuts, before and after background subtraction. The fit function is the sum of two gaussians plus a constant, and the spectrum is fit over the range 1300-2600 photoelectrons. (Bottom): Same as top, but zoomed around the fit region, showing the over-subtraction of background above the peak.

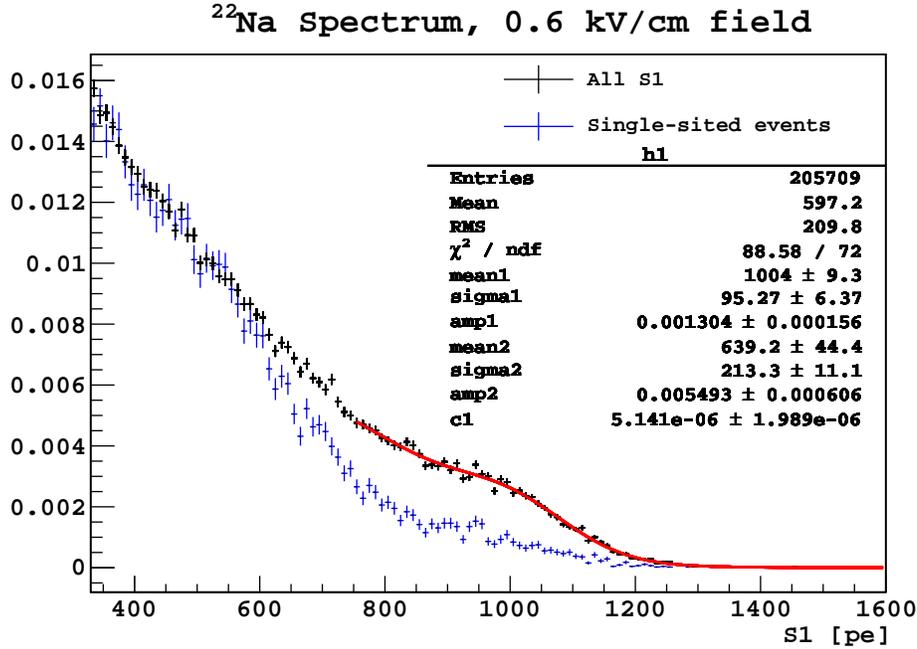


Figure 4.26: S1 spectrum of ²²Na events in the presence of a 0.6 keV/cm electric field. Entries in the black plot pass all single-pulse cuts, while those in the blue plot are well-defined single-sited events (two-pulse cuts). Both spectra are normalized to unit integral. The fit function is the sum of two Gaussians plus a constant.

With Electric Field

Figure 4.26 shows the S1 spectrum for ²²Na events from the run list with an electric field of 0.6 kV/cm present. The black spectrum is from those events that pass cuts to ensure a valid S1 signal (but not, in particular, any requirement on the number of pulses), while the blue spectrum is from events which pass two-pulse cuts, and so have exactly one well-defined S1 and one S2 pulse. The maximum energy that a gamma can deposit in a single scatter is backward scattering off an electron, which produces the Compton shoulder in a gamma spectrum at $E = 2E_\gamma^2 / (m_e + 2E_\gamma)$. To deposit more energy than the Compton scattering limit, a gamma must scatter multiple times in the argon. Therefore, we expect events above the Compton edge to be suppressed by removing multiple-sited (more than one S2) events, and this is shown to be the case in Figure 4.26.

Run	Height [cm]	511 keV fit, without 0,3,6 [pe]	511 keV fit, with 0,3,6 [pe]
1364	8	1758 ± 4	2279 ± 4
1365	4	1818 ± 3	2335 ± 3
1366	0	1867 ± 2	2360 ± 2
1367	-4	1899 ± 2	2361 ± 2
1368	-8	1918 ± 3	2349 ± 3

Table 4.9: Fit result for the location of the 511 keV peak in photoelectrons from collimated ^{22}Na runs with different heights relative to the center of the detector active volume. Spectra omitting channels 0, 3, and 6 were fit over the range 1300-3000 photoelectrons, while spectra including those channels were fit over the range 2000-2900 photoelectrons.

Using the same fit as in the previous section to the spectrum allowing multiple-sited events over the range 750 to 1600 p.e., the response to the 511 keV gamma is 1004 ± 9 photoelectrons, with a width of 95 ± 6 p.e (all uncertainties are statistical only). The light yield is therefore 1.96 ± 0.02 p.e./keV $_{ee}$, which gives a quenching factor of 0.530 ± 0.005 in the presence of an applied electric field of 0.6 kV/cm. The fit results are extremely sensitive to the range of the fit, giving a systematic uncertainty to the mean response of approximately 40 p.e. or 4%.

Dependence on height

Runs 1364-1368 were taken using the ^{22}Na source mostly enclosed using lead bricks so that it should produce a collimated beam with a vertical spread of a few centimeters. The location of the collimated source was scanned vertically across the active volume of the detector in 4 cm steps. Using these runs, we can attempt to evaluate the light yield in the detector as a function of the vertical position of the interaction. Table 4.9 lists the results of fitting the spectrum obtained for each run to the standard two gaussians plus a constant function. Results are given both for the case where the problematic channels (0,3, and 6) are skipped, in which case the spectrum was fit to the range 1300-3000 photoelectrons, and when all PMTs are included, where the spectrum was fit over the range 1800-4000 photoelectrons. Figure 4.27 graphs these same results, where all of the yields are scaled so that the center of the detector is 1. From the figure it is clear that there is a strong dependence of the light

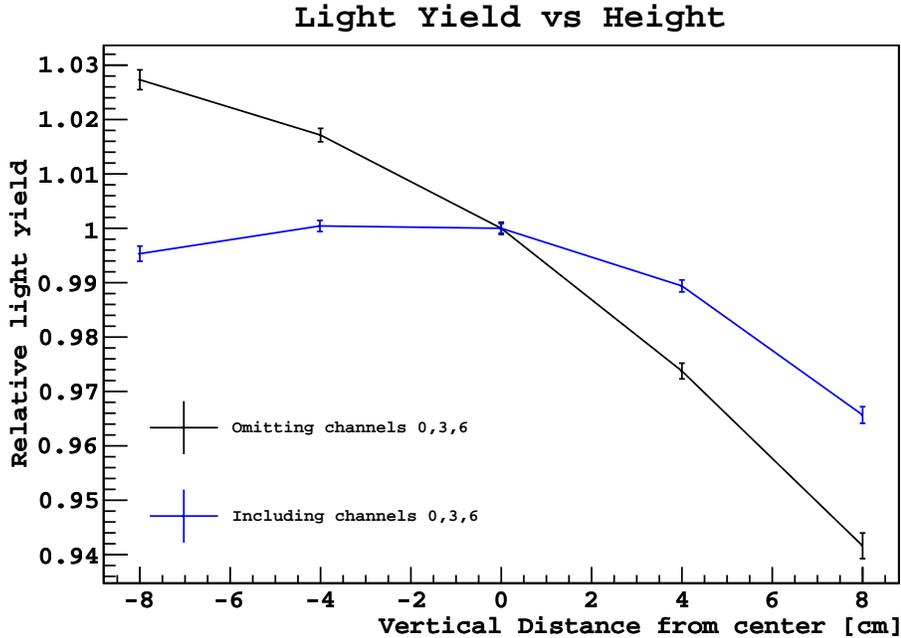


Figure 4.27: Relative light yield for collimated ^{22}Na runs as a function of the height of the collimated beam relative to the center of the detector. The effect is shown both with (blue) and without (black) the three problematic tubes (0, 3, and 6). Amplitudes are normalized to 1 at the center of the detector.

yield on the source height when not all of the top PMTs are used in the analysis, which is reduced but not eliminated by including all channels. The effect is especially pronounced near the top of the detector. The maximum deviation (from top to bottom of the detector) when the three bad channels are omitted is $\sim 8.5\%$; the max deviation including all channels is $\sim 3.5\%$.

Further Discussion

The expected luminescence quenching at 0.6 kV/cm is approximately 0.58 [105], compared to our measured value of 0.53. This indicates that we experience some 9-10% more light loss than expected when the field is turned on. Figure 4.28 shows the fraction of the total signal that comes from the bottom PMT for zero-field and 0.6 kV/cm field runs, and figure 4.29 plots the same quantity as a function of the drift time for with-field runs. The large difference in the fraction of light collected in the bottom tube with and without the presence of the

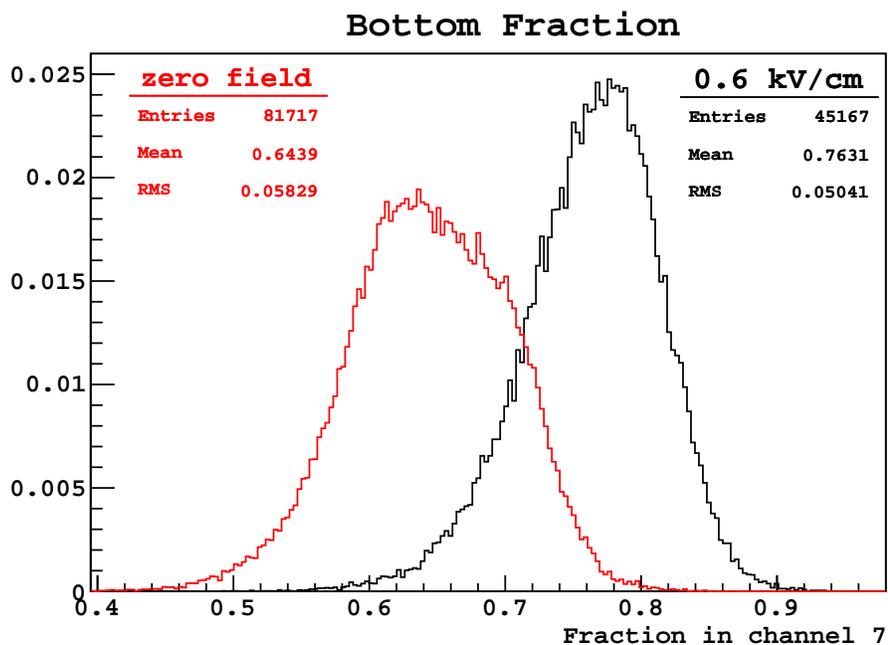


Figure 4.28: Fraction of S1 light collected by the bottom (8") PMT for ^{22}Na events, for runs with (black) and without (red) a drift field.

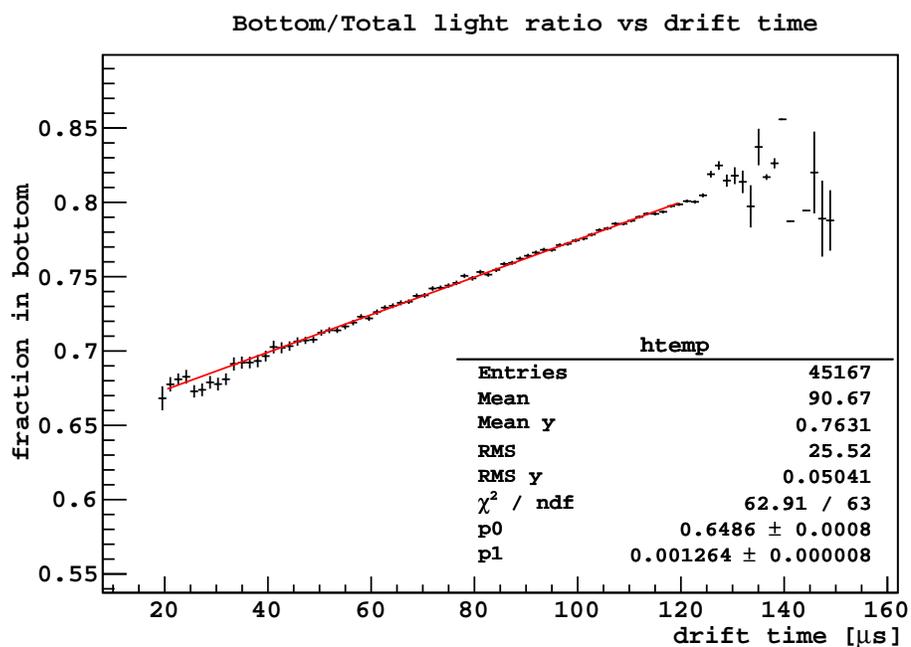


Figure 4.29: Fraction of light collected in the bottom (8") PMT (channel 7) as a function of drift time for ^{22}Na runs with drift field of 0.6 kV/cm.

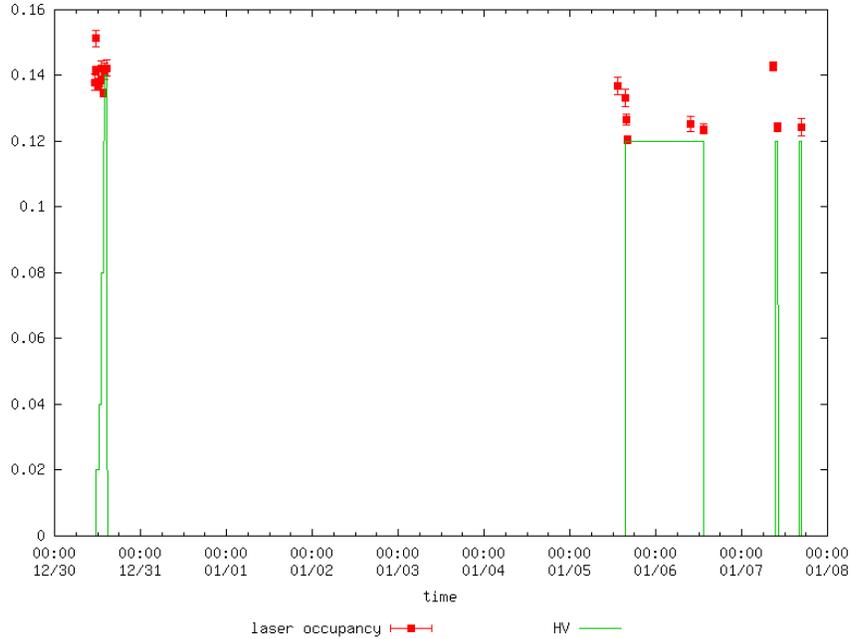


Figure 4.30: The laser occupancy (expected number of photons per trigger) for channel 1 versus time, showing also the measured current on one of the 50 kV HHV supplies (arbitrarily scaled) to show when the HHV fields are on. On 12/30, the extraction field was set at 3.0 kV/cm, and the drift field was ramped through several values. For the January runs, the fields were at (3.85,0.6) kV/cm. There seems to be a drop in occupancy when the fields are present, which seems to depend primarily on the extraction field and very little on the drift field.

drift field could be in part due to an effect of our selection criteria, which is biased towards events with longer drift time, i.e., toward the bottom of the detector (see section 4.5.3). However, this does not seem to be enough to account for the mean shift by itself.

Another likely explanation is that even the “good” channels experience a drop in efficiency similar to channels 0 and 3 due to the very high light levels present when the drift and extraction fields are on. Figure 4.30 shows the photon occupancy of channel 1 during laser pulsing as obtained from the single photoelectron fit and shows the recorded current in one of the 50 kV HHV power supplies as a way to indicate whether the drift and multiplication fields were enabled. Although not conclusive, there does seem to be evidence of a decrease in occupancy (which should be directly proportional to PMT efficiency) when the fields are turned on; the other top PMTs look similar. If the change in bottom/total ratio is due to a

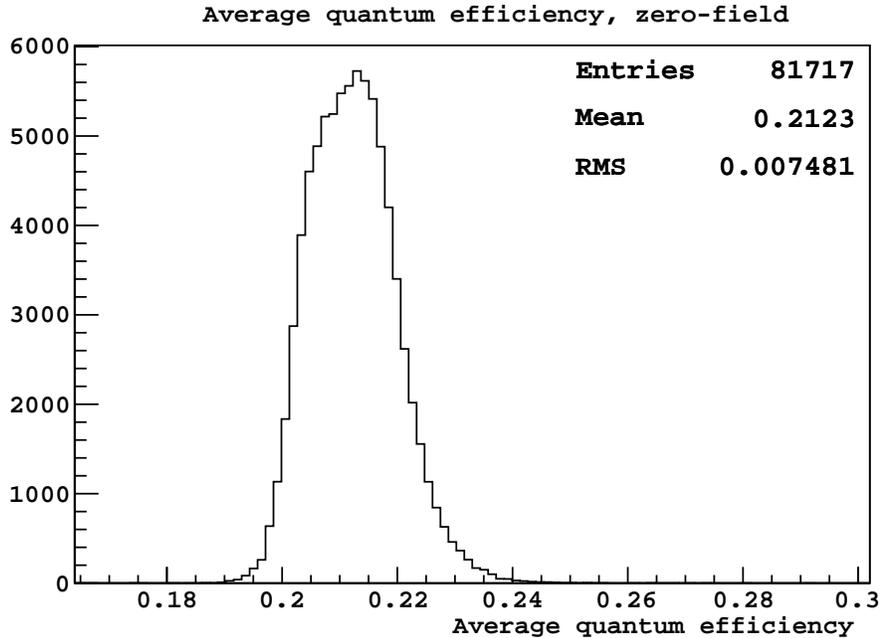


Figure 4.31: Average quantum efficiency calculated for S1 pulses with zero field. Average QE is calculated as $S1 \cdot (S1_{bot}/0.175 + S1_{top}/0.344)^{-1}$.

drop in efficiency of the top PMTs, the average efficiency drop can be calculated from the shift in ratio by $\eta = (1 - r_{field})/(1 - r_{null})$, which evaluates to 0.67 for $r_{field} = 0.76$ and $r_{null} = 0.64$. This would result in an overall loss of light of $(1 - \eta)(1 - r_{null}) = 12\%$. Since this is larger than the measured loss, it is reasonable to assume that the shift in observed bottom/total ratio is due to a combination of a drop in top PMT efficiency and selection effects.

We can also ask how much the difference in collection and efficiency between top and bottom affects the energy resolution of the detector. Figure 4.31 plots the average or effective quantum efficiency for events passing cuts in the zero-field runs. The total number of photons striking the photocathodes is approximated as the fraction of photoelectrons detected in the bottom PMT divided by its quantum efficiency (17.5%) plus the fraction detected in all of the top PMTs divided by their average quantum efficiency (34.4%). The effective quantum efficiency is then the total number of measured photoelectrons in S1 divided by the number of photons incident on the cathodes. Note that this ignores effects like differing collection

efficiencies or position-dependent light collection. The resulting distribution has a mean of 0.212 with a standard deviation of 0.0075, and is asymmetrical, with a longer tail towards higher values. This suggests that the differences in quantum efficiency between the top and bottom of the detector alone result in a 3.5% spread in the resolution of the energy response measurement, again ignoring effects such as photoelectron collection efficiency (which is expected to be worse for the 8" tube) and light collection issues.

This 3.5% spread is in reasonable agreement with the width of the 511 keV peak discussed in section 4.5.1 after photoelectron counting statistics are accounted for. The RMS of the peak was fit to be ~ 100 photoelectrons, compared to a mean of ~ 1890 photoelectrons, or 5.3%. From Poisson statistics of photoelectron counting, we expect a spread of \sqrt{N} photoelectrons, and from the spread of the single photoelectron in a PMT, we expect the N th photoelectron peak to have a width of approximately $\sigma_{1pe}\sqrt{N}$ (this approximation is only valid when the response to each single photoelectron in a pulse is independent). Therefore, the total spread from photo-statistics plus PMT response is $k\sqrt{N}$, where $k = \sqrt{1 + \sigma_{1pe}^2}$. This does not include electronics response. The average width of the single photoelectron is approximately 0.46 p.e. for the 3" PMTs and 0.58 p.e. for the 8" PMT, leading to values of k ranging from 1.10 to 1.15 (depending on whether the event is collected primarily in the top or bottom of the detector). So the expected width of this response at 1890 photoelectrons is approximately 48-50 photoelectrons, or 2.5-2.6%, leaving a further 87 p.e. or 4.6% spread contribution from other sources including the top/bottom collection efficiencies (assuming that the uncertainties are independent and add in quadrature).

Alternatively, one can note that the variation in light yield with interaction height is due to the combination of light collection efficiency and average quantum efficiency. From Figure 4.27, the spread in light yield is a few percent; again in good agreement with the average quantum efficiency spread.

Part of the reason for the 8" PMT to account for a large fraction of the signal is that light can be totally internally reflected at the liquid/gas interface, so light that might have hit the

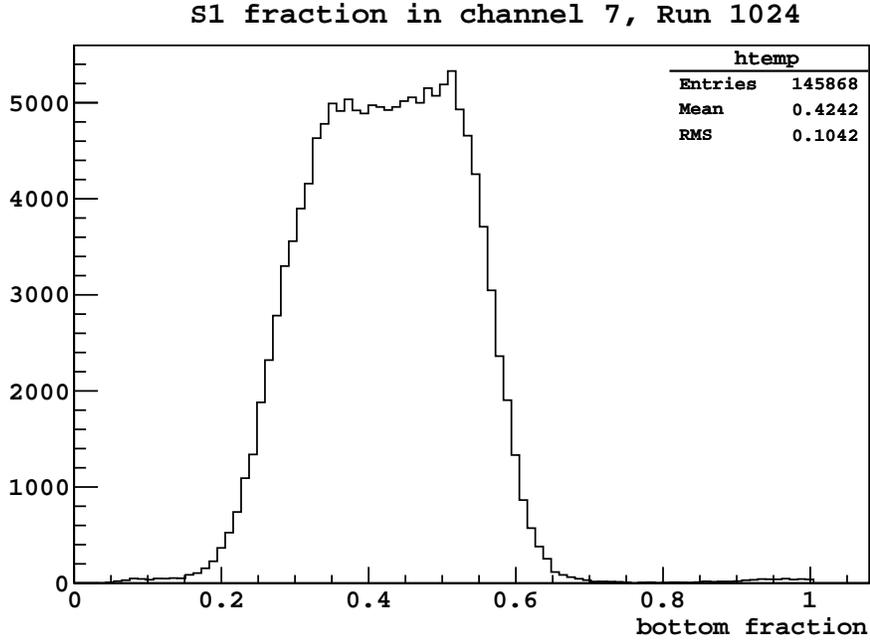


Figure 4.32: Fraction of S1 light collected by the 8” PMT in Run 1024, a zero-field background run without a gas pocket.

top PMTs is instead sent to the bottom of the detector. The 8” PMT has a total surface area roughly equal to all seven of the 3” PMTs, but a significantly lower quantum efficiency, so, when other conditions are equal, we expect the 8” PMT to account for a smaller fraction of the signal than the sum of all of the top PMTs. Figure 4.32 shows the fraction of S1 signal measured in channel 7 with all PMTs enabled for run 1024, which was a background run with no fields and no gas pocket. Here we see that the 8” PMT accounts for 42% of the S1 signal on average, which is roughly within expectations. For the future run of DarkSide-10 at LNGS, we will replace the single 8” PMT on the bottom with an array of seven 3” PMTs identical to the top arrangement. Since there are no effects like the gas pocket to differentiate top and bottom in Run 1024 other than the PMTs, we can estimate that the light yield will increase when we move to the new configuration by a factor $2(1 - \overline{f_{bot}})$, or roughly 15%. So if all other factors remain identical, we can expect the null-field light yield in the next campaign to increase to around 5.2 p.e./keV, and probably higher, since the present light yield estimate of 4.5 p.e./keV was obtained with the gas pocket in place.

4.5.2 F90

The primary discrimination parameter for argon detectors is the ratio of scintillation light emitted from singlet de-excitations (7 ns lifetime) to the total amount of light collected (singlet plus a triplet state with $\sim 1.6 \mu\text{s}$ lifetime), denoted by F_{singlet} . As an easier-to-calculate approximation to F_{singlet} , we typically use the variable $F90$, which is defined to be the ratio of scintillation light collected within the first 90 ns of a pulse to the total amount collected.

Figure 4.33 shows the distribution of $F90$ for the three sets of runs considered in this analysis, with a cut of $S1 < 200$ p.e. applied in addition to the standard cuts. The additional cut is to restrict the plot to approximately the energy regions favorable for WIMP searches; at 0.6 kV/cm, 200 p.e. corresponds to $\sim 200 \text{ keV}_{\text{nr}}$. Each of the histograms is normalized to have unit integral. From this plot, if we define a cut on $F90$ to reject beta-like events that accepts approximately 50% of nuclear recoils, $\sim 10^{-4}$ of the original ^{22}Na population remains for this dataset. Although this is likely limited due to the high background rate at surface, further conclusions about the background rejection cannot be drawn from this dataset.

It's difficult to directly compare runs with and without field due to the differing light yields, but the offset visible between those two histograms in Figure 4.33 does not appear to be solely an effect of comparing differing energies. Figure 4.34 is a scatterplot of the two distributions, showing both vs $S1$ as measured in photoelectrons and $S1$ measured in energy ($S1$ divided by the respective light yield, 3.69 for runs without field and 1.96 for runs with field). The energy scaling ignores any possible non-linear quenching effects in the light yield. As the figure shows, the difference in $F90$ distribution is persistent across a wide range of $S1$, whether comparing deposited energy or collected photoelectrons. Figure 4.35 plots the profiles of each distribution as a function of energy, and fits each profile to a function consisting of two exponentials plus a constant offset, which seems to describe the distribution relatively well. The sharp drop-off in the mean of the $F90$ distribution in the

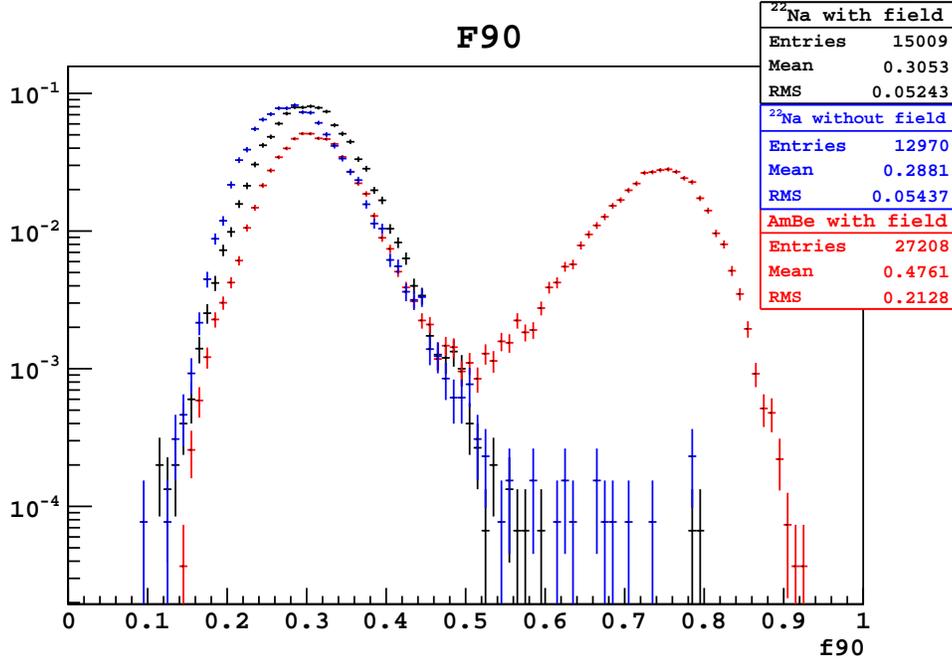


Figure 4.33: $F90$ distribution for the three groups of runs considered in this analysis: ^{22}Na runs with field (black), ^{22}Na runs without field (blue) and AmBe runs with field (red). Only events with $S1 < 200$ p.e. are included.

zero-field plot near 500 keV is a bias introduced by requiring that events not vertically saturate the 2 V range of the digitizer, which preferentially passes events with low $F90$ when near the saturation point. To attempt to avoid this bias, the profile is fit only up to 450 keV.

The “turn-up” below about 60 keV in Figure 4.35 has also been observed by Lippincott et. al. [75]. Figure 4.36 shows this same plot for the null-field runs zoomed into the low energy region below 35 keV_{ee}, including measurements their measurements for comparison. There is fairly good agreement between the two plots, but not perfect. Two likely causes for the discrepancy could be quenching of the triplet lifetime due to impurities such as nitrogen, or improper calibration of the energy scale. This feature of the $F90$ distribution is likely due to the increasing ionization density for electrons of that energy. Figure 4.37 plots the stopping power (or dE/dx) for electrons in argon, obtained from the NIST stopping power and range tables database [10], which shows a steep climb in stopping power below ~ 100 keV. Figure 4.38 shows $F90$ as a function of stopping power for runs with and without field (again

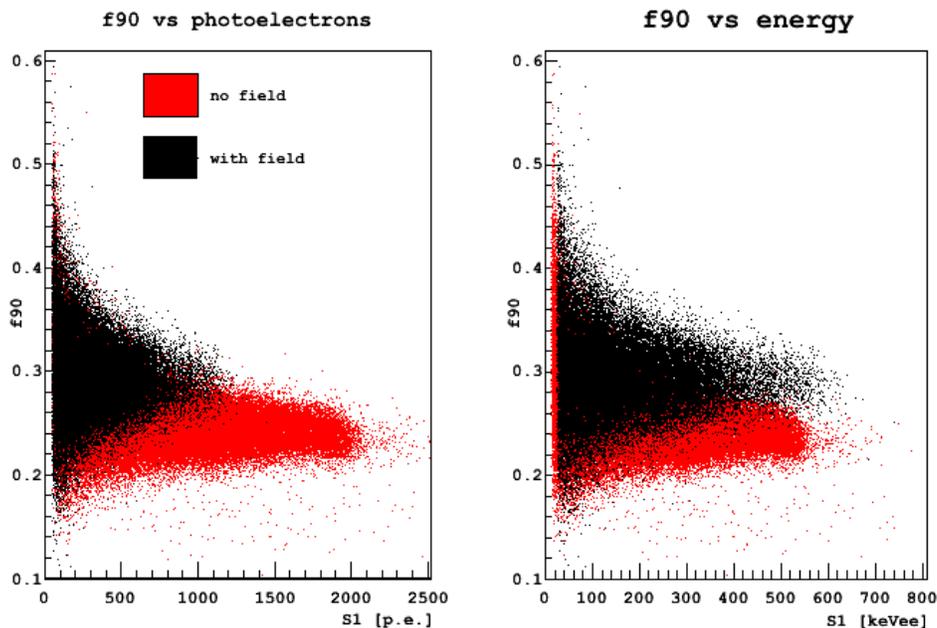


Figure 4.34: $F90$ distribution of events as a function of $S1$, where $S1$ is measured either in photoelectrons (left) or energy (right), for ^{22}Na runs with (black) and without (red) field.

assuming that the amount of light collected during a pulse is linear with energy down to the 50 keV threshold). Although both distributions appear roughly linear, the parameters are very different; a linear fit to both distributions is shown, despite being a poor fit overall, to give an estimation of these parameters.

Because of the change in performance of the photomultipliers due to the high light levels when the $S2$ signals are present, it is difficult to determine whether the observed shift in $F90$ in the presence of the electric fields is a physical or electronics effect. In order to verify that it is not an effect of analysis cuts, Figure 4.39 shows the distribution of “ F_{prompt} ” calculated as the ratio of the integral of ROI 2 to ROI 1 (see Section 4.3.4), with no explicit data cuts: the distributions are still quite different. One more thorough check which can be performed with the present data is to examine $F90$ as measured by individual PMTs. This can help to tell us, for instance, if the shift in $F90$ is directly caused by the increased fraction of light collected by the 8” PMT noted in section 4.5.1.

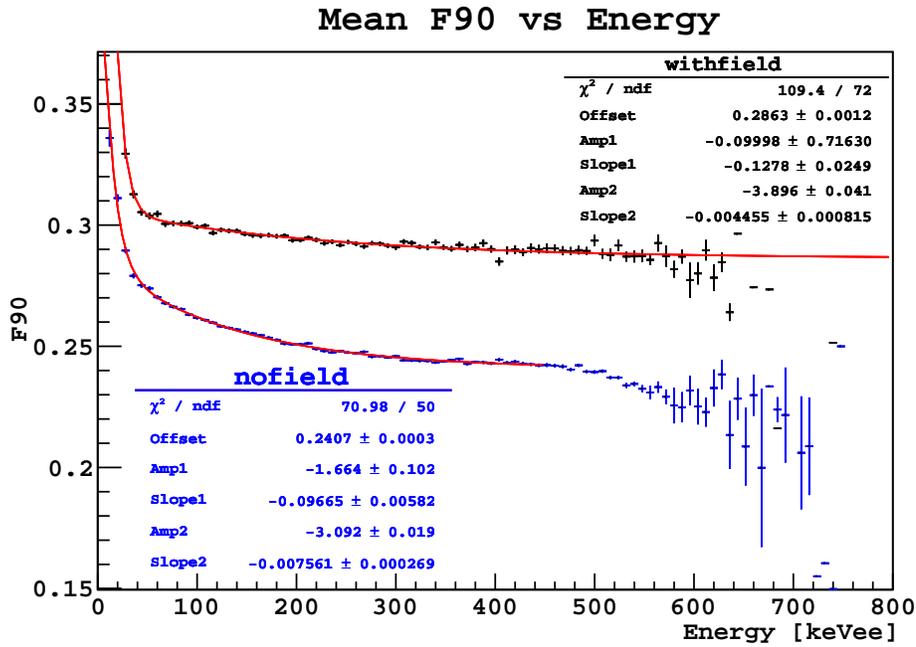


Figure 4.35: Profile of the mean of $F90$ as a function of deposited energy, for ^{22}Na runs with (black) and without (blue) field present. Each distribution is fit to a function consisting of two exponentials plus a constant offset.

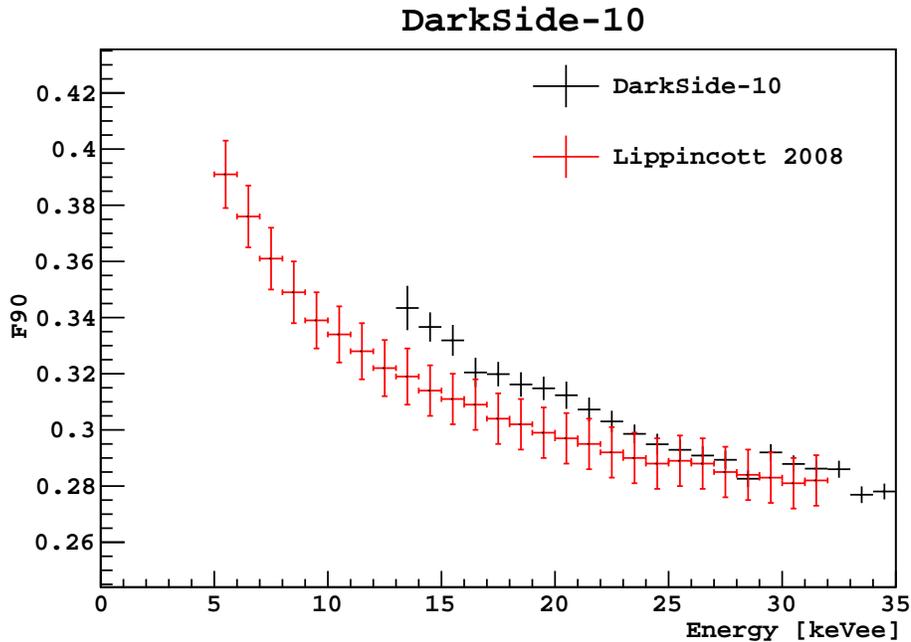


Figure 4.36: Profile of the mean of $F90$ as a function of energy for null-field runs, zoomed into the low energy region in order to compare to the published measurement from Lippincott et. al. [75].

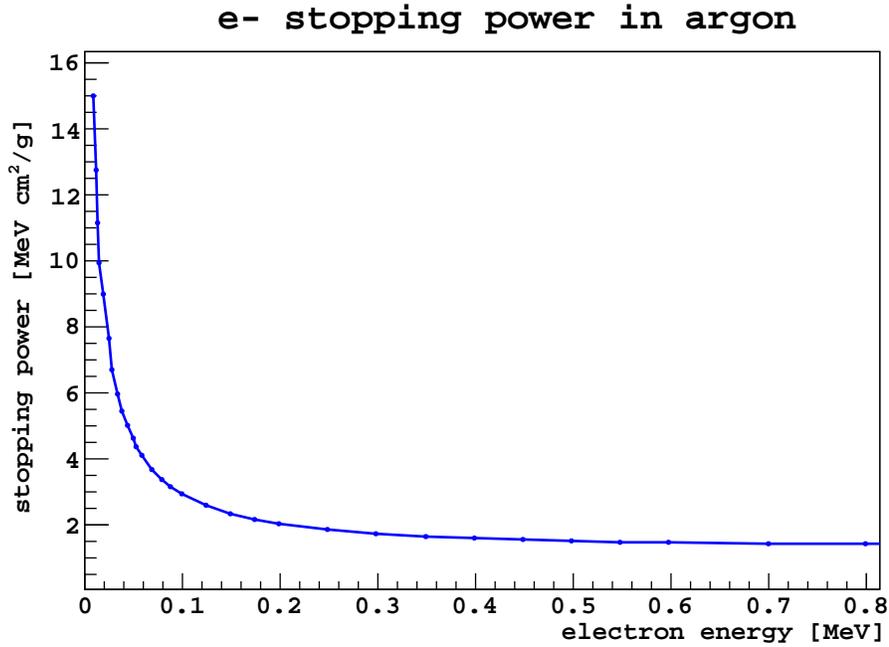


Figure 4.37: Electron stopping power in argon, obtained from NIST stopping power database [10].

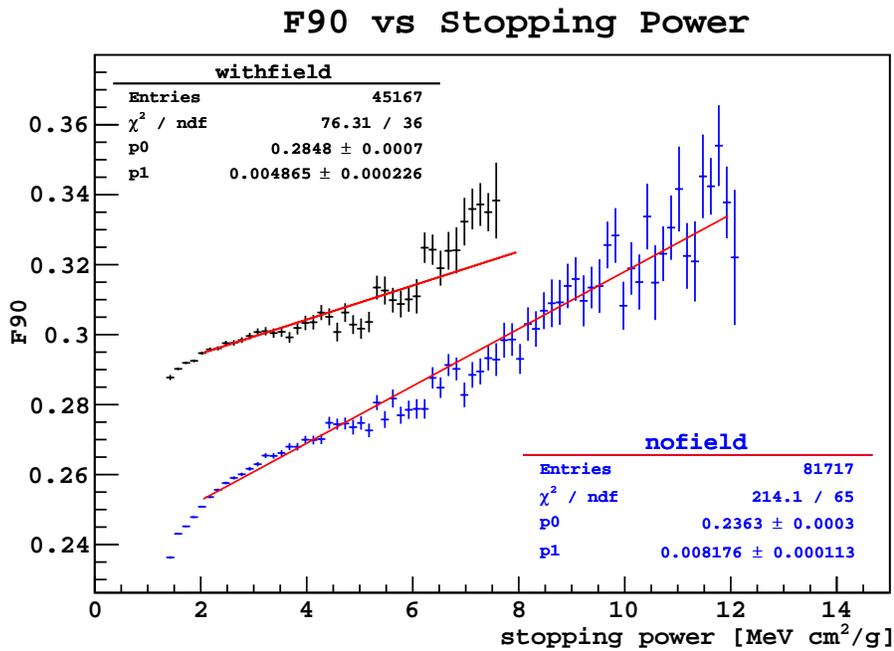


Figure 4.38: F_{90} as a function of electron stopping power, for runs with (black) and without (blue) electric field present.

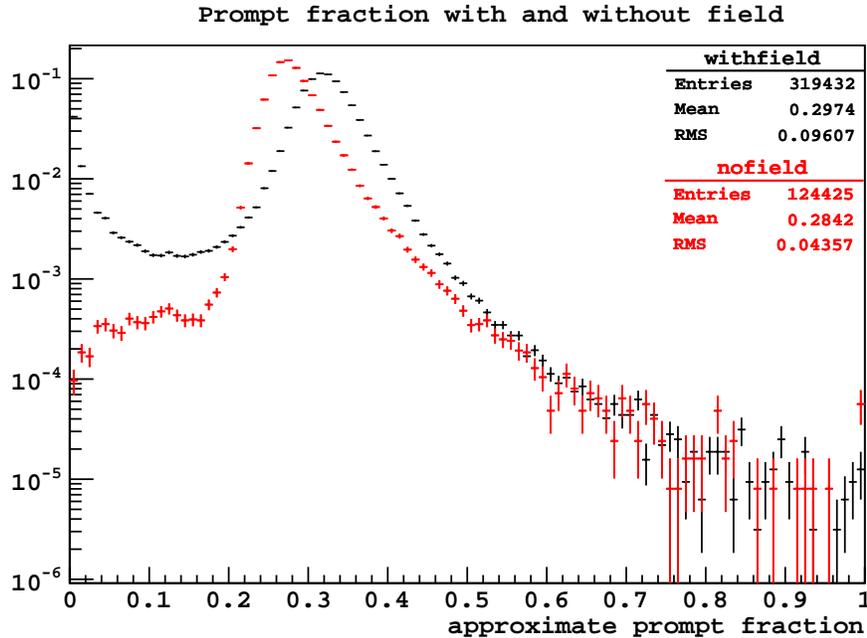


Figure 4.39: Distribution of the ratio of the integral over ROI 2 to ROI 1 (as an estimate for F_{prompt}) for runs with (black) and without (red) applied drift field, with no explicit cuts applied. The difference between the two distributions indicates that the difference in distributions in Figure 4.34 are not due to analysis cuts.

Figure 4.40 shows $F90$ as a function of S1 in photoelectrons for all channels as well as channels 7 (the 8" PMT) and 1 (one of the top 3" PMTs) individually, for runs with and without electric fields present. As the figure shows, $F90$ as measured by channel 7 is generally smaller than the value for all channels combined, while the value for channel 1 (and, in fact, all of the top PMTs) is larger. This is likely due to the ~ 10 -20 ns offset remaining in the signal for channel 7, which is visible in Figure 4.7. (A quick check offsetting channel 7 by 50 ns instead of the usual value of 30 ns did bring all the $F90$ distributions closer together, but did not significantly affect the total $F90$.) However, the mean for each channel shifts roughly the same amount when the electric field is applied, which indicates that the effect cannot be solely due to the increased fraction of signal measured by channel 7. Since all of the PMTs are exposed to more light than is ideal when S2 signals are present, it is quite possible that they all experience a similar shift in performance.

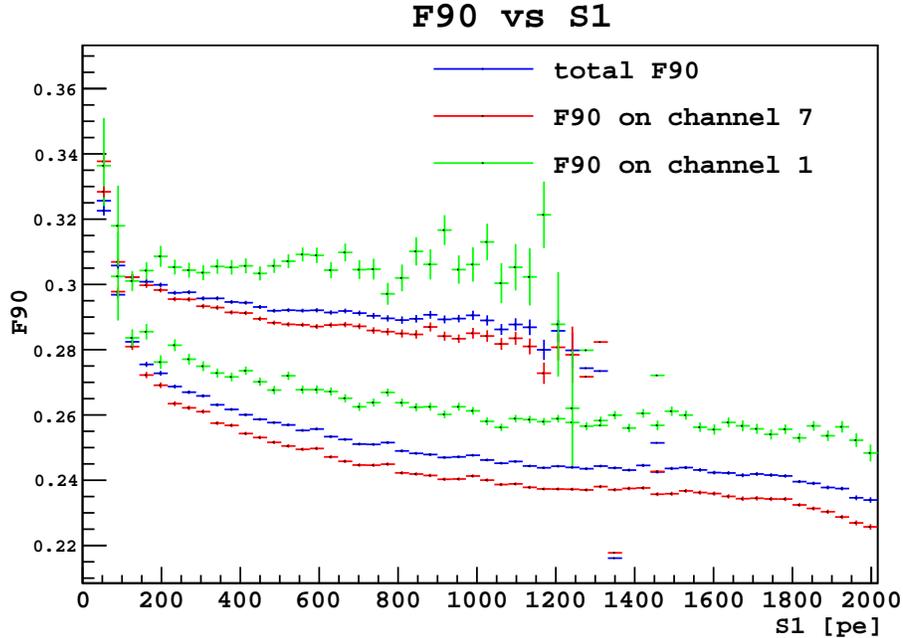


Figure 4.40: $F90$ vs $S1$ for all channels combined (blue), for channel 7 only (red) and for channel 1 only (green). The top three plots are from runs with electric field, the bottom three for runs without.

As further evidence that the shift in $F90$ is likely a physical effect, Figure 4.41 shows the distribution of $F90$ vs. $S1$ for two AmBe coincidence runs, one with no applied field and one with a drift field of 0.6 kV/cm. The gamma band (at low $F90$) shows the same shift with field as Figure 4.34, but the neutron distributions do not show any shift. (Incidentally, this also demonstrates that the absolute light yield for nuclear recoils does not change significantly in the presence of a 0.6 kV/cm field.) The fact that the two populations behave differently is strong evidence that the effect is physical, and most likely not an effect of electronics or analysis. Unfortunately, this indicates that the pulse shape discrimination power decreases when an electric field is applied, an effect which should be studied in more detail in the next campaign.

Finally, we can ask whether the spread of the $F90$ distribution is roughly consistent with what we expect from photon counting statistics and the noise inherent in the system. To evaluate this, we generate the $F90$ spectrum in four bins, 50 p.e. wide, over the range 100-

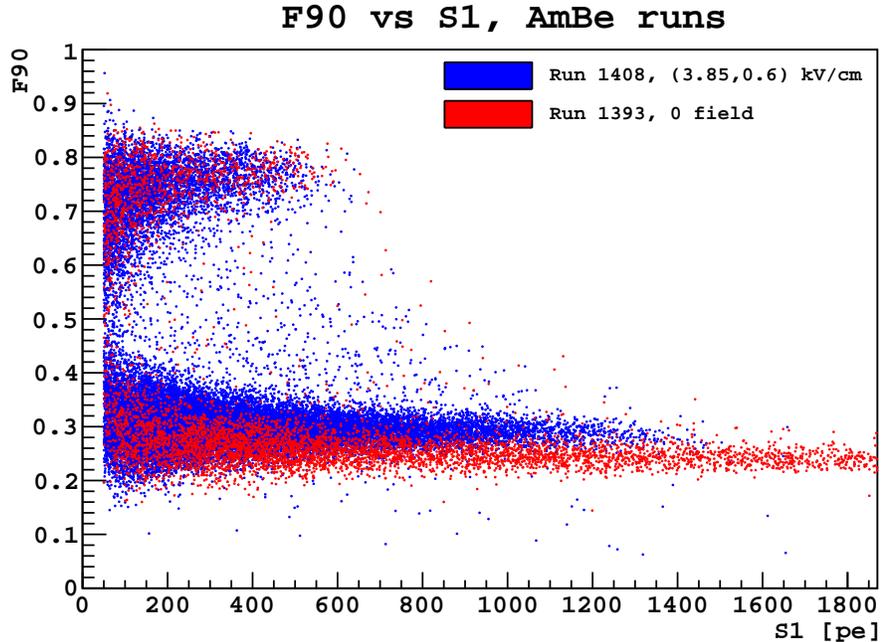


Figure 4.41: Distribution of $F90$ vs. $S1$ for two AmBe runs, with (blue) and without (red) electric field applied. The mean $F90$ of the gamma population shows a significant shift when the field is applied, but the nuclear recoil distributions are roughly the same.

300 p.e. For bins of this size in this energy range, we can approximate $F90$ as being constant over each 50 p.e.-wide bin. We then generated a Monte Carlo $F90$ distribution following the prescription in Section 2.4.4, with the mean value of $F90$ for each bin determined from the data. (The mean value of $F_{singlet}$ used in the simulation is found by inverting equation (4.5).) The energy spectrum for each bin is assumed to be flat, with the energy given by $S1/(\text{lightyield} \cdot \text{quenching})$, where the lightyield is 3.69 p.e./keV, and the quenching value is equal to 0.53 for runs with field and 1 for runs without field. In addition, the fractional width of the single photoelectron response is assumed to be 0.55 (dominated by the width of the 8" PMT at 0.58). The results of this simulation are shown in Figure 4.42 for the standard coincidence runs with field and Figure 4.43 for runs without field and summarized in Table 4.10.

Although the general shape of the distributions agree roughly by eye on the log scale shown, the width of the data distributions are consistently wider than predicted by Monte

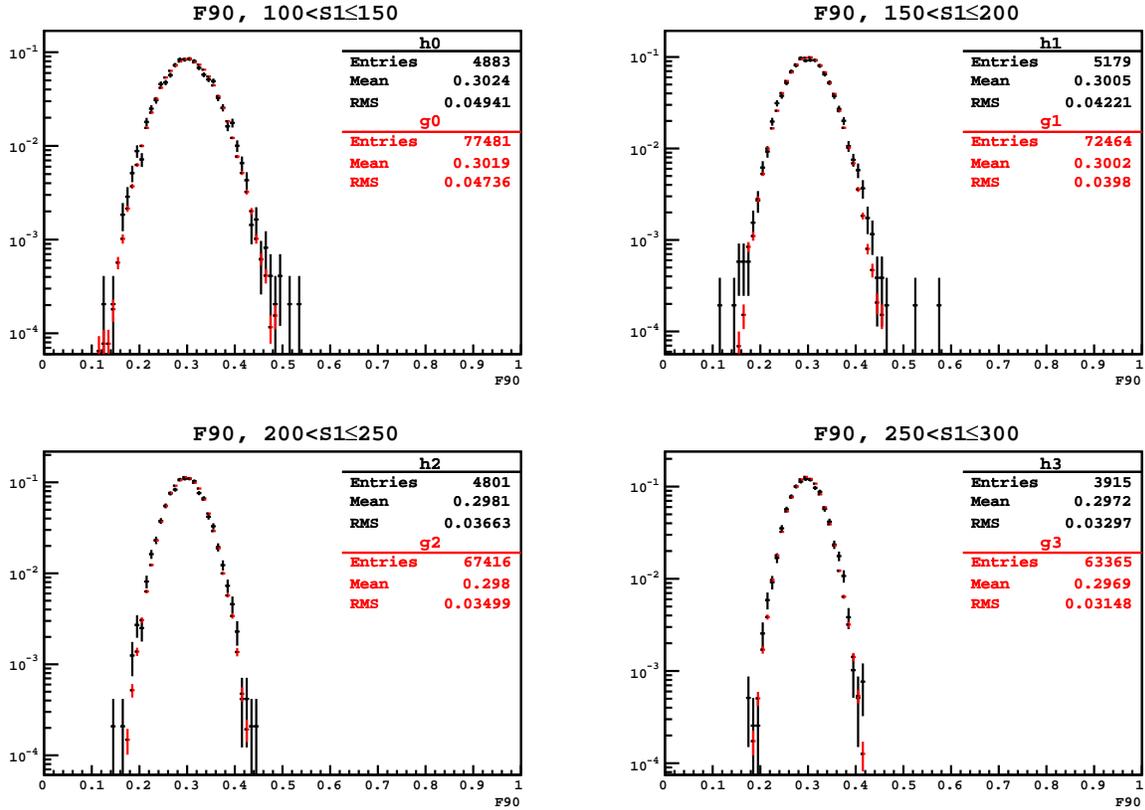


Figure 4.42: Comparison between $F90$ obtained from ^{22}Na coincidence runs with field and expected distribution generated by Monte Carlo.

Carlo; the extra spread is more pronounced in the runs without field. The most likely cause of this is a source of “noise” in the data which is unaccounted for in the Monte Carlo. For example, the variability of the 90 ns evaluation point due to PulseFinder’s identification of the start of the pulse, the timing offsets between the PMTs, the jitter in the arrival time of any given photon signal, the integration pedestal, and the finite width in time of the single photon pulse are all ignored in the Monte Carlo distribution. The assumptions that the mean value of $F90$ is constant and the spectrum is flat over the entire bin may also not be appropriate. Finally, the initial distribution of singlet and triplet states may not be well-described by a binomial distribution. Nevertheless, the agreement is good enough that it seems the majority of the spread of the $F90$ distribution is caused by photon statistics and the width of the single photoelectron response.

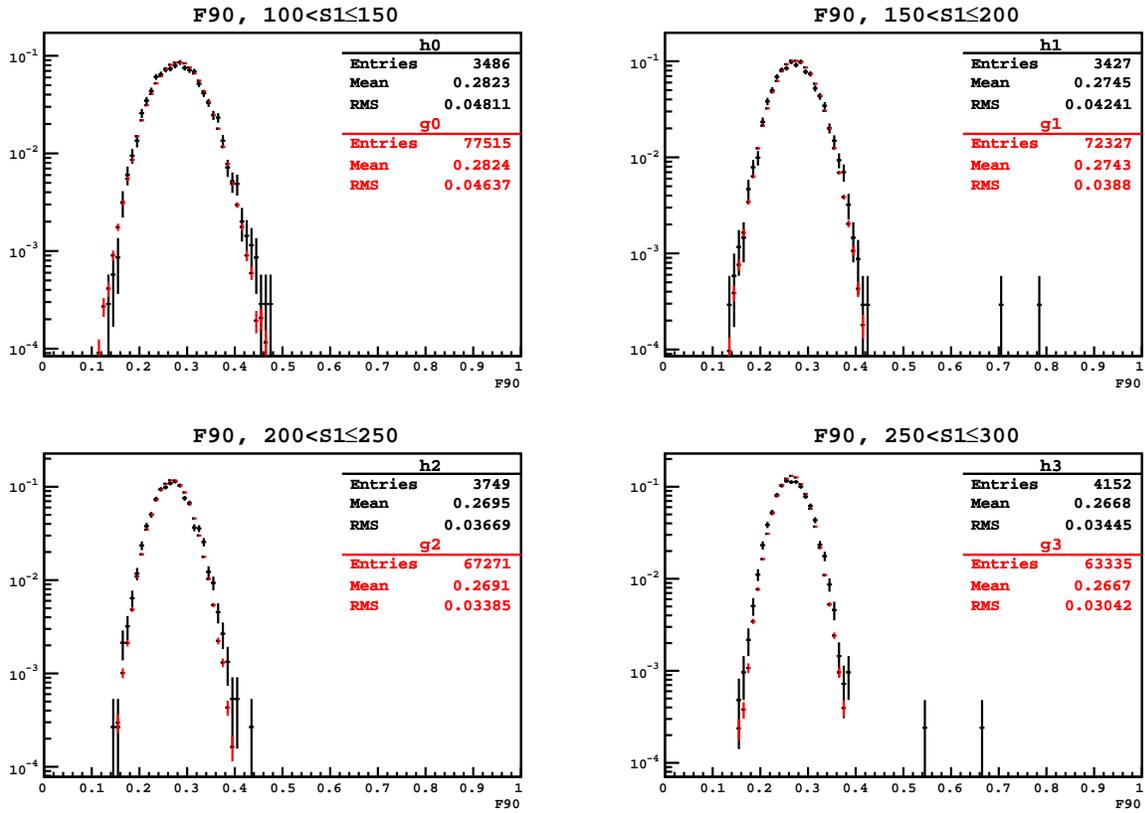


Figure 4.43: Comparison between $F90$ obtained from ^{22}Na coincidence runs without field and expected distribution generated by Monte Carlo.

Runs with field

Range (pe)	Data		Monte Carlo		χ^2/NDF
	Mean	RMS	Mean	RMS	
100-150	0.302	0.0494	0.302	0.0474	110/37
150-200	0.301	0.0422	0.300	0.0398	116/35
200-250	0.298	0.0366	0.298	0.0350	73.6/28
250-300	0.297	0.0330	0.297	0.0315	54.8/24

Runs without field

Range (pe)	Data		Monte Carlo		χ^2/NDF
	Mean	RMS	Mean	RMS	
100-150	0.282	0.0481	0.282	0.0464	50.3/34
150-200	0.274	0.0424	0.274	0.0388	87.3/31
200-250	0.270	0.0367	0.269	0.0338	121/27
250-300	0.267	0.0345	0.267	0.030	156/25

Table 4.10: Comparison of the distribution of $F90$ for 50 p.e. wide bins from data and Monte Carlo.

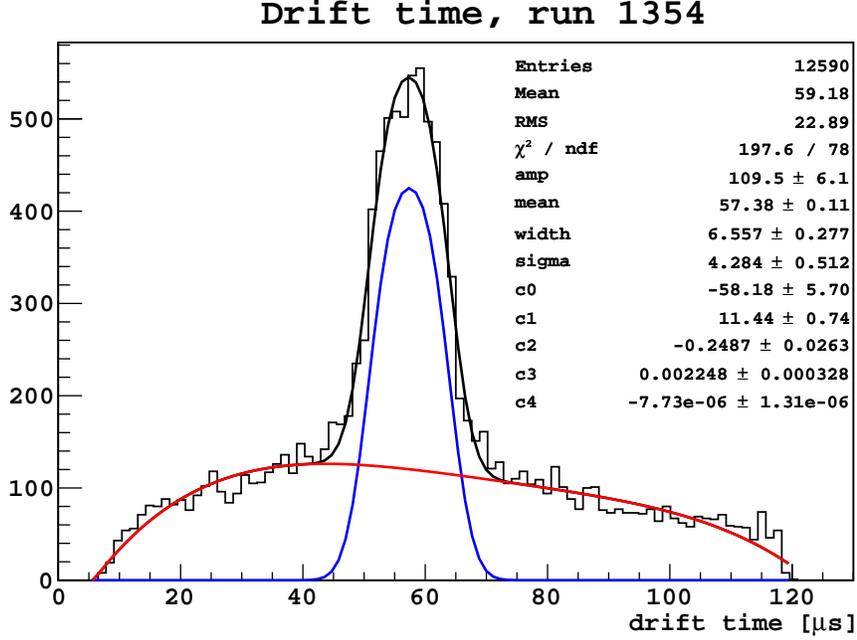


Figure 4.44: Drift time distribution from collimated ^{22}Na coincidence run 1354, showing the fit to Equation (4.6). The “signal” portion of the fit is shown separately in blue, with the “background” in red.

4.5.3 Ionization (S2)

Drift speed

The drift time was calculated using the ^{22}Na coincidence runs 1346, 1347, 1348, 1351, 1354, 1357, and 1360, which had a drift field of 0.7 kV/cm and an extraction field of 2.7 kV/cm. These runs are vertically collimated at different heights, and so present a narrow peak in the drift time distribution. In order to estimate the drift time for each source position, the drift time distribution was modeled as a “boxcar” function (i.e. constant over some range a – b , zero elsewhere) convolved with a gaussian, with a constant background modeled by a fourth-order polynomial:

$$y = a \cdot \text{erfc}\left(\frac{\mu - w - x}{\sigma}\right) \text{erfc}\left(\frac{x - \mu - w}{\sigma}\right) + p_4(x), \quad (4.6)$$

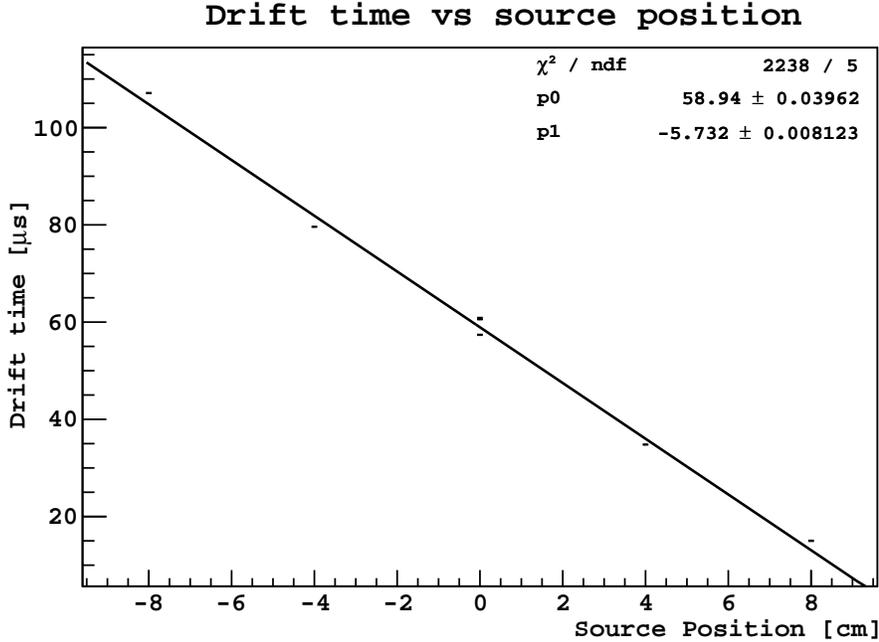


Figure 4.45: Drift time vs source height for collimated ^{22}Na coincidence runs. All the usual two-pulse cuts are applied except for the explicit cut on drift time. Vertical error bars are those directly reported by the fitter for the parameter μ in Equation (4.6). Drift speed extracted from the fit is $\sim 1.74 \text{ mm}/\mu\text{s}$.

where μ and w are the center and width of the boxcar function and σ is the width of the gaussian convolution; x , μ , w , and σ all have units of microseconds. See Figure 4.44 for an example fit to this function, for run 1354, which was nominally at the midplane of the detector. Likely because this is not a particularly accurate model for the distribution, the uncertainties returned by MINUIT are not very realistic.

Figure 4.45 is a plot of the estimated drift time versus height of the collimated beam relative to the approximate center of the active volume. Three of the points, runs 1346, 1347, and 1354, were all nominally at the center of the detector; the difference in the drift time center between the first two and 1354 is an indication of the precision to which the source was located. The points are fit to a straight line, which returns a slope of $5.732 \mu\text{s}/\text{cm}$, for a measured drift time of approximately $1.74 \text{ mm}/\mu\text{s}$, in very good agreement with other measurements made by the Icarus collaboration, reproduced in Figure 4.46 [11].

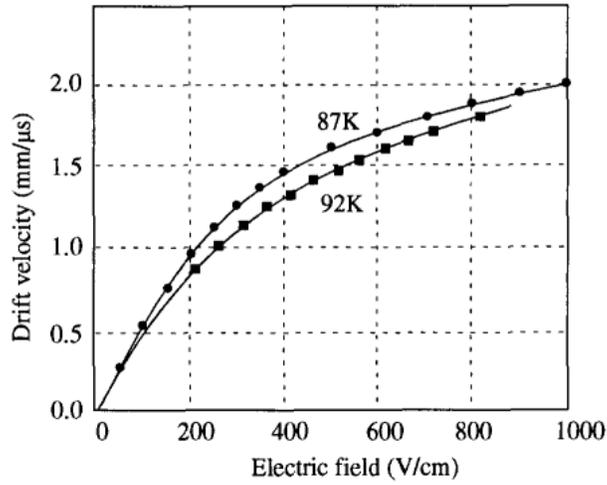


Figure 4.46: Electron drift speed in liquid argon vs. applied electric field [11].

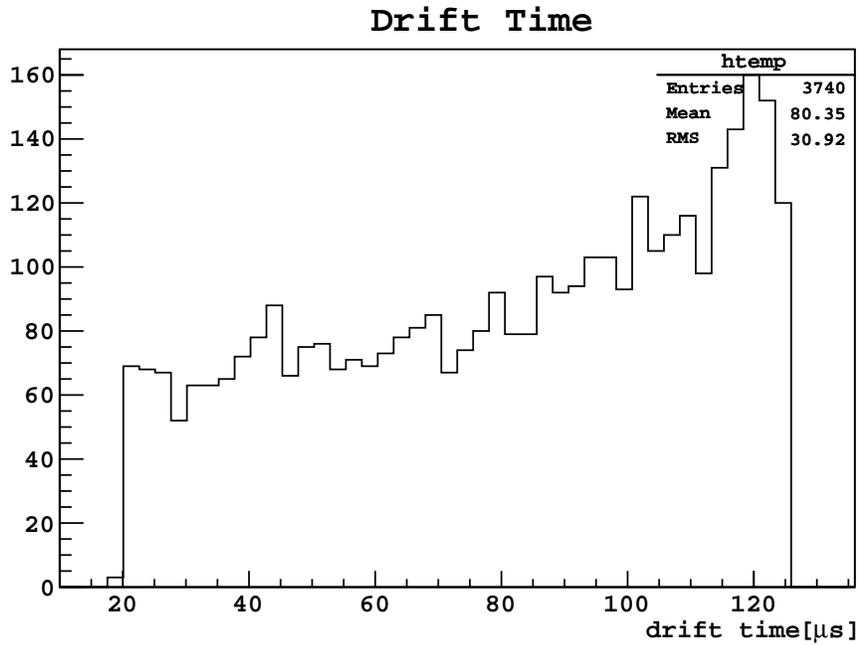


Figure 4.47: Drift time (time between S1 and S2) distribution for a background triggered run(Run1536)

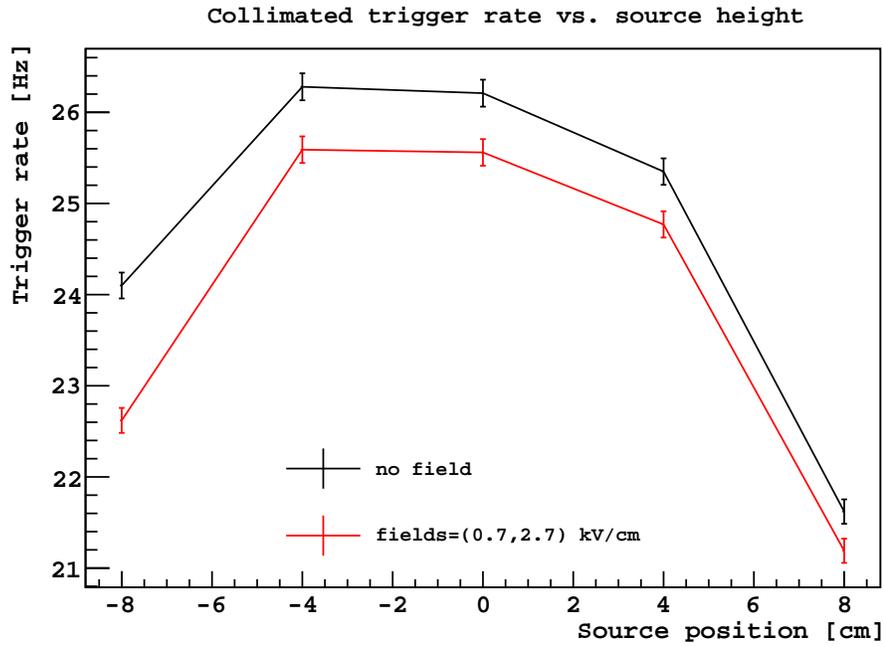


Figure 4.48: Trigger rate for collimated ^{22}Na coincidence runs as a function of the height of the source from the detector active volume midplane. Error bars shown are the square root of the number of triggers divided by the acquisition time.

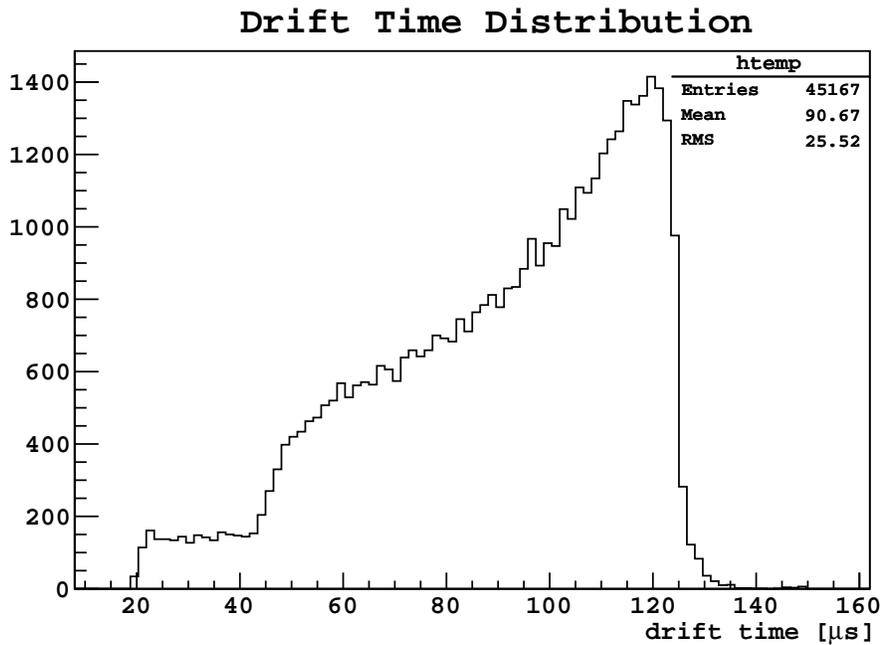


Figure 4.49: Drift time (time between S1 and S2) distribution for ^{22}Na runs in the memo dataset

Drift time

Figure 4.47 shows the drift time (time between S1 and S2) distribution of events passing cuts in run 1536, which was a background run taken with the drift field set to 0.6 kV/cm and extraction field of 3.3 kV/cm. In general we expect background events to be evenly spread throughout the detector, leading to a flat drift time distribution. The distribution in Figure 4.47, on the other hand, is peaked towards longer drift times. One possible explanation for this bias is that, because the trigger is formed solely by the 8" PMT, we trigger more favorably on events closer to the bottom of the detector. This seems unlikely because the 50 pe analysis threshold is significantly higher than the trigger threshold, and signals are generally seen on all of the (functioning) PMTs. An alternative explanation is that the rate of background decays is actually higher toward the bottom of the detector, possibly emanating from the 8" PMT itself, for instance, or from the floor. Finally, there could be some bias introduced by cuts that require only a single S2 (although it is not clear why this should be the case). To explore the question of trigger inefficiencies, Figure 4.48 shows the trigger rate for collimated ^{22}Na runs (roughly 1348-1368) as a function of the height of the source and NaI counter. (When comparing to Figure 4.47, recall that the x-axes are switched, i.e., long drift times correspond to negative source positions.) The sharp drop-off in trigger rate at ± 8 cm is most likely due to the collimated beam only partially intersecting the active volume and not trigger efficiencies, and the difference between the two graphs is likely due to the quenching effect of the electric field. The $\sim 3\%$ difference between the points at ± 4 cm could be the result of differing trigger efficiencies, but does not seem to be enough by itself to account for the shape of Figure 4.47.

Figure 4.49 shows the distribution of drift time for events in the standard ^{22}Na coincidence run set passing all cuts. If the active volume were evenly illuminated by the source, we would expect a drift time distribution similar to Figure 4.47; however, Figure 4.49 shows two notable differences: the prominent "step" around $50 \mu\text{s}$, and a more pronounced climb toward longer drift times. Both of these effects are likely due to the source not uniformly illuminating the

active volume, although other sources (differences in electronics performance, e.g.) cannot be ruled out at this point. Figure 4.50 illustrates the position of the source, detector, NaI counter, and shielding in order to indicate how the detector could be illuminated non-uniformly. The commercial NaI detector is surrounded on all sides by lead bricks, with an opening in the front (facing the Dewar). The bricks on the bottom extend slightly forward from the rest. The whole assembly sits on a hydraulic jack which allows it to be moved up and down.

For uncollimated runs, which make up the bulk of the dataset, the ^{22}Na source is taped to the center of a thin plastic box that has approximately same dimensions as the front of the NaI detector, and just fits into the opening in the lead. Almost all of the coincidence runs in this campaign were taken by either Ben Loer or Peter Meyers, and we recently discovered that we used different orientations of the box inside the opening. The orientation used by Ben Loer is that shown in Figure 4.50, with the box on its side between the source and the NaI detector. In this orientation, the height of the box puts the source at roughly the center of the NaI detector. The depth of the box is such that, if the top is placed flush against the NaI detector, the bottom (and the source) are slightly recessed into the opening in the lead. It was believed that the inner volume would not be eclipsed by any of the shielding in this configuration, but that may not be true. In runs taken by Peter Meyers, the box was laid flat, so that the source is on top, at roughly the center of the NaI. This orientation places the source closer to the NaI detector and more recessed into the shielding.

For collimated runs, the source is sandwiched between two half-bricks of lead top and bottom, held apart by copper bars approximately the same width as the source. A piece of scotch tape attached to the source is used to slide it into the slot, and is marked at the front of the slot to ensure it is pushed in far enough. Figure 4.51 is a photograph of the entire setup for a coincidence run with the top layer of shielding (and the top of the collimator) removed.

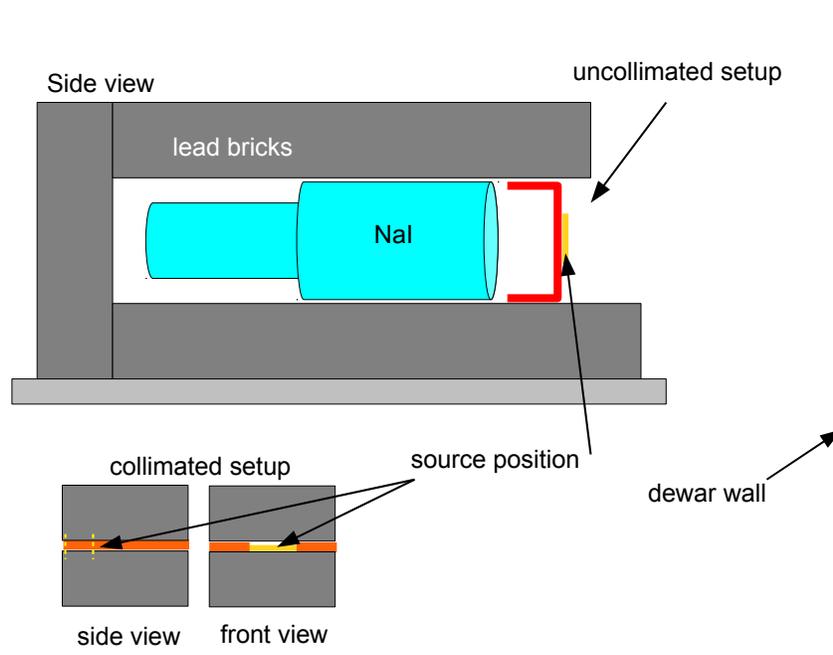


Figure 4.50: Cartoon of the setup of the ^{22}Na source with shielding and NaI detector for coincidence measurements. Dimensions are approximate, but not to scale.

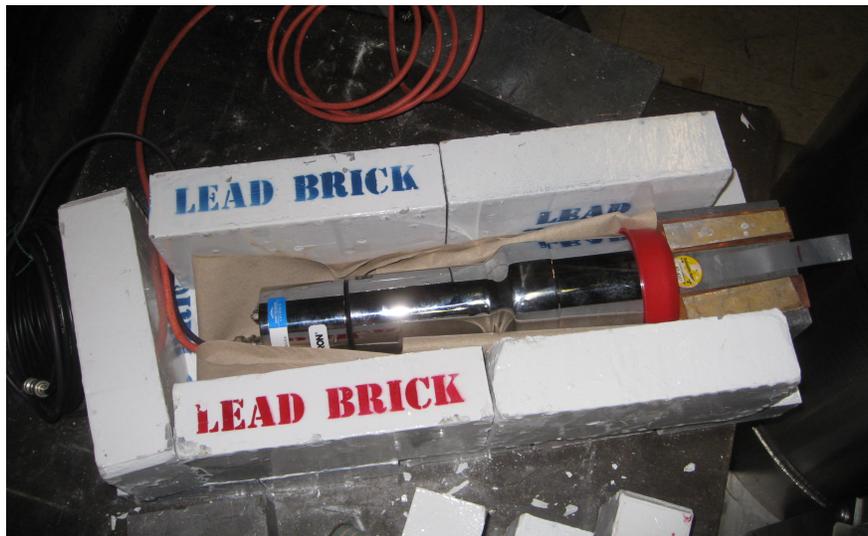


Figure 4.51: Photograph of the external NaI detector and shielding, set up for a collimated ^{22}Na coincidence measurement, with the top layer of shielding removed.

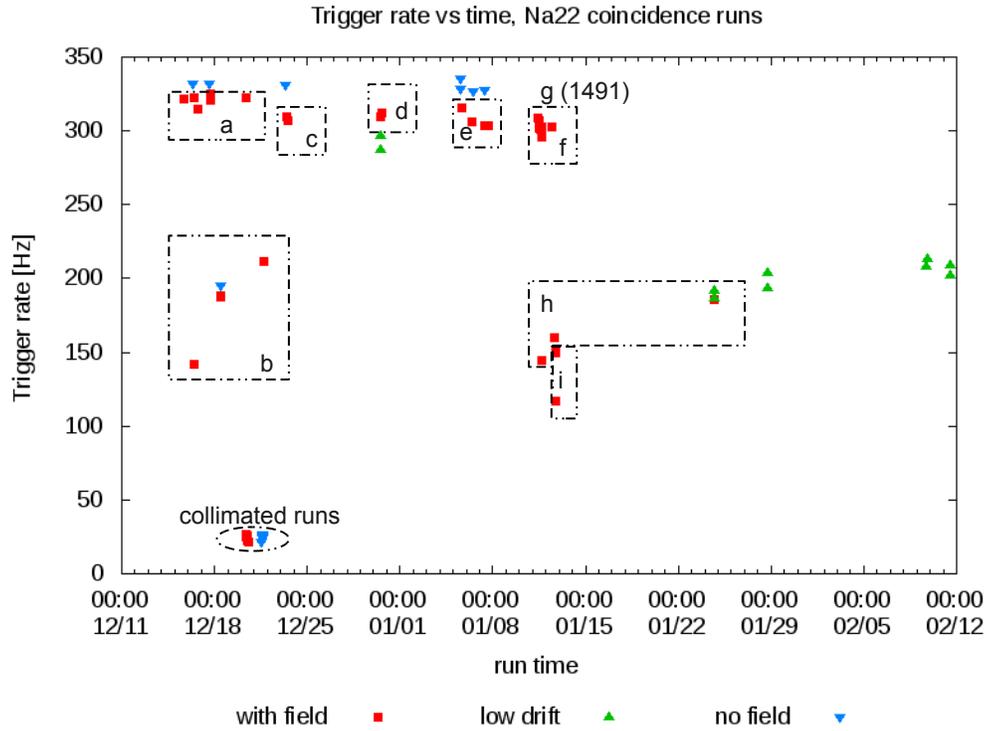


Figure 4.52: Trigger rate for ^{22}Na coincidence runs versus time. Representative drift time distributions are shown for each of the labeled groups in Figure 4.53 and described in the text.

To demonstrate the variability of drift time distributions and why at least part of it is likely due to exact positioning of the source, Figure 4.52 shows the trigger rate for ^{22}Na coincidence runs over time, since Run 1295, when the current trigger logic was implemented. The runs labeled “low drift” (green triangles) are special-purpose runs taken with very low drift field (see Section 4.5.3 for details). The runs with applied drift field greater than 0.2 kV/cm (the “usual” range) are divided into similar groups for further discussion. For each of the labeled groups in the figure, Figure 4.53 shows the drift time distribution for a representative set of those runs. The groups are as follows:

- (a) Runs 1295, 1305, 1314, 1328, 1330, and 1343 all had a extraction field of 2.7 kV/cm and drift field of 0.8 kV/cm except for 1295, which had 0.5 kV/cm drift field. Trigger rates for these runs were all around 320 Hz. The runs were set up variously by Alex

Wright, Ben Loer, or Peter Meyers. Runs in this group show a mostly flat drift time distribution.

- (b) Runs 1308, 1336, 1339, and 1369 had extraction fields at 2.7 kV/cm, and drift fields set at 0.8 kV/cm, except 1369, which had 0.7 kV/cm drift field. Trigger rates were 150-200 Hz. Run 1308 was set up by Alex Wright; the others were set up by Ben Loer. The drift time distribution for these runs shows a sharp climb with increasing trigger time, but no obvious “step”, although there may be some small hint of it.
- (c) Run 1396 and 1405 had extraction fields at 3.85 kV/cm; Run 1396 had a drift field of 0.5 kV/cm, while it was 0.6 kV/cm for 1405. Both runs had trigger rates around 300 Hz, and were set up by Peter Meyers. The drift time distribution for these runs shows a sharp climb for longer drift time.
- (d) Runs 1428 and 1433 both had a extraction field of 3 kV/cm, with drift fields of 0.4 kV/cm for Run 1428 and 0.7 kV/cm for Run 1433. Both had trigger rates around 300 Hz and were set up by Peter Meyers. The drift time distribution for these runs is essentially flat.
- (e) Runs 1444, 1450, 1457, and 1460 all had extraction fields at 3.85 kV/cm and drift fields at 0.6 kV/cm (these are in the primary run set). All had trigger rates around 300 Hz and were set up by Peter Meyers. The drift time distribution for these runs show a very clear “step” at around 50 ns and a rise at longer drift time.
- (f) Runs 1465, 1468, 1471, 1476, 1483, 1488, 1498, 1501, and 1506 all had extraction field at 3.3 kV/cm and drift field at 0.6 kV/cm. All had trigger rates around 300 Hz and were set up by Peter Meyers. Runs in this set show a slight rise at longer drift times, but no “step.”
- (g) Run 1491 had extraction field at 2.7 kV/cm and drift field at 0.6 kV/cm. This run had a trigger rate of 295 Hz, and was set up by Peter Meyers. Note that it takes place between the runs in group (f). This run has a flat drift time distribution.

- (h) Runs 1501, 1511, 1512, and 1539 all had extraction fields of 3.3 kV/cm and drift fields of 0.6 kV/cm. Trigger rates for these runs were between 140 and 185 Hz. All were set up by Ben Loer. Runs in this group have a drift time distribution which increases very steeply with increasing drift time, but no evidence for the “step.”
- (i) Runs 1519, 1522, and 1523 all had extraction fields of 3.85 kV/cm and drift fields at 0.6 kV/cm (and are in the primary run set). Run 1519 had a trigger rate of about 117 Hz, while the others had around 150 Hz trigger rate. All were set up by Ben Loer. Runs in this group show both a steep rise toward higher drift times and a “step” at around 50 ns.

Very few patterns can be found in the preceding evidence. The total trigger rate seems to depend only on who set the source in place for the run. There does not seem to be any strong correlation between the shape of the drift time distribution and the trigger rate. The “step” feature around 50 ns in the drift time distribution only appears clearly when the extraction field is set to 3.85 kV/cm (the maximum value at which data was acquired), but there are possible hints of it at lower fields. Perhaps most telling are groups (a) and (b), which were acquired under near identical conditions, but which have very different drift time distributions. So, while we cannot rule out whether the field settings had some effect on the drift time distribution, it seems not to be the most important determining factor.

S2 spectrum

Figure 4.54 shows a scatterplot of S2 vs S1 for ^{22}Na and AmBe runs, and figure 4.55 shows the ratio of S2/S1 for those same runs. For the AmBe runs in Figure 4.55, the S2/S1 spectra are shown separately for events with $F90$ greater or less than 0.5, which largely selects between gamma and nuclear recoil events. The long tail into low S2/S1 of the gamma events limits the discrimination power of the S2/S1 statistic. If we assume that the AmBe spectrum with $F90 > 0.5$ is representative of the WIMP-induced spectrum, Figure 4.56 shows the leakage of ^{22}Na gamma events into the signal region when considering only the S2/S1 cut, as a function

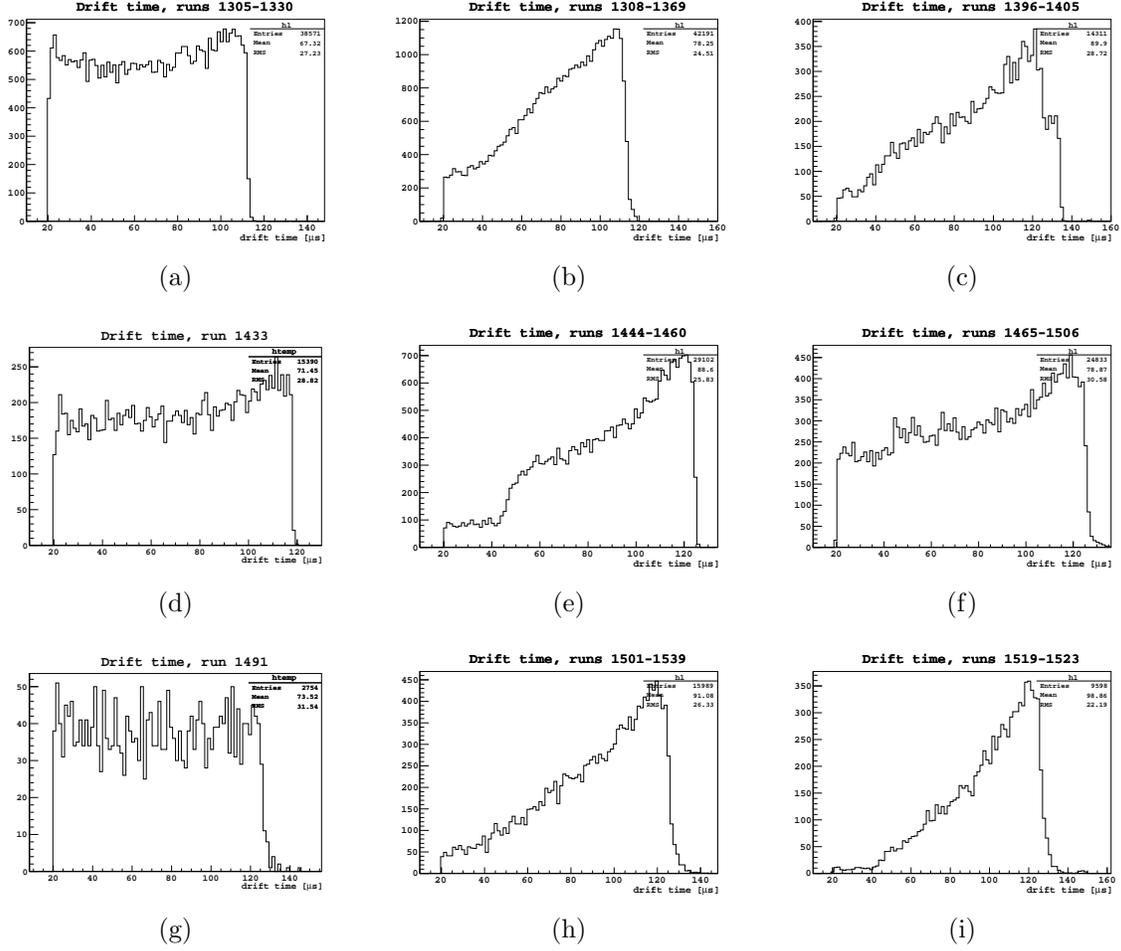


Figure 4.53: Representative drift time distributions for runs grouped in Figure 4.52. In each group, only runs with the most common drift and extraction field settings are included.

of the acceptance fraction for nuclear recoils. Because the spread of the distribution is wider at lower energies (due to photo-statistics), Figure 4.56 also shows the acceptance fraction for events with $S1 < 200$ p.e., which roughly corresponds to the region of interest for WIMP searches (recall that with field, $1 \text{ p.e.} \sim 1 \text{ keV}_{nr}$). For all triggers, the leakage of ^{22}Na events into the nuclear recoil acceptance region is approximately 1%, 3%, and 6% for nuclear recoil acceptances of approximately 50%, 90%, and 95%. Considering only events with $S1 < 200$ p.e., the corresponding leakage fractions are 2%, 6% and 10%. Figure 4.57 shows the change in $S2/S1$ for ^{22}Na events when $S2$ is corrected for the drift time (see section 4.5.3). The fact

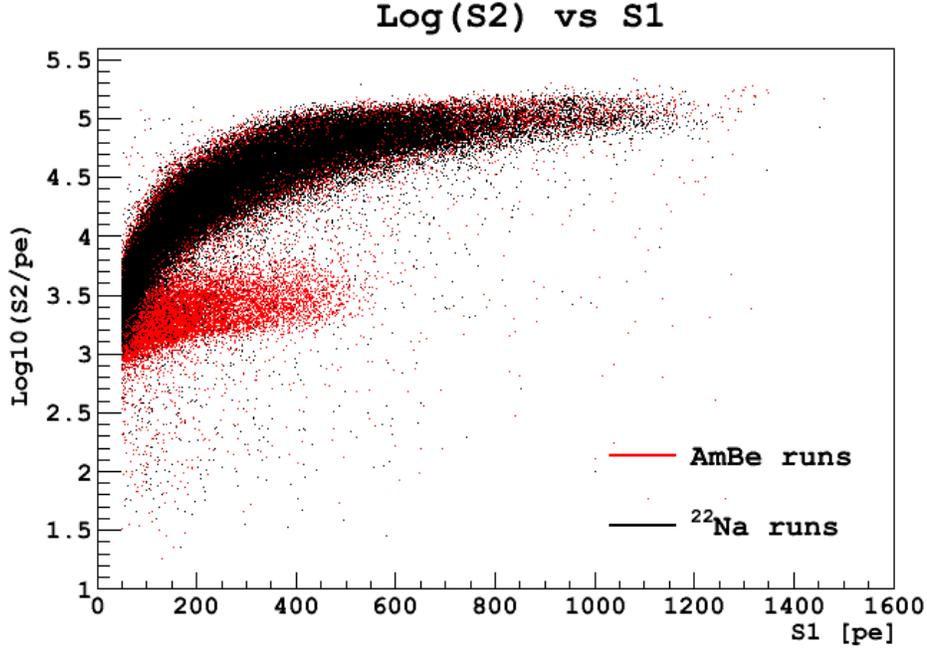


Figure 4.54: Scatterplot of the logarithm of S2 as a function of the event energy (S1) in photoelectrons, for ^{22}Na and AmBe runs. The population that appears at low S2 ($\sim 10^{3.5}$) in the AmBe runs but not ^{22}Na is assumed to be due to argon nuclear recoils.

that the width of the distribution does not change significantly is indication that the electron lifetime is not a large factor in determining the spread for these runs.

Although the mean values of the $S2/S1$ and $F90$ parameters both depend on the ionization density of the event, it is believed that the fluctuations in these parameters for a given event are not correlated. If this is true, then the parameters are independent, and the total discrimination power is the product of the power of each parameter. In order to test this assertion, Figure 4.58 shows the $S2/S1$ distribution for ^{22}Na events, split into two spectra, one for events drawn from the center bulk of the $F90$ distribution ($0.25 < F90 < 0.35$) and one for events drawn from the nuclear recoil-like (i.e., high) tail of the $F90$ distribution ($F90 > 0.35$). We take the additional cut that $80 \text{ p.e.} < S1 < 150 \text{ p.e.}$ in order to restrict the analysis to a region where the mean of the $F90$ distribution is essentially flat with respect to S1. If $F90$ and $S2/S1$ are independent, then these two distributions should be the same aside from statistical fluctuations. Visually, the agreement is quite good; the χ^2 value be-

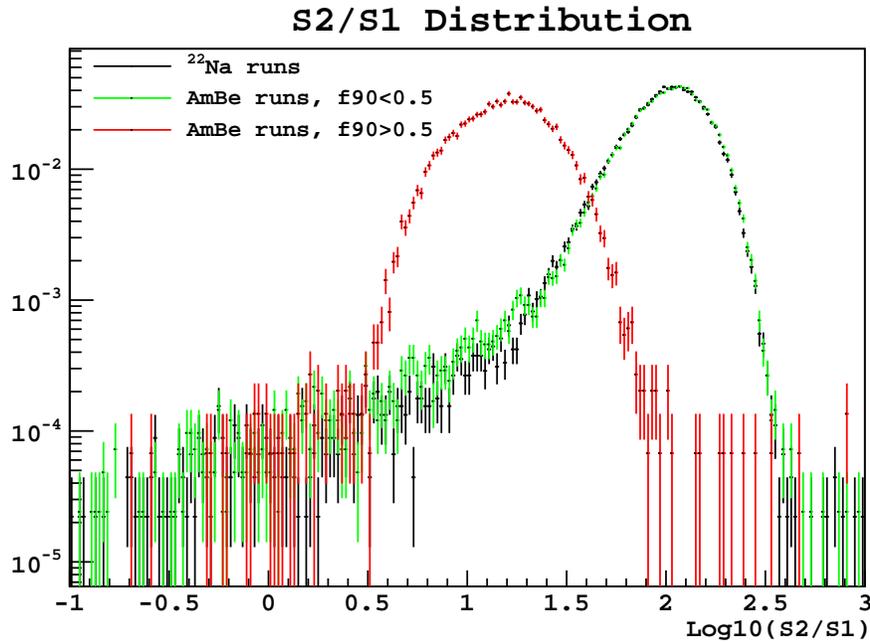


Figure 4.55: Distribution of the logarithm of S2/S1, the scintillation-to-ionization discrimination parameter. The black spectrum is all events passing cuts from the ^{22}Na runs; the green spectrum is events from AmBe runs with $F_{90} > 0.5$, assumed to be mostly gammas, and the red spectrum is events with $F_{90} < 0.5$, assumed to be mostly nuclear recoils. Each spectrum is separately normalized to unit integral.

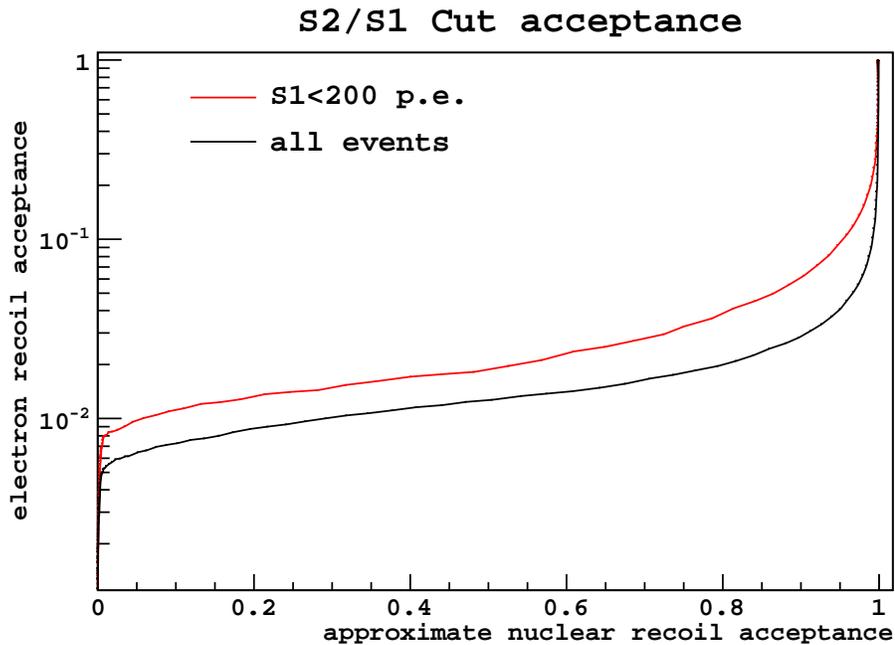


Figure 4.56: Leakage of S2/S1 for ^{22}Na events into the nuclear recoil signal region, as a function of the acceptance fraction for nuclear recoils.

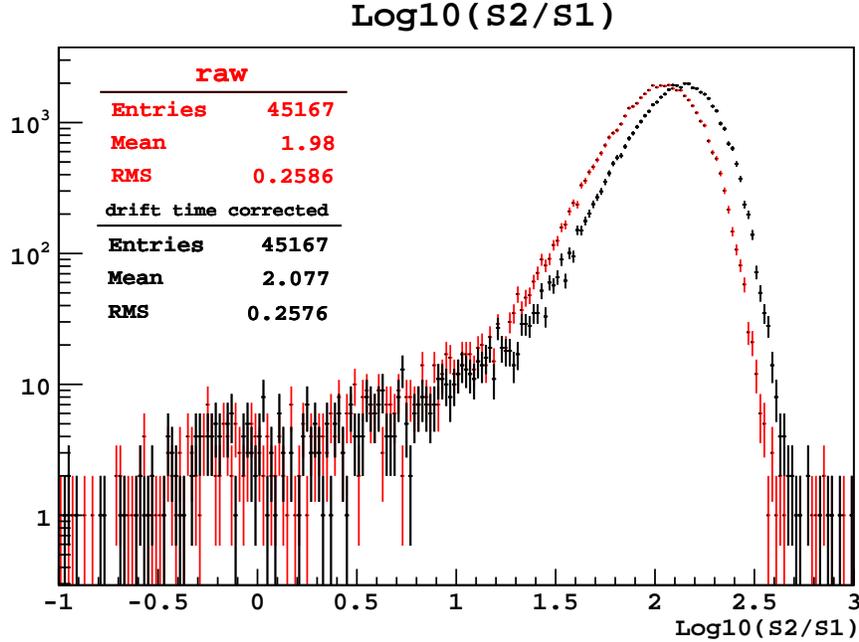


Figure 4.57: Distribution of the log of S2/S1 with and without correcting for the exponential decay with drift time.

tween the two is 45.4 with 68 degrees of freedom, indicating quite good agreement (and that the number of entries in several bins is low enough that the \sqrt{N} uncertainty approximation used to perform the χ^2 test is not very valid). As a secondary check, Figure 4.59 shows the fraction of events with $\text{Log}_{10}(\text{S2/S1}) < 1.5$ for each of several bins in $F90$ 0.05 wide, with the same cuts on S1. The uncertainty shown for each bin is the square root of the number of events with $\text{Log}_{10}(\text{S2/S1}) < 1.5$ over the total number of events in that $F90$ bin. Although there is some variation in the bins, a fit to a constant returns a χ^2/NDF of 4.76/5, indicating that the distribution is consistent with being flat, i.e., that S2/S1 is independent of $F90$ in this energy range. However, it should be noted that a fit to a quadratic function yields a much better fit, with a χ^2/NDF of 1.71/3.

The total discrimination power of the detector is represented by the combination of the drift-time-corrected S2/S1 variable and $F90$. Figure 4.60 shows the usual scatterplot of events in $\log(\text{S2/S1})$ vs $F90$ space, showing the separation between gamma and nuclear recoil events. The two points in blue are the only events from the ^{22}Na runs that pass all

cuts and have $F90 > 0.6$. Depending on the exact nuclear recoil acceptance bounds chosen, this gives 2 background events from a population of 45167 events which pass all cuts. It's impossible to tell with the high rate of backgrounds in the detector whether these events are due to real nuclear recoils, pileup, or statistical fluctuations of normal gamma events. Visual inspection of the two events did not indicate any anomalies.

As an alternative way to view these events, figure 4.61 shows the distribution of $S2/S1$ vs $F90$, with the event energy ($S1$) as the Z-axis and color scale, for events from the ^{22}Na runs. From this plot it is obvious that the two background events are very low energy, close to threshold; the actual values are 73 and 80 photoelectrons. Again, its difficult to draw any conclusions from this. Figure 4.62 plots $F90$ vs $S1$, with the log of $S2$ on the z-axis. The most interesting feature of this plot is the population of events with an intermediate $F90$ value. Because the AmBe source releases both neutrons and a high rate of gammas (both 60 keV gammas from the ^{241}Am decay and ~ 4 MeV gammas from the (α, n) process), it is likely that a number of events in this region would be those where both a gamma and neutron (producing a nuclear recoil) interacted in the detector, but were not flagged as multiple-sited by cuts.

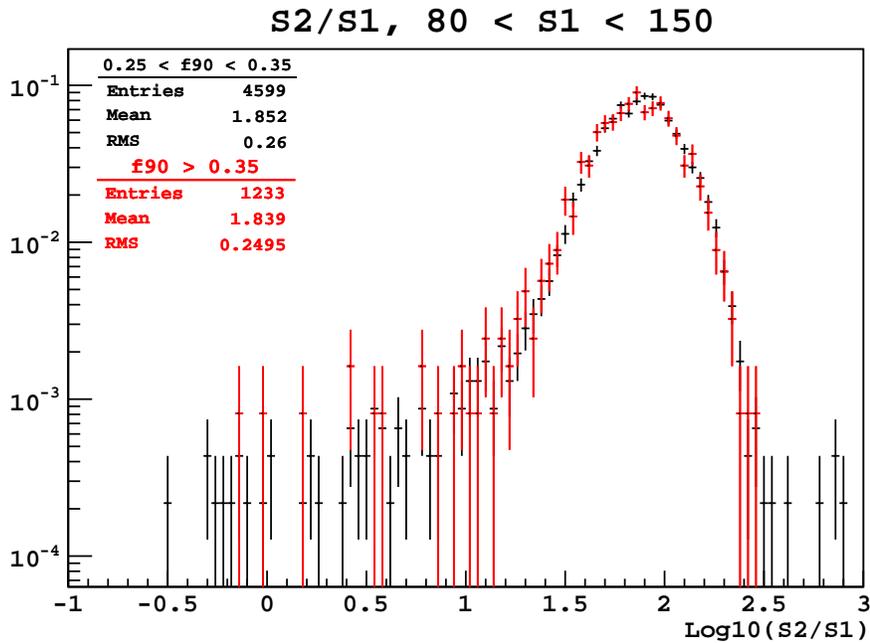


Figure 4.58: Distribution of the log of $S2/S1$ after correcting for drift time for events with $80 < S1 < 150$ photoelectrons, with different cuts taken on the $F90$ parameter. Each spectrum is separately normalized to unit integral. The close agreement between the two spectra demonstrates the independence of the $F90$ and $S2/S1$ parameters.

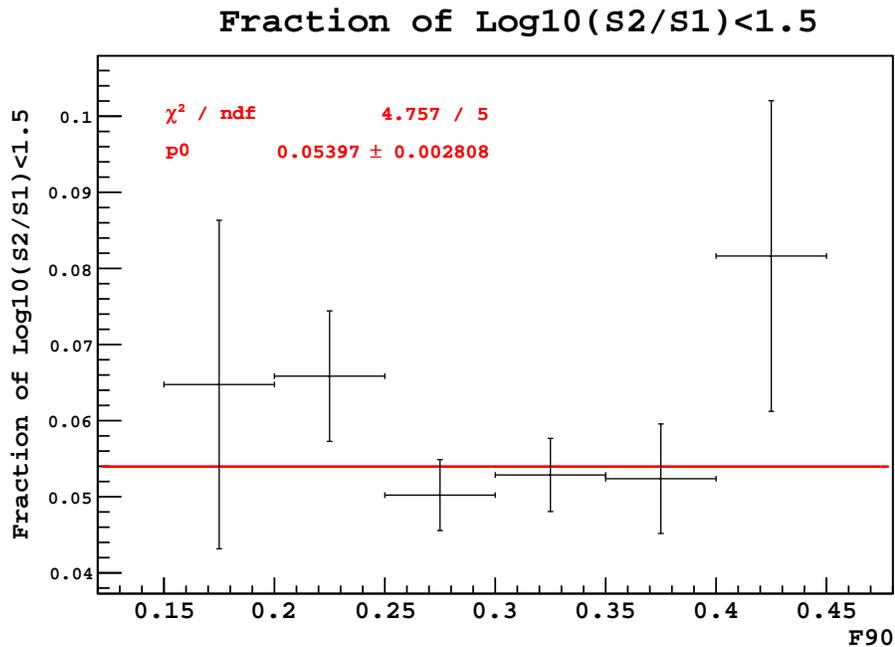


Figure 4.59: Fraction of events with $\text{Log}_{10}(S2/S1) < 1.5$ as a function of $F90$ for events with $80 < S1 < 150$ p.e.

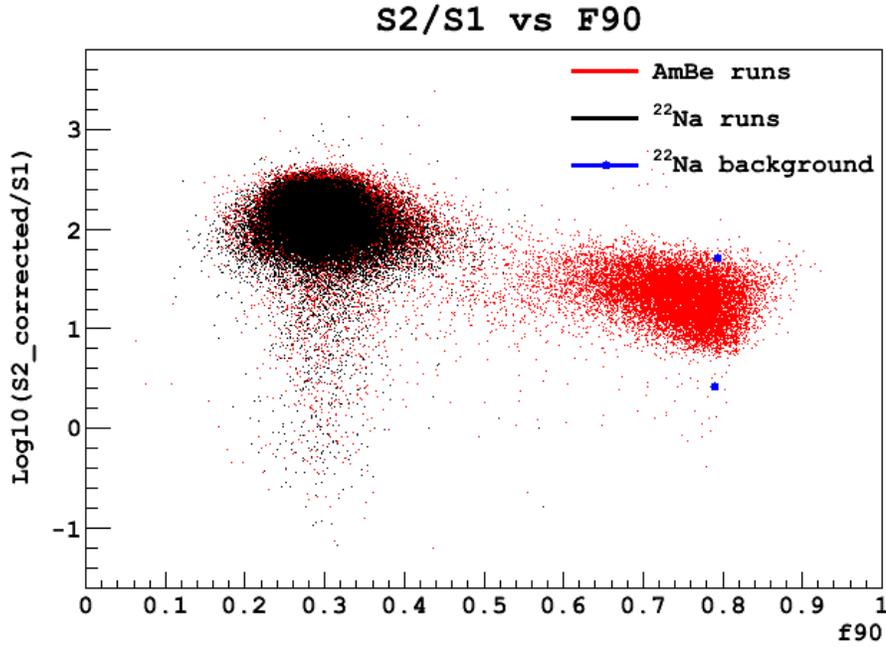


Figure 4.60: Scatterplot showing the distribution of the two main discrimination parameters, $S2/S1$ and $F90$, for ^{22}Na and AmBe runs. The two blue points represent potential background events for a dark matter search; see text for details.

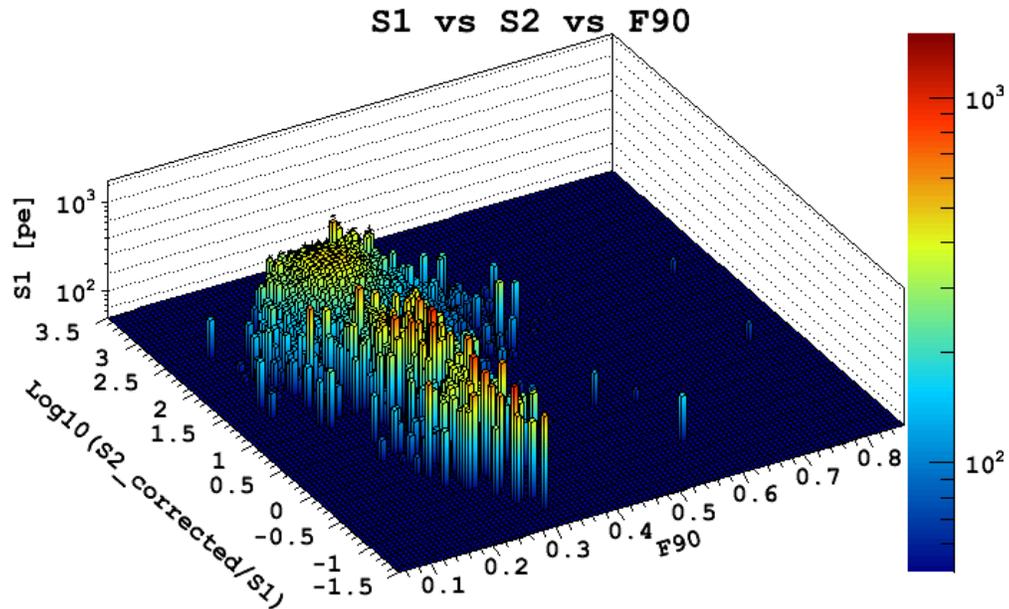


Figure 4.61: Distribution of the two discrimination parameters $S2/S1$ and $F90$ for the standard run set, with the average energy ($S1$, in photoelectrons) for each bin given by the Z axis and color scale. This shows that the two background events of Figure 4.60 are at low energy, near threshold.

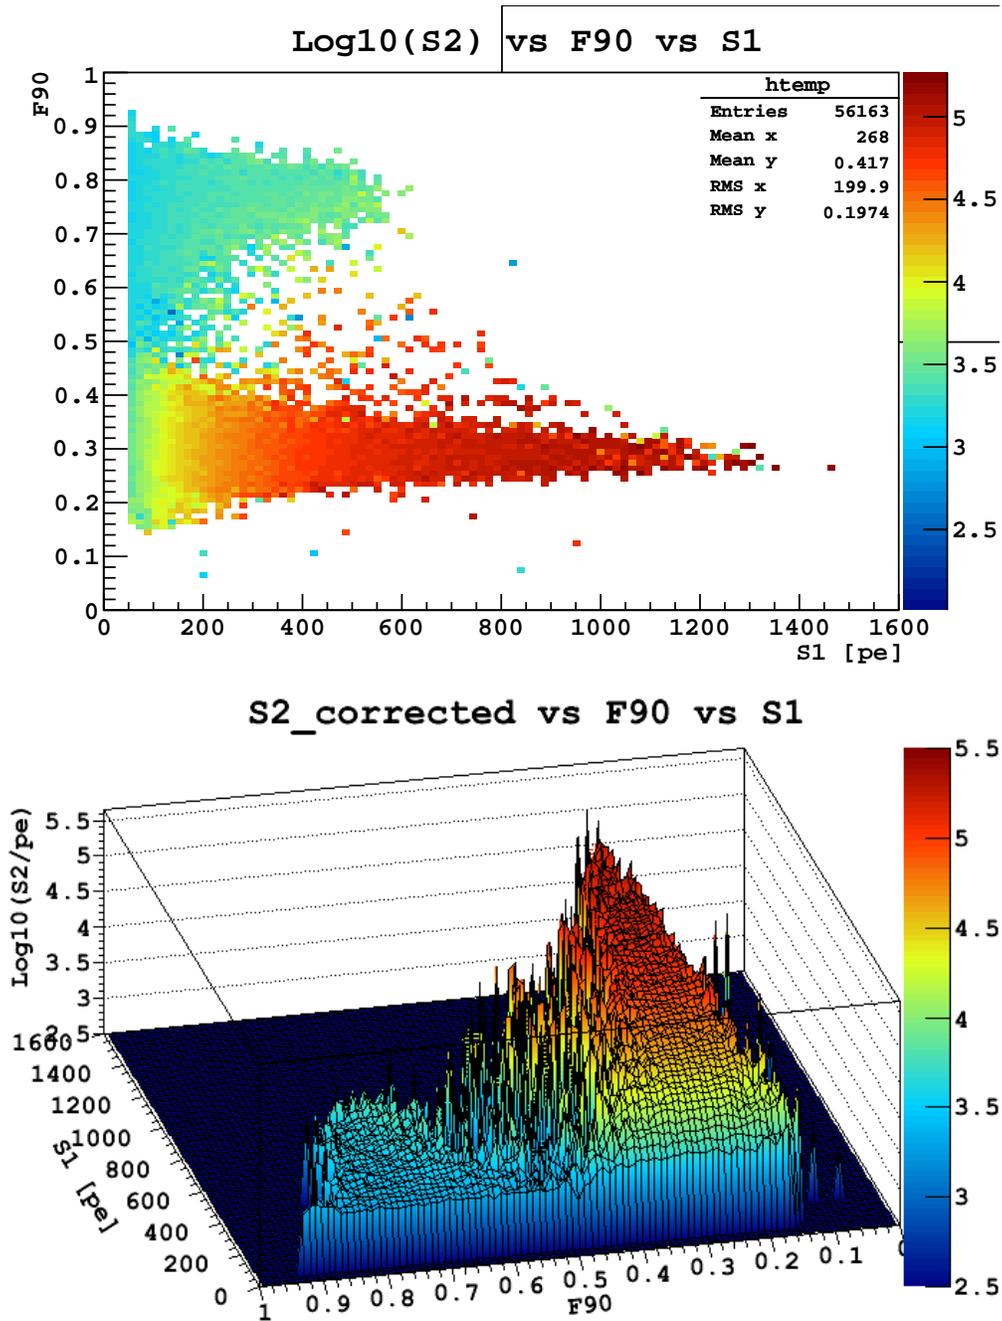


Figure 4.62: Distribution of $F90$ as a function of event energy ($S1$) for the standard run set, showing the log of $S2$ on the Z axis (bottom) and color scale (top and bottom). The events with intermediate $F90$ are likely events in which both a gamma and neutron deposited energy in a small enough volume to be seen as a single-sited event.

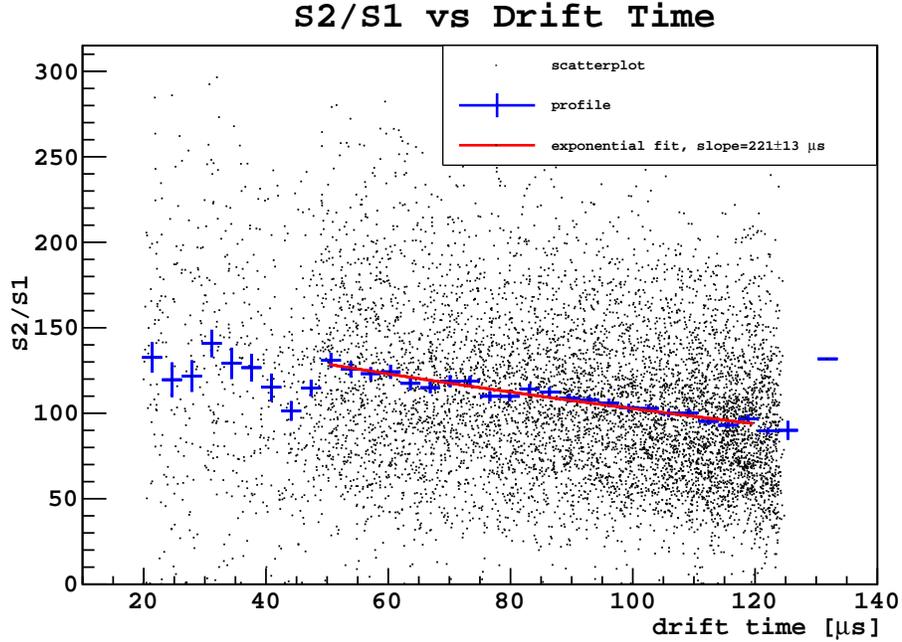


Figure 4.63: S2/S1 event distribution as a function of drift time for run 1460, showing the wide scatter of values for individual events and the profile (mean \pm uncertainty) fit to an exponential.

Electron lifetime

As a sensitive way to evaluate the purity of the argon in the detector, we can measure the dependence of the ratio S2/S1 as a function of the drift time of the event. If there is some contamination of electronegative impurities in the detector, free electrons will have a finite lifetime, as they will eventually encounter and be captured by these impurities. Thus, the population of free electrons will tend to decrease exponentially with time. Although there is a large spread in the ratio of S2/S1 over a population of events, the mean value of S2/S1 measured as a function of drift time should follow the same exponentially decreasing trend.

To attempt to quantify this value, we first bin events according to their measured drift time. The uncertainty in the drift time is primarily a function of the down-sampling of the PulseFinder, and so is on the order of (slightly better than) 1 μ s. (See Section 4.3.5 for a more complete description.) For each bin, we find the mean of S2/S1 for all the events in the bin, and assign it an uncertainty of the RMS of those events divided by \sqrt{N} , where N is the

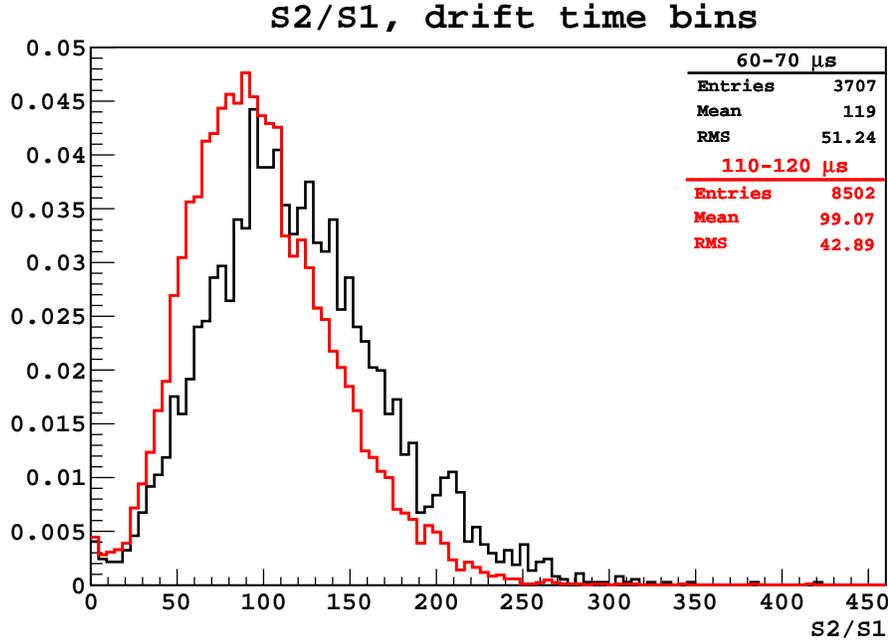


Figure 4.64: Example S2/S1 distributions from ^{22}Na runs in two drift time bins: 60–70 μs and 110–120 μs .

number of events in that bin. Figure 4.64 shows the S2/S1 distribution from two example bins covering 60–70 μs and 110–120 μs ; as expected, the mean value of S2/S1 is lower at longer drift time.

This mean vs. drift time distribution is then fit to an exponential function. For the standard runs in this analysis, we fit between drift times of 50 μs to 120 μs , in order to attempt to include only events which seem to come from true coincidence triggers (in particular avoiding the “step” at 50 ns visible in the drift time distributions). Figure 4.63 shows as an example the fit from run 1460. To estimate the systematic uncertainty due to the fit range, Run 1460 was fit using a range of start points from 20 to 60 μs and end points from 90 to 130 μs , the results of which are reported in table 4.11. The RMS spread about the mean for all results is 22%; the RMS for only those ranges where $\chi^2/\text{NDF} < 1.5$ is 11%. Note that the lifetime returned from the fit decreases with the high endpoint of the fit. For the fits that extend below the “step” in the drift time plot (around 50 μs), this is likely because the

Min (μs)	Max (μs)	τ (μs)	χ^2/NDF
20	90	417 ± 75	38.9/20
20	100	352 ± 41	42/23
20	110	313 ± 25	45.5/26
20	120	271 ± 15	59/29
20	130	258 ± 13	65/31
30	90	388 ± 73	37.3/17
30	100	331 ± 40	39.6/20
30	110	297 ± 25	42.2/23
30	120	259 ± 15	53.7/26
30	130	247 ± 13	58.8/28
40	90	450 ± 112	34.2/14
40	100	350 ± 49	37.5/17
40	110	304 ± 28	40.7/20
40	120	260 ± 16	52.6/23
40	130	246 ± 13	57.7/25
50	90	252 ± 42	9.29/11
50	100	248 ± 28	9.65/14
50	110	242 ± 20	10.1/17
50	120	221 ± 13	16.3/20
50	130	213 ± 11	18.6/22
60	90	305 ± 89	7.76/8
60	100	269 ± 44	8.47/11
60	110	251 ± 27	9.17/14
60	120	221 ± 15	15.7/17
60	130	212 ± 13	18.1/19
Avg (all fits)		287 ± 62	
Avg ($\chi^2/\text{NDF} < 1.5$)		244 ± 27	

Table 4.11: Result of fits to the electron lifetime for run 1460, for various fit start and end points.

distribution is obviously not well-described by a single exponential; the fact that the pattern holds also for fits starting above the “step” is harder to explain.

Another possible contribution to the systematic uncertainty of the fit is the dependence on height of the S1 response (shown in Figure 4.27). Since the light yield for S1 is highest towards the bottom of the detector (i.e., at longer drift times), this will tend on average to decrease S2/S1 at long drift times, and increase it at short drift times, possibly mimicking an exponential decay, which could cause us to underestimate the electron lifetime. To attempt to quantify this, we attempt to correct S1 in run 1460 for the height dependence and fit

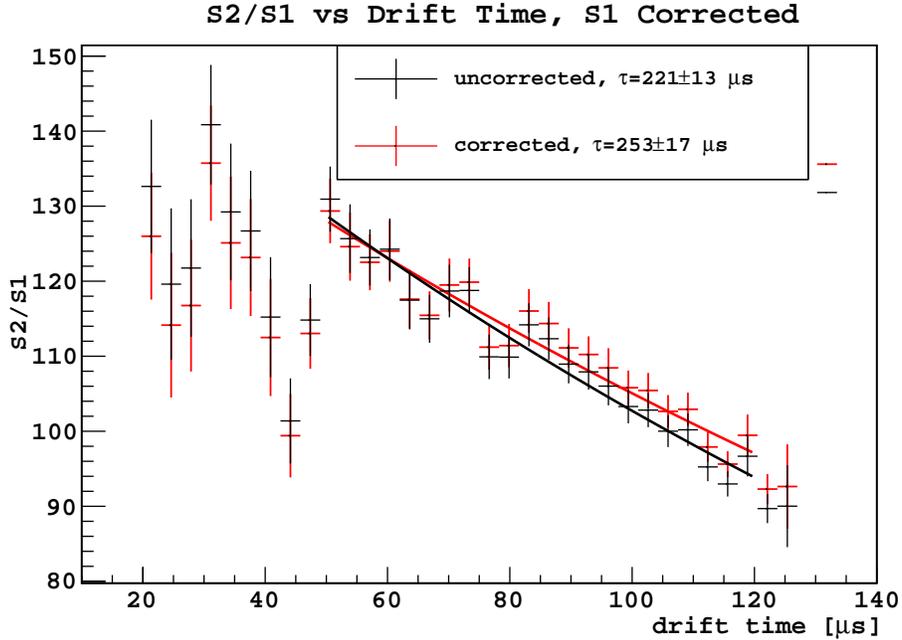


Figure 4.65: $S2/S1$ vs drift time for run 1460, with $S1$ corrected for the z -dependence of light yield. See text for details.

for lifetime. To do this, we fit the graph in Figure 4.27 to a quadratic function with fixed intercept of 1, which returns $y = 1 - 5.38 \times 10^{-3}z - 2.485 \times 10^{-4}z^2$, where z is height in cm relative to the detector midplane. Then we assume a drift speed of $0.173 \text{ cm}/\mu\text{s}$ (from Figure 4.49), and that a drift time of zero corresponds to a height of 10.8 cm (so $z = 10.8 - 0.1733t$ where t is the drift time). Figure 4.65 is a plot of $(S2/S1) \times y$ vs. drift time (that is, $S2/S1$ corrected for z -dependent light yield) as well as the same plot of uncorrected $S2/S1$ from Figure 4.63 for comparison. The exponential fit to the “corrected” graph yields a lifetime of $253 \pm 17 \mu\text{s}$, which indicates that the z -dependence on the light yield causes us to underestimate the lifetime (for this run) by $\sim 13\%$.

In order to show the behavior of the electron lifetime over time, and how it is affected by recirculation (see section 4.2), we present the results of the $S2/S1$ vs. drift time fit over a much larger sampling of runs than the dataset used in the rest of this analysis. Because these runs were taken at a variety of different conditions, especially electric field settings, there will likely be some systematic effects on the fits, which we have not attempted to account for

here. The runs are divided into two groups; all are ^{22}Na coincidence triggers. The first group is similar to the default set for this memo, but the drift field ranges from 0.5-0.8 kV/cm, and the extraction field from 2.7-3.85 kV/cm. Runs in the second set were taken with much lower drift fields (0.1-0.2 kV/cm), and the trigger acquisition window was extended to 250 or 450 μs , respectively.

The intention of this latter group of runs was to have a more sensitive measurement of the electron lifetime due to having a longer drift time over which to observe the decay. However, because analysis cuts necessarily remove events with very low S2 ($S2 < 10$ p.e.), the mean S2/S1 value is biased upwards at longer drift times. In this way all of the fits tend to overestimate the drift lifetime, but the effect is more pronounced in the runs with longer drift time, as the bias becomes more pronounced as the true average S2 signal gets smaller. Determination of the magnitude of this systematic bias will likely require a Monte Carlo approach, which has not yet been implemented.

Figure 4.66 plots the electron lifetime as determined by the fit as a function of time since the institution of the overnight vigorous recirculation strategy. Standard runs are represented by the black points, while the long drift time runs are shown in red. Vertical error bars are only those reported by MINUIT. The green shading of the background shows periods when active, vigorous recirculation took place, and the blue box marks a period of so-called gentle recirculation, where purified gas was allowed to flow into both the inner and outer volumes, rather than being all forced into the inner volume. At some point during the last long recirculation period shown on this plot (over the beginning of February), a failure of the SAES getter occurred, which resulted in blowing unpurified gas into the inner volume. Unfortunately, the time of this cannot be pinpointed exactly, so the graph shows purification as continuing until the valves were closed. The large drop in lifetime after this period is evidence that the fault likely occurred quite early during the period. Table 4.12 lists the values of the fits for all the runs in the figure, as well as providing the intercept (S2/S1 at zero drift time) and the χ^2/NDF values for each fit.

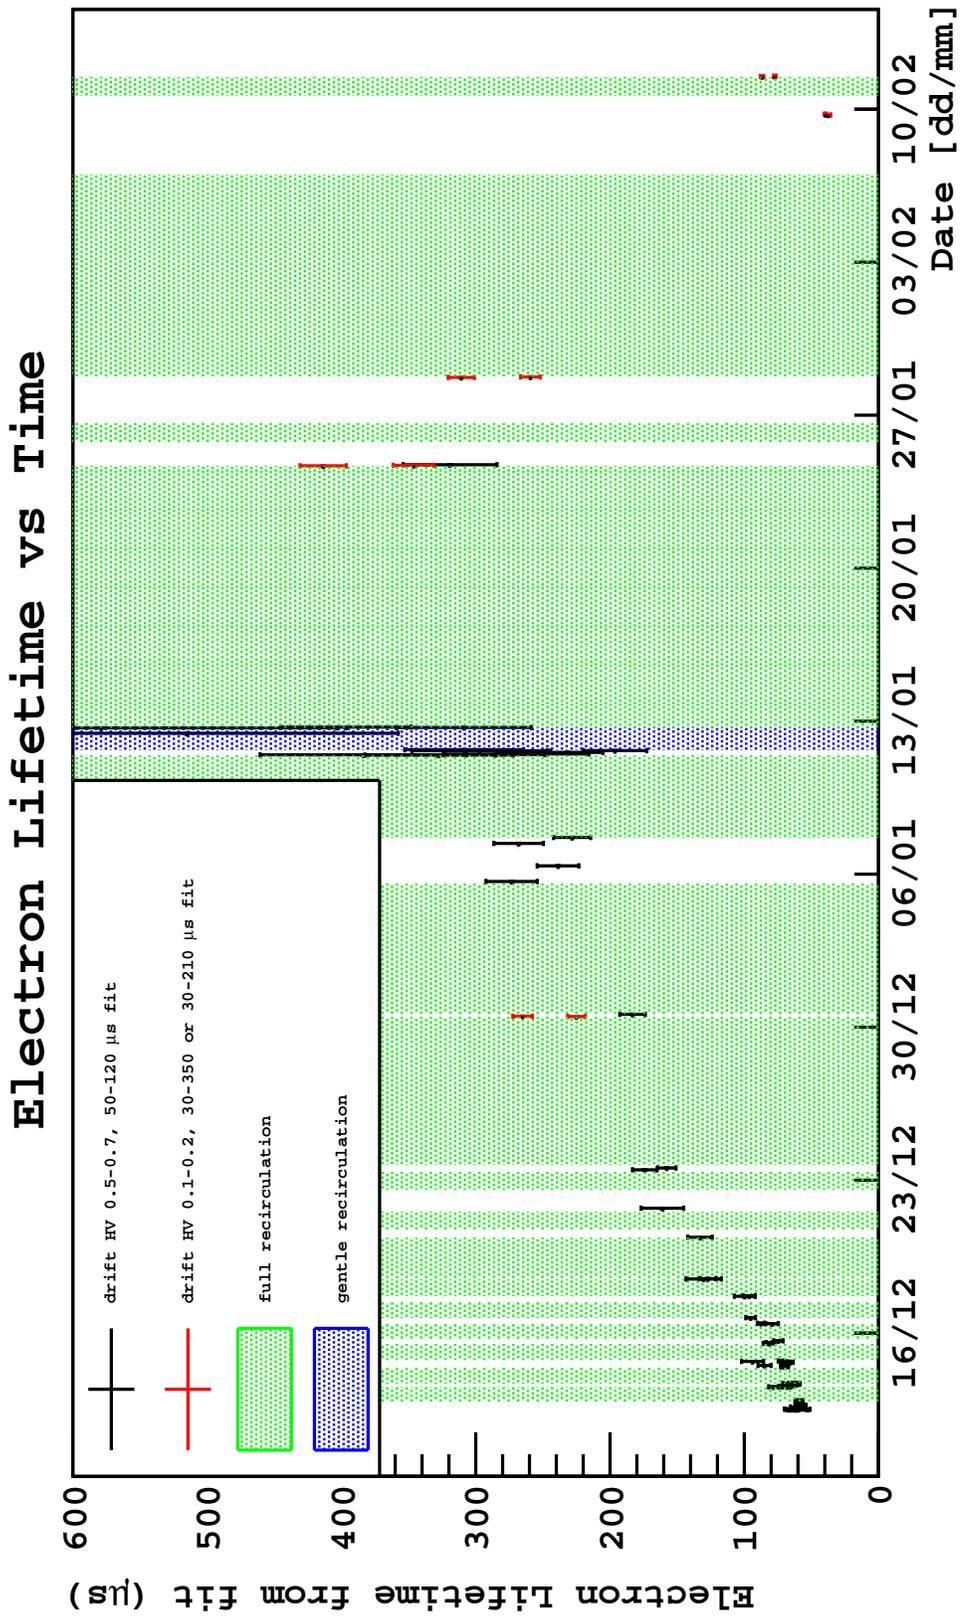


Figure 4.66: Electron lifetime in microseconds from fit as a function of time for multiple runs. See text for details.

Run	Date	Drift HV	Mult. HV	Lifetime(μ s)	Intercept	χ^2 /NDF
Fit over 50-120 μ s						
1179	2010-12-12	0.7	2.7	62 \pm 2.3	11 \pm 0.59	58/50
1180	2010-12-12	0.7	2.7	66 \pm 2.5	9.2 \pm 0.47	78/55
1183	2010-12-12	0.7	2.7	53 \pm 1.9	14 \pm 0.75	92/58
1184	2010-12-12	0.7	2.7	62 \pm 2.8	11 \pm 0.75	41/44
1185	2010-12-12	0.7	2.7	67 \pm 2.9	10 \pm 0.59	75/58
1186	2010-12-12	0.7	2.7	56 \pm 2.6	12 \pm 0.79	71/55
1187	2010-12-12	0.7	2.7	61 \pm 2.4	12 \pm 0.69	59/50
1190	2010-12-12	0.7	2.7	60 \pm 1.3	13 \pm 0.40	70/48
1193	2010-12-12	0.7	2.7	58 \pm 2.1	13 \pm 0.71	78/58
1194	2010-12-12	0.7	2.7	60 \pm 1.2	13 \pm 0.39	47/46
1195	2010-12-12	0.7	2.7	64 \pm 2.2	11 \pm 0.56	98/57
1196	2010-12-12	0.7	2.7	59 \pm 0.65	14 \pm 0.24	161/48
1200	2010-12-12	0.7	2.7	60 \pm 2.0	13 \pm 0.67	56/55
1201	2010-12-12	0.7	2.7	56 \pm 1.9	15 \pm 0.80	45/48
1202	2010-12-12	0.7	2.7	56 \pm 1.6	15 \pm 0.68	68/55
1206	2010-12-12	0.7	2.7	60 \pm 1.8	15 \pm 0.68	68/55
1209	2010-12-12	0.7	2.7	58 \pm 1.0	16 \pm 0.43	95/45
1214	2010-12-12	0.7	2.8	58 \pm 2.2	24 \pm 1.4	81/55
1215	2010-12-12	0.7	2.9	60 \pm 2.0	31 \pm 1.6	66/55
1216	2010-12-12	0.7	3	59 \pm 1.9	45 \pm 2.2	75/50
1225	2010-13-13	0.5	3.2	78 \pm 3.5	42 \pm 2.3	33/49
1226	2010-13-13	0.5	3.2	69 \pm 3.1	47 \pm 2.8	51/49
1227	2010-13-13	0.5	3.2	68 \pm 3.1	46 \pm 2.9	47/49
1228	2010-13-13	0.5	3.2	71 \pm 3.1	45 \pm 2.6	43/50

Continued on next page

Run	Date	Drift HV	Mult. HV	Lifetime(μ s)	Intercept	χ^2 /NDF
1236	2010-13-13	0.5	3.2	63 ± 2.4	48 ± 2.5	83/49
1238	2010-13-13	0.5	3.2	61 ± 2.3	70 ± 3.8	48/44
1239	2010-13-13	0.5	3.2	64 ± 2.6	69 ± 4.0	39/44
1243	2010-13-13	0.5	3.2	68 ± 2.8	63 ± 3.4	72/44
1244	2010-13-13	0.5	2.7	68 ± 3.5	13 ± 0.86	41/44
1247	2010-13-13	0.5	2.7	61 ± 3.0	15 ± 1.1	46/44
1248	2010-13-13	0.5	3.2	64 ± 2.6	70 ± 4.1	62/44
1251	2010-13-13	0.5	3.2	65 ± 2.4	69 ± 3.4	58/45
1260	2010-14-14	0.5	3.2	70 ± 2.9	48 ± 2.5	66/49
1263	2010-14-14	0.5	3.2	71 ± 1.1	50 ± 0.94	62/48
1264	2010-14-14	0.5	3.3	85 ± 4.9	55 ± 3.0	62/49
1267	2010-14-14	0.5	3.3	70 ± 2.9	64 ± 3.4	47/48
1268	2010-14-14	0.5	3.3	71 ± 1.3	64 ± 1.5	58/46
1269	2010-14-14	0.5	3.3	70 ± 3.0	68 ± 3.8	46/47
1273	2010-14-14	0.5	3.3	66 ± 2.6	72 ± 3.8	44/48
1274	2010-14-14	0.5	3.6	69 ± 2.7	130 ± 6.6	69/49
1277	2010-14-14	0.5	3.6	70 ± 1.5	130 ± 3.5	34/44
1278	2010-14-14	0.5	3.85	94 ± 8.3	130 ± 11	120/49
1282	2010-14-14	0.5	3.85	73 ± 2.0	170 ± 6.0	34/47
1288	2010-15-15	0.5	2.7	82 ± 4.3	9.4 ± 0.55	89/49
1292	2010-15-15	0.5	2.7	81 ± 2.4	10 ± 0.32	48/48
1295	2010-15-15	0.5	2.7	75 ± 3.5	10 ± 0.59	57/49
1305	2010-16-16	0.8	2.7	80 ± 4.7	22 ± 1.4	40/45
1308	2010-16-16	0.8	2.7	88 ± 2.6	27 ± 0.83	53/41
1314	2010-16-16	0.8	2.7	95 ± 3.6	11 ± 0.38	83/47

Continued on next page

Run	Date	Drift HV	Mult. HV	Lifetime(μ s)	Intercept	χ^2 /NDF
1328	2010-17-17	0.8	2.7	100 \pm 7.8	12 \pm 0.76	77/55
1330	2010-17-17	0.8	2.7	96 \pm 4.5	13 \pm 0.52	42/42
1336	2010-18-18	0.8	2.7	130 \pm 13	12 \pm 0.76	66/55
1339	2010-18-18	0.8	2.7	127 \pm 5.9	13 \pm 0.39	75/41
1343	2010-20-20	0.7	2.7	133 \pm 9.1	12 \pm 0.53	64/57
1369	2010-21-21	0.7	2.7	160 \pm 16	9.7 \pm 0.52	93/55
1396	2010-23-23	0.5	3.85	174 \pm 9.1	140 \pm 3.9	27/44
1405	2010-23-23	0.6	3.85	158 \pm 6.9	160 \pm 4.1	35/44
1433	2010-30-30	0.7	3	183 \pm 9.5	31 \pm 0.76	218/57
1444	2011-05-05	0.6	3.85	273 \pm 19	150 \pm 3.5	52/55
1450	2011-06-06	0.6	3.85	238 \pm 15.6	164 \pm 4.2	60/53
1457	2011-07-07	0.6	3.85	267 \pm 18	147 \pm 3.5	50/55
1460	2011-07-07	0.6	3.85	228 \pm 14	159 \pm 4.0	45/50
1465	2011-11-11	0.6	3.3	328 \pm 56	49 \pm 2.2	68/53
1468	2011-11-11	0.6	3.3	383 \pm 77	48 \pm 2.2	75/53
1471	2011-11-11	0.6	3.3	285 \pm 39	53 \pm 2.3	61/53
1476	2011-11-11	0.6	3.3	249 \pm 33	57 \pm 2.7	56/53
1483	2011-11-11	0.6	3.3	232 \pm 26	68 \pm 2.9	75/51
1488	2011-11-11	0.6	3.3	297 \pm 50	72 \pm 3.7	47/48
1491	2011-11-11	0.6	2.7	196 \pm 24	15 \pm 0.78	47/46
1498	2011-11-11	0.6	3.3	283 \pm 39	75 \pm 3.3	62/51
1501	2011-11-11	0.6	3.3	309 \pm 44	66 \pm 2.9	62/49
1506	2011-12-12	0.6	3.3	516 \pm 158	69 \pm 3.6	92/48
1511	2011-12-12	0.6	3.3	712 \pm 237	62 \pm 2.8	68/44
1512	2011-12-12	0.6	3.3	577 \pm 174	64 \pm 3.1	56/46

Continued on next page

Run	Date	Drift HV	Mult. HV	Lifetime(μs)	Intercept	χ^2 /NDF
1519	2011-12-12	0.6	3.85	305 ± 49	159 ± 7.9	30/44
1522	2011-12-12	0.6	3.85	889 ± 290	135 ± 3.3	92/47
1523	2011-12-12	0.6	3.85	397 ± 49	152 ± 4.6	61/44
1539	2011-24-24	0.6	3.3	319 ± 35	54 ± 1.7	46/53
Low drift HV runs, fit from 30-210 or 30-350 μs						
1418	2010-12-30	0.1	3	265 ± 7.2	8.3 ± 0.2	108/75
1423	2010-12-30	0.2	3	224 ± 6.5	13 ± 0.2	125/76
1528	2011-01-24	0.1	3.3	414 ± 17	6.6 ± 0.2	108/76
1533	2011-01-24	0.2	3	347 ± 15	11 ± 0.2	8776
1582	2011-01-28	0.1	3	311 ± 9.7	8.5 ± 0.2	85/76
1587	2011-01-28	0.1	3	259 ± 7.3	15 ± 0.2	91/76
1605	2011-02-09	0.1	2.7	38 ± 1.9	2.1 ± 0.3	230/28
1611	2011-02-09	0.2	2.7	40 ± 0.6	3.1 ± 0.1	128/69
1614	2011-02-11	0.1	2.7	87 ± 1.3	2.4 ± 0.1	89/73
1615	2011-02-11	0.2	2.7	77 ± 1.0	4.5 ± 0.1	111/76

Table 4.12: Summary of fit results for electron lifetime for large set of ^{22}Na coincidence runs. Runs in **bold** are the standard run set for this memo. Runs in the first group were all taken with the drift field set between 0.5 and 0.7 kV/cm and fit over the range 50-120 μs . Runs in the second set are either taken with a drift field of 0.1 kV/cm and fit over 30-350 μs , or with a drift field of 0.2 kV/cm and fit over 30-210 μs .

4.6 Conclusions

We have developed algorithms that can competently identify the baseline offset of PMT waveforms, determine the single photoelectron response, identify individual scintillation pulses,

and determine the total amount of light collected, height, pulse shape, and ratio of scintillation to ionization for events in the detector. We focused our analysis on a subset of ^{22}Na coincidence runs with drift fields of 0.6 kV/cm and extraction fields of 3.85 kV/cm. The event acceptance for analysis cuts of this subset was $14 \pm 1\%$. The acceptance for all runs in the entire approximately two month dataset was $24 \pm 6\%$ for runs with fields and $76 \pm 14\%$ for runs without fields. We believe that the majority of rejected events are due to the high background rate of around 500 Hz in the detector; testing the false acceptance/rejection of these cuts will require a study with Monte Carlo events.

The zero-field light yield with all PMTs enabled measured both by 511 keV gammas from ^{22}Na and 356 keV gammas from ^{133}Ba has an average value of 4.5 pe/keV_{ee}; individual runs showed light yields as high as 4.76 pe/keV_{ee}. The width of the 511 keV peak is approximately 4%. The light yield when channels 0, 3, and 6, which showed problems during some runs, are removed is 3.69 pe/keV_{ee}. In the presence of an electric field of 0.6 kV/cm, the gamma light yield is quenched to approximately 0.53 times the zero-field value, compared to an expectation of 0.58. The extra apparent quenching observed is likely due to a decrease in PMT performance due to the high light levels present when the drift and multiplication fields are enabled. If we replace the single 8" PMT on the bottom of the detector with an array of seven 3" PMTs, we expect the light yield to increase by approximately 15%.

The only two major sources of systematic uncertainty in the determination of the energy scale so far identified are the difference between using a moving versus a fixed baseline, which creates an uncertainty of approximately +10% near threshold and +1.5% at 511 keV, and the difference between performing integration of the waveforms before or after summation, which contributes an uncertainty of approximately -2.5% near threshold and negligibly at 511 keV. Uncertainty is estimated as half the difference between the two methods. Our choice of PDF to model the single photoelectron response is non-standard, and results in a light yield estimate that is 3-4% higher than if the conventional PDF were used, but this would not affect the energy scale in terms of, for example, determining threshold energies. Other

sources so far considered have all contributed less than 1% to uncertainty. Several potentially large sources of systematic uncertainty that we have not yet attempted to quantify remain :

- We use a simple function fit over a limited range rather than attempting to model the full energy spectrum for gammas.
- The digitizer samples at 4 ns, which is comparable to the time of a single photoelectron pulse.
- Most events in the 511 keV peak of the ^{22}Na spectrum with zero field, which is the primary determinant of light yield, are slightly outside the digitizer range for channel 7.
- The distributions of drift times for events in several runs are not well understood. These may indicate severe biases in the trigger or cuts, and it is unclear how the light yield might be affected.

We have observed a characteristic exponential decay of the S2/S1 value as a function of drift time between the two pulses, which is interpreted as the free electron lifetime. By recirculating the argon in the detector through a getter, we have measured an increase in the electron lifetime from an initial value of approximately 60 μs to around 400 μs before a failure of the getter occurred. Selection of the drift time range over which to fit the S2/S1 ratio leads to a systematic uncertainty of around 11% in the fit value. There is also an upward bias in the fit due to very small S2 pulses being removed from analysis, unquantified at this time. When using only four of the top PMTs, the few percent spread in effective light yield as a function of source height tends to bias the measurement to smaller lifetimes, estimated on one run to be $\sim 13\%$.

The mean of the pulse shape discrimination parameter $F90$ as a function of energy has roughly the expected distribution when compared to published references, although there is some indication that pulses tend to have a prompt fraction slightly larger than references. This could indicate either a bias due to analysis methods, or a quenching of the slow scintillation component due to impurities. At this time, we are not able to reliably measure directly the slow scintillation lifetime. The distribution of the $F90$ parameter in a given energy bin

is largely consistent with that predicted by photon counting statistics and the width of the PMT single photoelectron response. The measured $F90$ distributions are slightly wider than predicted, however, indicating that there is some small noise component not accounted for in the prediction. The average value of $F90$ does seem to show a dependence on whether the electric fields (and S2 pulses) are present. Although not conclusive, various tests seem to indicate that the difference is indeed a physical effect and not due solely to electronics or analysis.

We have observed discrimination between gamma-like events from a ^{22}Na source and nuclear recoil events from an AmBe source. We have performed tests designed to test whether statistical fluctuations of the S2/S1 and $F90$ discrimination parameters are independent and found no significant evidence against independence in a small sample in the tested energy range of $80 \text{ p.e.} < S1 < 150 \text{ p.e.}$. Combining the two parameters, we find two events out of 45167 ^{22}Na coincidence events that pass all cuts near the nuclear recoil acceptance window. Due to the high neutron rate at surface, it is not possible at this time to determine whether these events are gammas in the tails of the distributions or background neutron events.

For the next analysis, particular attention should be paid to the parameters used in the BaselineFinder and PulseFinder modules, as, at present, these parameters have been chosen without systematic study, and all of the analysis depends on the functioning of these modules. We should make a major effort to develop a method to reliably determine the lifetime of the slow scintillation component.

At the time of this writing, DarkSide-10 has been rebuilt and redeployed in the Laboratori Nazionali del Gran Sasso (LNGS) for a third data collection campaign. Several upgrades have been performed for this effort. The single bottom phototube has been replaced by an array of seven 3" PMTs (identical to the top array), for a total of 14 channels. The top and bottom acrylic windows have been replaced by fused silica, which allows for much thinner ITO layers; the same acrylic cylinder is used for the side walls of the inner vessel. The heat exchanger that cools and liquifies incoming argon gas from the purifier has been redesigned to

have more surface area and therefore deliver more purified liquid argon directly to the inner volume. Finally, the HHV feed-throughs have been redesigned to minimize the breakdowns observed in the second campaign. Finally, 1 m water shield and location in the LNGS cavern should significantly reduce the background. The detector is currently being commissioned; early data looks very promising, but is unfortunately beyond the scope of this work. One of the primary goals of this campaign will be to investigate the behavior of $F90$ as a function of the applied drift field in order to determine the source of the apparent dependence observed in the analysis presented in this chapter.

Overall, these results indicate that DarkSide-10 was very successful as a prototype. With it, we have demonstrated high light yield, electron lifetimes, and environmental stability which are adequate for the operation of a background-free 50 kg TPC. Although we were unable to demonstrate directly the required level of background rejection, this was not one of the design goals for the prototype, and was likely limited by the high background rate at surface. The success of this operation allows us to move forward with the final design and construction of DarkSide-50, our first physics-capable detector sensitive WIMP interactions at current experimental sensitivities.

Chapter 5

Monte Carlo Simulations

In Chapter 2, we outlined in a very general way the classes of background that must be reduced in order to perform a search for WIMP dark matter interactions. Using basic order-of-magnitude arguments, we showed that a two-phase argon time projection chamber (TPC), utilizing argon with a depleted radioactive ^{39}Ar content and employing external active neutron and muon vetoes, should be able to adequately reduce these backgrounds. We presented argon found in underground wells as a source for depleted argon, and introduced the technologies employed to extract and purify it for use in detectors, and showed that early samples of this argon had ^{39}Ar content reduced by at least a factor of 25 with respect to atmospheric argon. In Chapter 4, we outlined some of the technical requirements for a dual-phase argon TPC, and demonstrated that we could achieve several of those requirements by analyzing results from DarkSide-10, a 10 kg prototype. In particular, we demonstrated with the 10 kg prototype:

- Creation of a “pocket” of gaseous argon using a semi-sealed inner vessel. Very stable temperature, pressure, and level for the liquid-gas interface.
- Functioning TPC. Observation of both S2 and S1 signals.
- Low level of impurities in argon. Free electron lifetimes on the order of hundreds of microseconds.

- High light collection efficiency. Zero-field light yield of 4.5 photoelectrons/keV deposited by electron-like signals.

With these results in hand, we are ready to expand DarkSide-10 to a full-scale radiopure shielded TPC, capable of detecting WIMP-induced signals at the current limits of experimental sensitivity with zero background. Section 5.1 will describe the detector and shielding for DarkSide-50, the first generation of a series of WIMP-sensitive DarkSide detectors, which will be constructed in 2011-2012. The shielding for DarkSide-50 is designed to be large enough to accommodate an upgrade of the inner detector to DarkSide-5T, which will contain a 5 ton target mass. Particular attention will be paid to the radioactive backgrounds, both internal and external, that will be present. In the remainder of this chapter, we will revisit the question of background reduction and provide quantitative estimates of the remaining backgrounds to a WIMP search using Monte Carlo methods. Section 5.2 will describe the methodology employed for these simulations, and Section 5.3 will show the simulation results and projected sensitivities for DarkSide-50.

5.1 Full-scale DarkSide detectors: DarkSide-50

Figure 5.1 shows the proposed layout of the DarkSide-50 inner detector, neutron veto, and muon veto [106]. The basic layout of the inner detector is similar to DarkSide-10. The active volume will contain a target mass of 50 kg of depleted liquid argon, approximately 36 cm high by 36 cm diameter. Rather than the acrylic vessel used in DarkSide-10, the top and bottom windows of the inner vessel will be constructed of fused silica with evaporated ITO forming the anode and cathode of the drift and multiplication fields, and the side walls constructed from PTFE (Teflon). Both PTFE and fused silica are much less brittle than acrylic at cryogenic temperatures, which will simplify some of the mechanical constraints. The added mechanical stability also allows for much thinner ITO layers than acrylic (~ 15 nm compared to the ~ 100 nm used in DarkSide-10). PTFE and fused silica also outgas less than

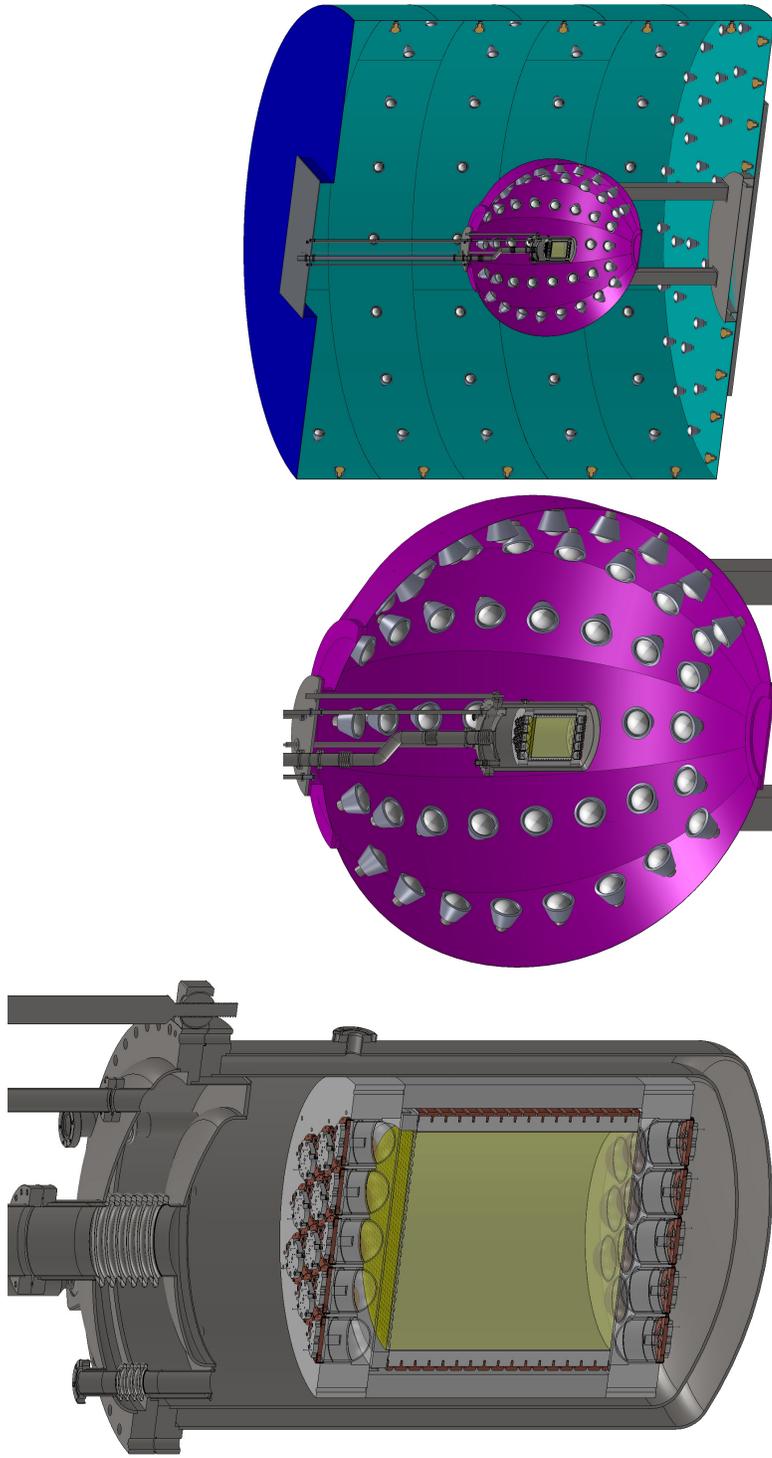


Figure 5.1: The planned layout of the DarkSide-50 detector. Left: Inner detector (shown instrumented with QUPIDS), support structure, and cryostat. Center: Inner detector inside the active neutron veto. Right: Neutron veto placed inside the CTF tank, which will serve as cosmic muon veto.

acrylic, leading to less contamination of the argon. They also contain no significant amounts of hydrogen, so they will attenuate neutrons less than acrylic, which aids the effectiveness of the neutron veto. The top fused silica window will be joined to a short cylindrical lip to form a “diving bell” that will hold the gas pocket for S2 generation. The gas will be maintained by boiling some of the argon in an exterior arm connected to the gas region as in DarkSide-10. A 100 μm thick stainless steel or copper grid a few mm below the liquid surface will allow the discontinuity between drift and extraction/multiplication fields.

The first run of DarkSide-50 will use two hex-packed arrays of 19 Hamamatsu R11065 3 inch diameter PMTs (the same used for the top face of DarkSide-10) above and below the active volume outside the fused silica windows. The PMTs will be mounted to a ~ 1 cm thick copper plate and held in place by a spring-tightened “hold-down” plate similar to the mounting scheme in DarkSide-10. The PTFE active volume side walls will extend above and below the windows and provide mechanical support for the whole assembly. The inner vessel will be immersed in a buffer volume of ~ 50 additional kilograms of depleted liquid argon, and the whole thing contained in a vacuum-walled stainless steel Dewar. In a possible second phase of DarkSide-50, the PMTs will be replaced with QUPIDs, hybrid photodetectors developed at UCLA that detect photons by focusing and accelerating photoelectrons emitted from a hemispherical photocathode onto an avalanche photodiode (APD) [107]. The QUPIDs have similar quantum efficiencies to the R11065 PMTs, but better signal resolution and significantly lower radioactive backgrounds, and will be one of the key technologies utilized in future DarkSide detectors. However, as of summer 2011, only a few prototypes are available, so the first phase of DarkSide-50 will run with PMTs.

The cooling system will be removed outside the Dewar, but purified, condensed liquid argon will still be delivered directly to the active volume. The Dewar will be contained inside the active neutron veto, a 4 meter diameter steel sphere filled with a 50% w/w mixture of pseudocumene (PC) and trimethyl borate (TMB), doped with a small (~ 1 g/L) concentration of diphenyl oxylate (PPO), a fluorescent compound that shifts the ultraviolet PC

scintillation emission into visible wavelengths. The veto will be viewed by ~ 100 8" PMTs mounted to the inside of the sphere. The PC/PPO mixture is the scintillator employed by the Borexino experiment and has a light output of ~ 11000 photons/MeV [108]. According to studies conducted by the collaboration, the TMB-loading should not significantly reduce the scintillation output. The high neutron cross section of boron will, however, greatly decrease the neutron attenuation length and capture times relative to pure scintillator, increasing its effectiveness as a neutron veto. The expected neutron detection performance of the veto has been extensively studied through simulation [83].

Finally, the neutron veto will be inserted into the Borexino Counting Test Facility (CTF) tank [109], a steel cylinder roughly 10 m high by 11 m diameter water tank. The inside of the CTF tank is instrumented with PMTs to detect the Cherenkov light from cosmic muons passing through the water and so serves as an active muon veto in addition to the passive shielding granted by the thickness of the water layer.

As mentioned, the DarkSide-50 neutron veto is designed to accommodate a future upgrade with a 5 ton target. This will increase the size of the active volume to ~ 165 cm height and diameter, with ~ 400 QUPIDs each above and below. This will leave approximately 1 m thickness of neutron veto around the enlarged cryostat.

5.2 Simulation Details

5.2.1 Monte Carlo Software

To simulate the passage of the neutrons and gammas through the various layers of shielding and measure their interactions in the active volume, we use the **GEANT4** Monte Carlo framework [110]. A **GEANT4** simulation requires the user to specify the detector geometry (size, location, density, and composition of each component) and a “primary” particle to simulate (particle type, initial position and momentum). In large part, the accuracy of the results obtained in a simulation depends on the accuracy with which the detector geometry

is specified. There is a tradeoff, however, as specifying every single small component greatly increases the necessary time involved in writing the simulation as well as the processing time for each simulated event, often without significantly affecting the result.

Since the exact geometry of the DarkSide detectors has not yet been finalized, the background estimation can be at best an approximation. For this, we use the program `s4det`, which is an implementation of `GEANT4` developed at Princeton. `s4det` was developed to rapidly test the backgrounds for different liquid argon TPC detector and shield configurations in different underground sites (initially for the DUSEL S4 proposal period, from which the name derives). All of the setup for a simulation is specified in simple text configuration files.

Detector geometry is greatly simplified by assuming that each detector component is a simple cylindrical volume. The geometry is most precisely specified near the active volume. The material composition of each volume can also be changed via the configuration files. The active volume is a cylinder of liquid argon whose height and diameter can be specified, below a cylindrical gas region whose height can be defined. Each array of PMTs, since they have little total mass and therefore negligible effect on particle tracking, is approximated as a single empty (vacuum) cylinder, above and below the active volume. The metal mounting plate is also defined; both its thickness and the length of the PMTs can be specified. The active volume itself is nested inside a container with thickness that can be specified separately for top, bottom, and sides. (In DarkSide-10, this was the acrylic inner vessel; for DarkSide-50, there are only fused silica windows at the top and bottom, so the side thickness is set to 0.) This container, the PMTs, and hold-down plate are all nested inside a support structure with variable thickness. The support structure, and each other volume outside it, is nested at the center of the containing volume, and the user only has to specify the thickness of each layer. The support structure is immersed in a buffer volume (usually liquid argon) contained in a Dewar (modeled as a single wall, ignoring the vacuum insulation). Finally, outside the Dewar, the user can specify any number of shield layers by specifying only a thickness and

Layer	Thickness cm	Material	Approx. Mass kg
Active Volume	18	liquid argon	51.3
Gas Pocket	2	gas argon	0.01
Inner vessel windows	2	fused silica	5×2
PMTs	8	vacuum	0
Metal plate	1	copper	10×2
Inner vessel support	6	PTFE	110
Buffer	2	liquid argon	40
Dewar	2	stainless steel	245
Neutron veto	160	50% w/w PC + TMB	38,000
Veto tank	2	ship steel	11,000
Muon veto	300	water	700,000
CTF tank	2	ship steel	72,000

Table 5.1: Simulation geometry input to `s4det` for DarkSide-50. The volumes are illustrated in Figure 5.2.

material. If one wants to simulate an experimental hall or cavern rock, these can be specified as additional “shield” layers.

Table 5.1 reports the configuration settings input to `s4det` to define the DarkSide-50 detector and shielding. Figure 5.2 illustrates the arrangement of the volumes. The largest approximation to the current plans is that the neutron veto is modeled as a cylinder rather than a sphere.

The greatest utility of `s4det` is in specifying the background sources. Any of the particle types defined by `GEANT4` can be added to any given volume by name, either distributed evenly throughout the volume or evenly over the outside surface (this is one reason volumes are restricted to simple cylinders). Both the spectrum and angular distribution of each source can be defined in text or `ROOT` files. Background rates can be specified absolutely in Bequerel, or in Bq/kg for volume distributions or Bq/cm² for surface distributions, and will be automatically scaled by the total mass or surface area of the appropriate volume. The total live time covered by each simulation is then automatically calculated from the specified rates and total number of primaries generated.

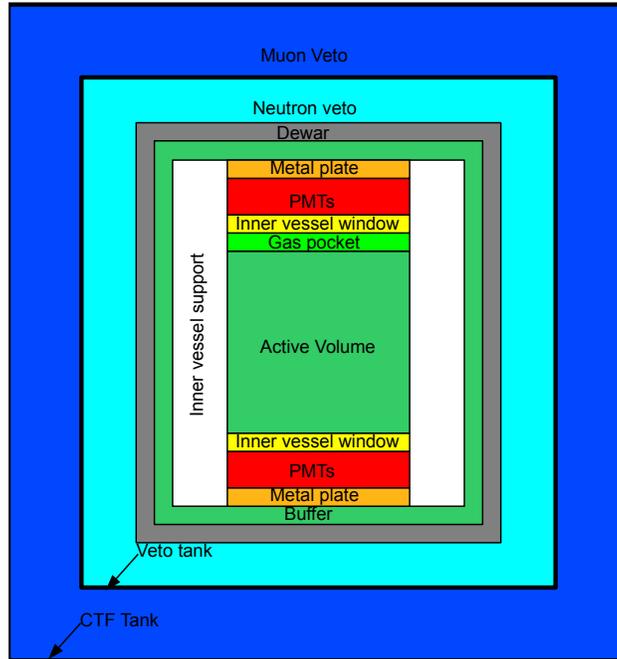


Figure 5.2: Cartoon representation of the detector volumes specified in `s4det`. Not to scale.

In addition to fundamental particles, `GEANT4` can also simulate nuclear decay chains via the `G4RadioactiveDecay` process. This is particularly useful in situations when secular equilibrium can be assumed; the entire decay chain can be simulated by spawning the parent isotope. (For a decay chain in secular equilibrium, the rate of decay for each isotope in the chain is equal to that of the parent isotope.) By default in `GEANT4`, all of the decays in the chain will be included in a single `GEANT4` “event,” even though the individual decays might be separated by hours (or years!). In `s4det`, any radioactive isotope created with a half-life longer than a specified value (default $2 \mu\text{s}$), once it has dissipated any kinetic energy imparted from the decay, is removed from that event and becomes the primary particle for the subsequent event. For the automatic livetime calculation, only parent isotopes are tallied, so the pushed-back daughters do not cause an overestimation of the simulated livetime. This removal of long-lived isotopes also occurs when they are created by processes other than radioactive decay, such as neutron capture. In this way we ensure that we do not over count the energy deposited during a single event by including decay energies that actually happen much later.

Isotope or chain	Gammas per equilibrium decay	Value used in simulations
^{238}U	2.2	5
^{235}U	2.8	5
^{232}Th	2.7	5
^{210}Pb	0.04	0.1
^{40}K	0.11	0.2

Table 5.2: The average number of gammas emitted from the entire chain for each equilibrium decay of the listed parent isotope. To be conservative, we simulate approximately twice the measured rate, which is given in the third column.

Gamma spectra are simulated in three different ways. For the volumes inside the detector cryostat, the parent isotope for each of the decays or chains (^{238}U , ^{235}U , ^{232}Th , ^{40}K , ^{210}Pb) is spawned, and the decays are modeled by the G4RadioactiveDecay process. This method will produce the correct spectrum not only of single decays, but also “cascades,” where multiple gammas are emitted simultaneously, and Bremsstrahlung X-rays. Although this is the most exact method, it is also quite slow, as all decay products and daughters are simulated, even ones, such as alpha decays, that cannot reach the active volume. Therefore, in the shield layers outside the cryostat, we directly simulate the gamma products. rather than the whole decay chain. To determine the rate and spectrum of gammas produced from each parent isotope, we run a separate simulation of each chain, and tally the gammas produced in the G4RadiocativeDecay process. Table 5.2 lists the average number of gammas emitted for each isotope, assuming secular equilibrium. In order to ensure that the simulation is conservative, we assume approximately twice the rate obtained by this method; the actual rate assumed is also reported in Table 5.2. The gamma spectra for the ^{238}U , ^{235}U , and ^{232}Th chains obtained by this method are presented in Figure 5.3. ^{210}Pb and ^{40}K produce monoenergetic gammas at 47 keV and 1.46 MeV, respectively. Finally, for external gammas, we make the very conservative assumption that the entire flux is at the highest natural decay energy, the 2.6 MeV ^{208}Tl line. This is conservative because these high-energy gammas are more penetrating than most of the true spectrum. The extra penetration effect is larger than the reduction of the rate from energy cuts.

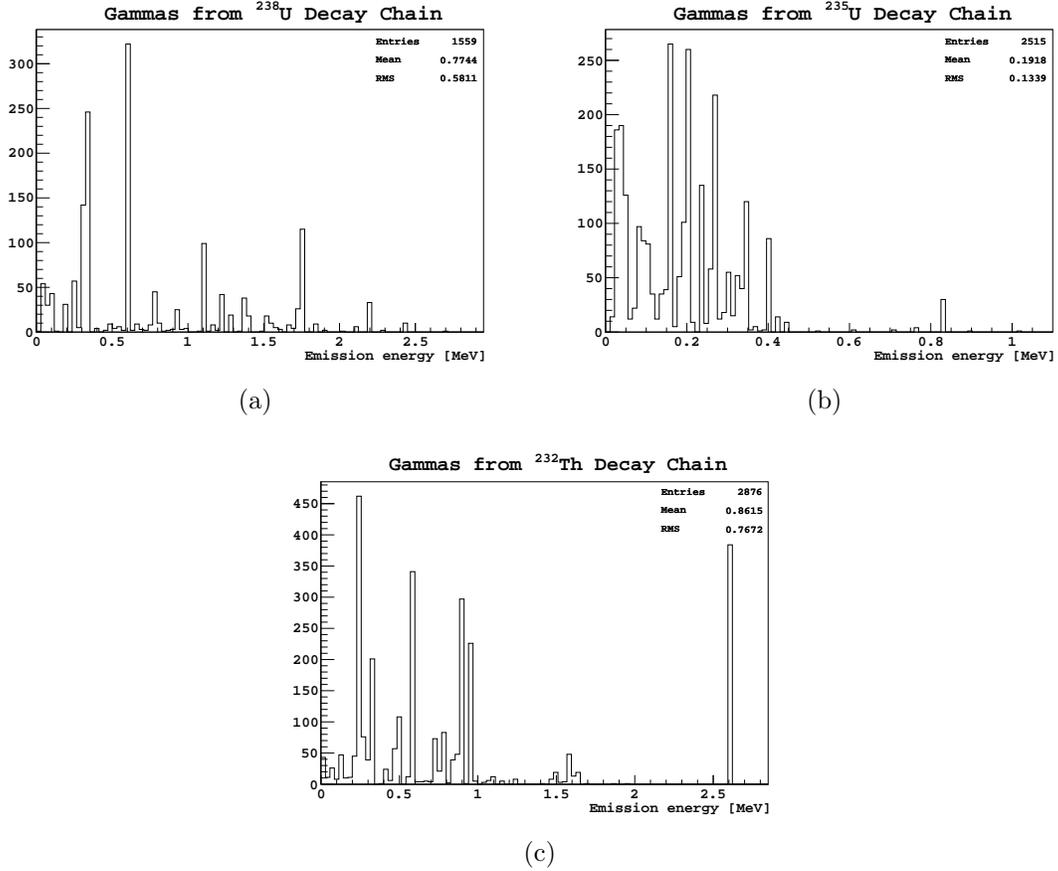


Figure 5.3: Spectrum of gammas emitted for each equilibrium decay of ^{238}U (a), ^{235}U (b) and ^{232}Th (c), obtained by simulating the decays of the parent isotope in GEANT4. ^{210}Pb and ^{40}K produce monoenergetic gammas at 47 keV and 1.46 MeV, respectively.

Radiogenic neutrons, i.e., those emitted by spontaneous fission or (α, n) processes, can also be modeled. For a given material composition with specified ^{238}U , ^{235}U , and ^{232}Th contaminations, we determine the rate and spectrum of radiogenic neutrons using the program SOURCES4A [111].

Cosmogenic neutron fluxes can be added in `s4det` by specifying only the site depth in km of water equivalent (km.w.e.) (normalized to a flat overburden) following the parametrizations of Mei and Hime [5]. All of the parametrizations provided by Mei and Hime were generated by Monte Carlo simulation (using the FLUKA framework) of propagation of cosmic ray muons, for which the flux and spectra are well measured, in large (20 m cubed) rock volumes. According to their results, the total neutron flux emitted from the cavern walls is

Depth (km.w.e.)	a_0	a_1	a_2	a_3
1.59	6.86	2.10	2.97×10^{-13}	2.46
1.95	7.33	2.11	-5.35×10^{-15}	2.89
2.05	7.55	2.12	-1.26×10^{-14}	2.76
2.81	7.88	2.21	-2.34×10^{-14}	2.61
3.00	7.83	2.23	-7.51×10^{-15}	2.83
6.01	7.77	2.13	-2.94×10^{-16}	2.86

Table 5.3: Fit parameters as a function of site depth used in Equation (5.2) [5]

assumed to be

$$\phi_n = 4.0 \times 10^{-7} \left(\frac{0.8}{h_0} \right) e^{-h_0/0.86} \text{ cm}^{-2}\text{s}^{-1}, \quad (5.1)$$

where h_0 is the depth in km.w.e. The spectrum of cosmogenic neutrons is given by

$$\frac{dN}{dE_n} = A_\mu \left(\frac{e^{-a_0 E_n}}{E_n} + B_\mu(E_\mu) e^{-a_1 E_n} \right) + a_2 E_n^{-a_3}, \quad (5.2)$$

where A_μ is a normalization constant, $B_\mu(E_\mu)$ is a function of the mean muon energy:

$$B_\mu(E_\mu) = 0.234 - 0.6441 e^{-0.014 \langle E_\mu \rangle}, \quad (5.3)$$

with $\langle E_\mu \rangle$ given by

$$\langle E_\mu \rangle = 391.5(1 - e^{-0.4h_0}) \text{ GeV}. \quad (5.4)$$

The a parameters at several different depths are given by Mei and Hime (reproduced in Table 5.3); values for other depths are obtained by linear interpolation. Equation (5.2) is only valid above ~ 10 MeV, and in practice we ignore the component of the flux below that energy, since, except for very shallow sites, it is significantly less than the radiogenic flux.

The angular distribution of neutron directions is highly correlated with the parent muon:

$$\frac{dN}{d \cos \theta} = \frac{A}{(1 - \cos \theta)^{0.482 \langle E_\mu \rangle^{0.045}} + 0.842 \langle E_\mu \rangle^{-0.152}}, \quad (5.5)$$

where θ is the angle between the direction of the neutron and its parent muon, and A is a normalization constant. The muon angular distribution is itself directed primarily downward:

$$\frac{d\mu}{d \cos \theta_z} = A \frac{(8.6 \times 10^{-6})e^{-h_0/0.45 \cos \theta_z} + (4.4 \times 10^{-7})e^{-h_0/0.87 \cos \theta_z}}{\cos \theta_z}, \quad (5.6)$$

where θ_z is the zenith angle (relative to straight down) and A is again a normalization constant. The average angular distribution of neutrons relative to laboratory coordinates is generated by convolving Equation (5.5) with Equation (5.6) by a simple Monte Carlo.

Muons that pass through detector materials also generate cosmogenic neutrons. The neutron production rate per through-going muon for different elements is given by [5]

$$\langle n \rangle = 1.27 \times 10^{-4} \left(\frac{Z^2}{A} \right)^{0.92} n / (\mu \text{ g cm}^{-2}), \quad (5.7)$$

where the muon flux is given by

$$\phi_\mu = (67.97 \times 10^{-6})e^{-h_0/0.285} + (2.071 \times 10^{-6})e^{-h_0/0.698} \text{ cm}^{-2}\text{s}^{-1}. \quad (5.8)$$

The spectrum and angular distribution used for internal cosmogenics is the same as for externals.

For both the environmental cosmogenic neutron flux and the cosmogenic neutrons produced inside detector materials, we simulate only a single neutron per event, at the total production rate. This is in general extremely conservative, since most neutrons are generated in showers with other neutrons and/or charged particles, which greatly increase the chances of vetoing the event. We also do not simulate the parent muon; in reality, for many neutron events, the parent muon will pass through some part of the veto and allow us to veto the event. Other DarkSide collaborators are exploring the process of tracking entire showers to estimate the fraction of events that can be vetoed in this way.

5.2.2 Radioactive Background Rates and Spectra

Internal Sources

Internal backgrounds refer to those emitted from the detector materials due to contamination with radioactive isotopes, the most common of which are the uranium and thorium chains and ^{40}K . Cesium, cobalt, and manganese are other common sources of decays, but the rates are generally less than the uranium, thorium, and potassium chains. Figure 5.4 shows the decay chains for ^{238}U , ^{235}U , and ^{232}Th . In principle, to accurately predict the rate of backgrounds emanating from a material, the individual rates of each of the daughters in the chain must be measured. In practice, we usually assume that all of the products of the chain are in secular equilibrium; that is, the rate of decays of each of the daughters is equal to the parent rate. This assumption is usually valid for the entire ^{232}Th and ^{235}U chains, and the ^{238}U chain up to ^{210}Pb . The reason for the break is ^{222}Rn , which is a noble gas, and therefore highly diffusive. Unless special effort is taken to remove radon from a manufacturing environment, increased concentrations of ^{222}Rn and its daughters can build up in materials. These daughters all quickly decay away until ^{210}Pb , which has a half-life of 22.6 years [112]. The ^{210}Pb then attains secular equilibrium with its daughters, ^{210}Bi and ^{210}Po . The other isotopes of radon do not have any long-lived daughters in their decay chains, and so any contaminations due to them quickly decay away.

Radioactive contaminations can be present both on the surfaces of materials (particularly radon daughters) and in the bulk. For our purposes, we can generally assume that the rate from bulk contamination is much higher than the surface rate for most materials, except possibly for the very “cleanest” materials, such as PTFE, fused silica, and OFHC copper, that are close to the active volume. On the other hand, these pieces will be the most carefully prepared, and so their surface rates should also be much smaller than bulk rates. The one exception to this general rule of thumb is surface contamination on the inside surface of the active volume, which will generate alphas that can mimic nuclear recoils. However, as

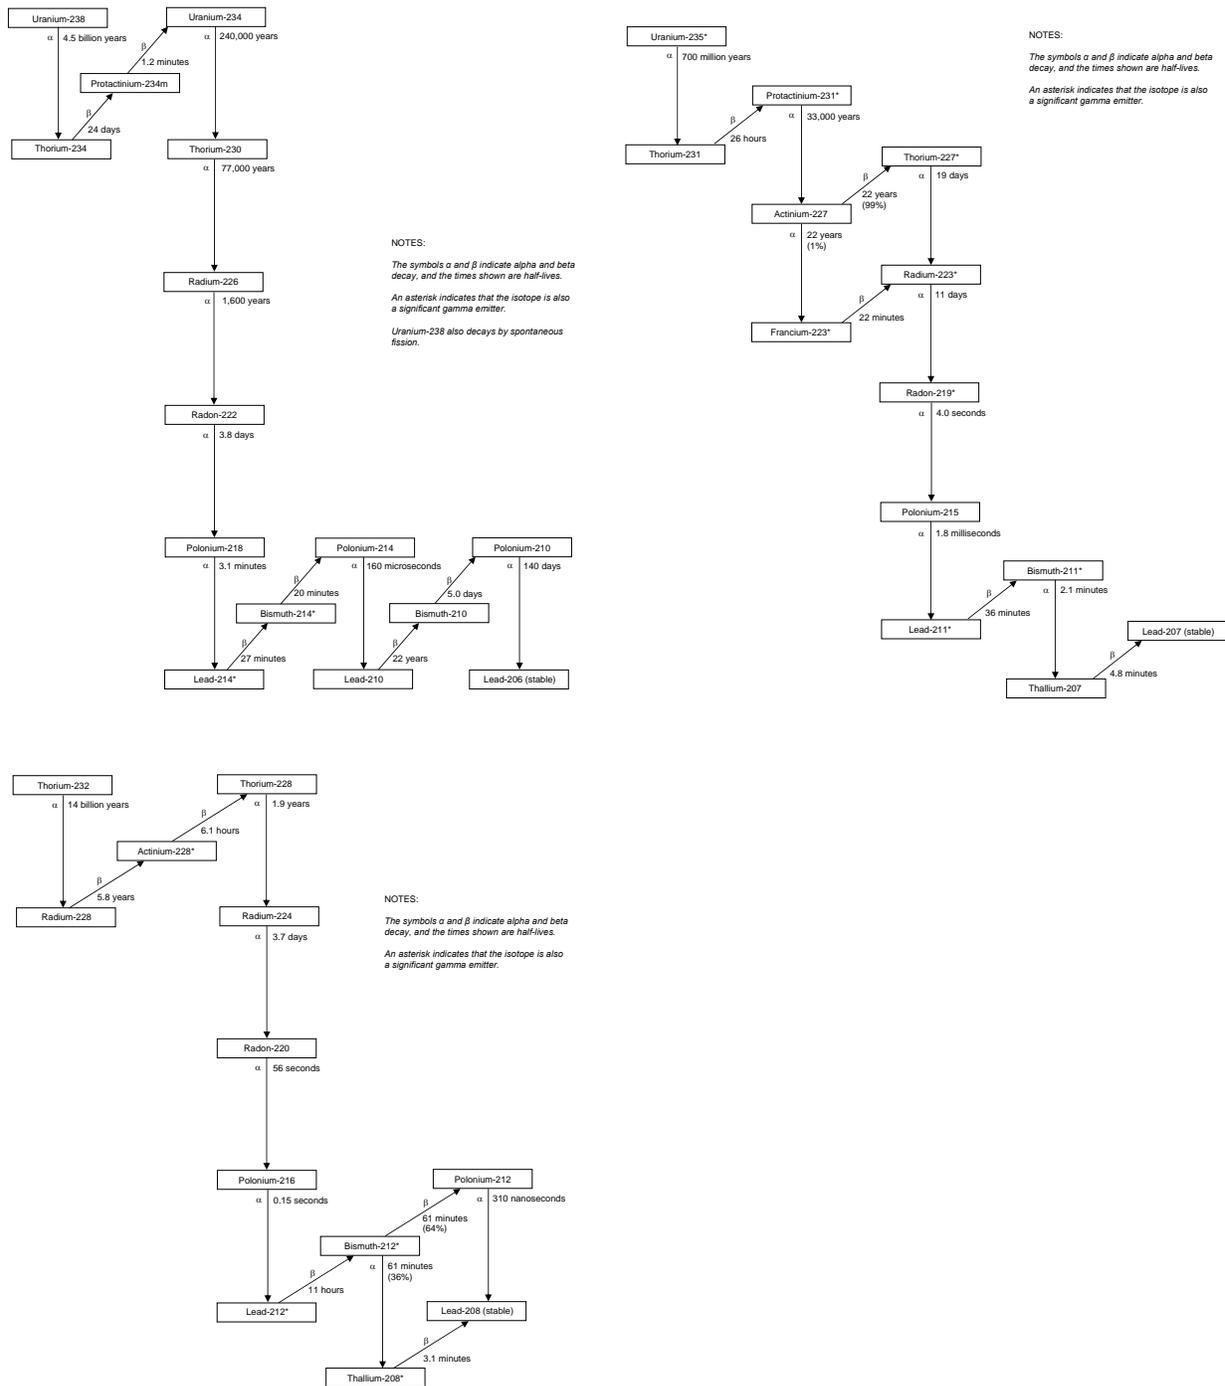


Figure 5.4: Decay chains for ^{238}U , ^{235}U and ^{232}Th , the most common contributors to intrinsic radioactive backgrounds [12].

discussed in Chapter 2, these surface events can generally be removed by reconstructing their position and taking a fiducial cut.

Table 5.4 lists the assumed bulk decay rates per unit mass for the materials that will be used in the construction of DarkSide-50. The ^{210}Pb rate is presented separately from ^{238}U where measurements exist. For convenience, the conversion rate between concentration and decay rate of the common parent isotopes is given in Table 5.5. The rates quoted in Table 5.4 are actual measurements made by various groups, but obtaining materials with those low levels will require a careful survey of sources before beginning construction.

In several of the materials listed in Table 5.4, in particular PTFE, the reported ^{210}Pb decay rate is much higher, sometimes by orders of magnitude, than the ^{238}U rate. Because direct measurements of the ^{210}Pb rate are not as commonly made as ^{238}U , ^{232}Th , and ^{40}K rates, it is possible that the high ^{210}Pb rates were measured in different, “dirtier” samples; unfortunately this question is not addressed in the literature. Because of this large uncertainty and its impact on the final results, in the remaining analyses we will provide two numbers, one assuming the reported ^{210}Pb rates, and one assuming ^{210}Pb in secular equilibrium with ^{238}U . We will see in Section 5.3 that the distinction is only relevant for the PTFE of the inner vessel support structure. By careful control of the manufacturing process, in particular the reduction of radon, it should be possible to obtain PTFE with ^{210}Pb rate approximately in secular equilibrium with the ^{238}U chain (although it may be prohibitively costly, and require cooperation with the manufacturer). We therefore adopt the equilibrium rate for our default estimate, and provide the higher rate as a conservative limit.

Most of the isotopes listed above decay by emission of an electron (beta decay) or an alpha particle (alpha decay), and some of the beta decays are accompanied by gamma ray emission. Of these decay products, only the gammas interest us directly, since the beta and alpha particles do not in general have enough range to escape the parent volume in which the decay occurred. Some higher-energy beta decays close to the active volume can be detected,

Component	^{238}U (Bq/kg)	^{235}U (Bq/kg)	^{232}Th (Bq/kg)	^{210}Pb (Bq/kg)	^{40}K (Bq/kg)	Comments	Ref
PTFE (Dupont Teflon)	$<9.6 \times 10^{-6}$	$<4.4 \times 10^{-7}$	$<1.1 \times 10^{-6}$	~ 0.20	$\sim 1.0 \times 10^{-4}$	Cs, Co	[113, 114]
Copper (OFHC)	$<3.6 \times 10^{-5}$	$<1.6 \times 10^{-6}$	$<9.8 \times 10^{-6}$	~ 0.30	$<1.6 \times 10^{-3}$	Cs, Co	[113, 114]
Fused Silica (Suprasil)	$\sim 2.5 \times 10^{-3}$	$\sim 1.1 \times 10^{-4}$	$\sim 8.1 \times 10^{-4}$	<0.30	$\sim 9.3 \times 10^{-4}$	Cs, Co, Mn	[114]
Stainless Steel (304SS)	<0.002	$<9 \times 10^{-5}$	<0.001		<0.004		[114]
Ship Steel	$<7.4 \times 10^{-4}$	$<3.4 \times 10^{-5}$	~ 0.0011		<0.013	Mn	[115]
TPB	<0.026	~ 0.0012	<0.023		~ 0.41		[116]
Boron Loaded PC	$\sim 1.9 \times 10^{-8}$	$\sim 8.8 \times 10^{-10}$	$\sim 1.2 \times 10^{-10}$	$<1.0 \times 10^{-4}$	$<3.1 \times 10^{-10}$		[117, 118]
Water	$<8.0 \times 10^{-6}$	$<3.7 \times 10^{-10}$	$\sim 2.8 \times 10^{-8}$		$\sim 3.1 \times 10^{-7}$		[119]
Argon	(Bq/tube)	(Bq/tube)	(Bq/tube)	(Bq/tube)	(Bq/tube)		
3" QUPID	$<4.9 \times 10^{-4}$	$<2.3 \times 10^{-5}$	$<4.0 \times 10^{-4}$		$<2.4 \times 10^{-3}$	Co	[120]
R11065 PMT	~ 0.0147	$\sim 6.8 \times 10^{-4}$	$\sim 6.8 \times 10^{-4}$	~ 0.0157	~ 0.0420	Co	[121]
8" PMT	~ 0.43	0.02	~ 0.020	~ 0.13	~ 0.0420		[121, 122]

Table 5.4: Assumed rates of radioactive decays of the most common isotopes for the materials used in DarkSide TPC and shielding construction. For the uranium and thorium chains, the rate specified is for *each* step in the decay, assuming secular equilibrium. Very few direct measurements of the ^{235}U chain rates are available, and so the rates given here are all assuming the natural abundance, or $\sim 4.6\%$ of the ^{238}U rate. For argon, the decay rate due to contaminants should be negligible compared to the ^{39}Ar decay rate. For references where separate measurements of ^{210}Pb are not available, we assume secular equilibrium with the ^{238}U chain. Where a measurement rather than a limit was reported (denoted by $<$), the table reports the $+1\sigma$ rate. These numbers were compiled by A. Wright and P. Mosteiro.

Concentration	Decay rate
1 ppb ^{238}U	= 0.0123 Bq/kg ^{238}U
1 ppb ^{232}Th	= 0.00407 Bq/kg ^{232}Th
1 ppm $^{\text{nat}}\text{K}$	= 0.031 Bq/kg ^{40}K

Table 5.5: Conversion between decay rate and concentration for ^{238}U , ^{232}Th , and ^{40}K [6].

either by the beta itself reaching the active volume or by the production of secondary x-rays through Bremsstrahlung.

Although the alphas do not directly generate backgrounds, they can produce neutrons through (α, n) reactions on light nuclei. These neutrons have long range and cannot be removed using the discrimination capabilities of argon, and so make up the dominant background to a WIMP search. The total (α, n) cross section depends both on the energy of the alpha particle and the reacting isotope, and so the estimated (α, n) production rates will differ both for different decay chains and different materials. Table 5.6 gives the (α, n) rate for various detector materials per equilibrium decay of each of the alpha-producing decay chains, as calculated using the SOURCES4A code [111]. (So, the rate of (α, n) decays in a given material is the product of the numbers in Tables 5.6 and 5.4 for each chain). Some of the heavier elements can also decay by spontaneous fission, which releases one or more neutrons directly. Table 5.7 gives the fractional rate of spontaneous fission for each of the decay chains, also obtained from SOURCES4A.

As one of the primary background sources, the neutron emission rate for the PMTs has been studied in detail; we expect 3.4 n/y/PMT for the 3" Hamamatsu R11065 PMTs and, conservatively, ~ 300 n/y/PMT for 8" PMTs [82]. QUPIDs, with much lower intrinsic backgrounds, should emit < 0.05 n/y/QUPID.

Cosmogenic neutrons are generated for every volume outside the Dewar. Within the Dewar, we assume that we can detect the parent muon with essentially 100% efficiency. For the volumes in which cosmogenic neutrons are generated, the rate is determined from the site depth at LNGS (3.1 km.w.e normalized to a flat overburden [5]) and the material composition, according to Equation (5.7). Table 5.8 lists the assumed composition and

Material	^{238}U	^{235}U	^{232}Th	^{210}Pb
Copper	2.9×10^{-8}	1.5×10^{-8}	2.5×10^{-7}	0
Stainless Steel	2.7×10^{-7}	2.2×10^{-7}	1.0×10^{-6}	7.0×10^{-11}
Ship Steel	3.9×10^{-7}	3.2×10^{-7}	1.5×10^{-6}	5.0×10^{-10}
10% Boron PC	7.3×10^{-6}	8.4×10^{-6}	7.7×10^{-6}	8.8×10^{-7}
TPB	1.2×10^{-6}	1.6×10^{-6}	1.6×10^{-6}	1.0×10^{-7}
PTFE	5.7×10^{-5}	7.4×10^{-5}	7.3×10^{-5}	4.9×10^{-6}
Fused Silica	1.2×10^{-6}	1.6×10^{-6}	1.8×10^{-6}	8.9×10^{-8}
Water	4.3×10^{-7}	5.3×10^{-7}	2.0×10^{-10}	4.7×10^{-8}

Table 5.6: Rate of (α , n) reactions for different detector materials as a fraction of total decay rate for the isotopes listed. For uranium and thorium, this rate includes equilibrium decays of the entire chain along with the parent. Rates were calculated using SOURCES4A [111] by P. Mosteiro and A. Wright.

Decay Chain	Fission neutrons per equilibrium decay
^{238}U	1.10×10^{-6}
^{235}U	3.75×10^{-9}
^{232}Th	3.03×10^{-11}

Table 5.7: Rate of fission neutrons emitted per equilibrium decay of uranium and thorium, calculated from SOURCES4A.

resulting cosmogenic neutron rate per kilogram of material for the shielding materials. The spectrum and angular distributions are determined from Equations (5.2), (5.5), and (5.6). Figures 5.5 and 5.6 show the resulting spectrum and angular distribution obtained for LNGS depth. Note that this assumes a flat overburden, which is not true of LNGS, so the angular distribution is only very approximate.

Material	Composition	Cosmo. Neutron rate n/y/kg
Ship steel	99.8% Fe, 0.2% C	0.65
Water	88.9% O, 11.1% H	0.22
Borated scintillator	62.3% C, 23.1% O, 9.4% H, 5.2% B	0.17
Stainless steel	68%Fe, 20% Cr, 12% Ni	0.65

Table 5.8: Rate of cosmogenic neutrons produced in detector shielding materials.

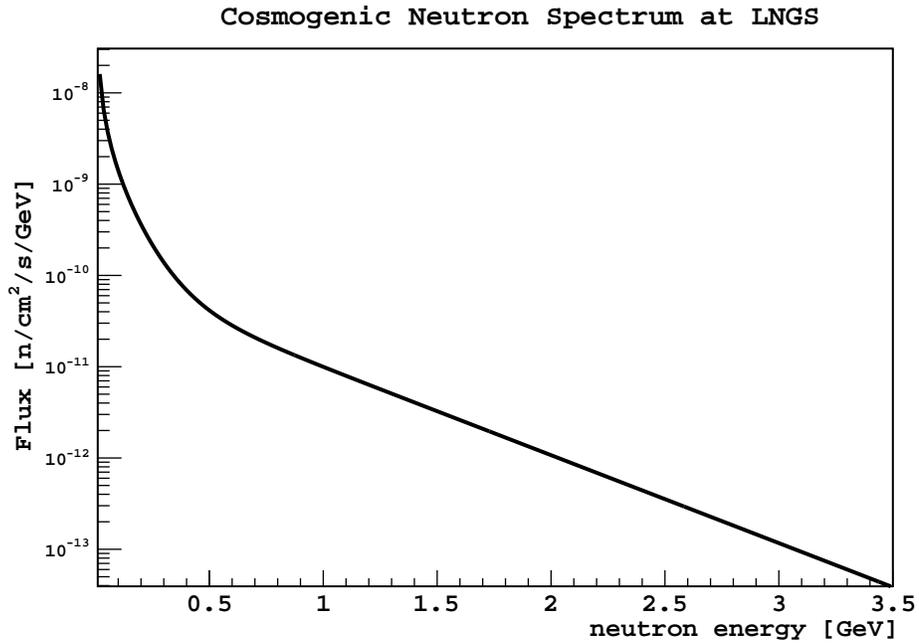


Figure 5.5: Spectrum of cosmogenic neutrons assumed for LNGS.

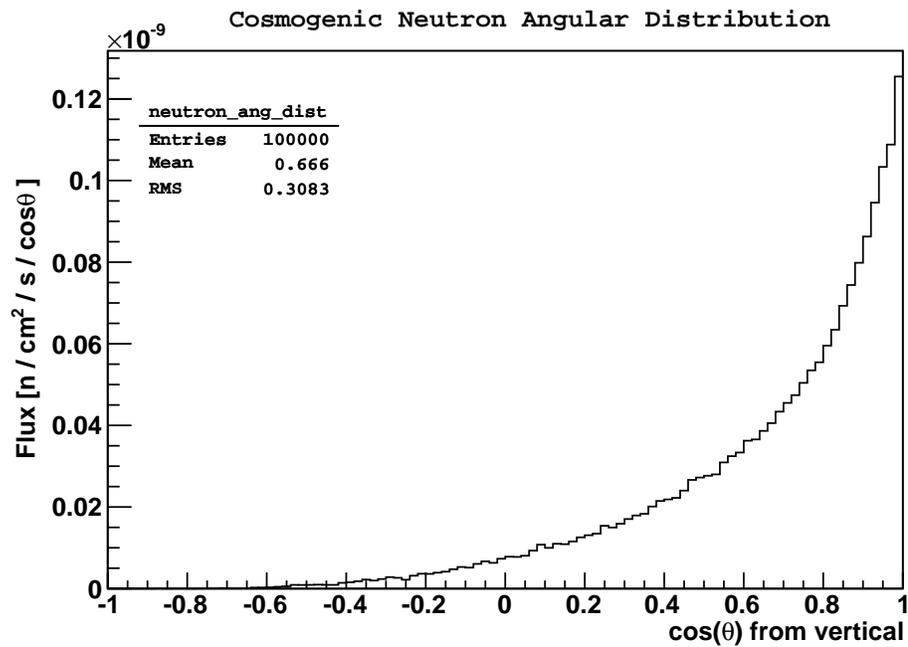


Figure 5.6: Zenith angular distribution of cosmogenic neutrons (relative to straight down) assumed for LNGS.

External Sources

By external sources, we refer to gammas and neutrons emanating from the rock walls of the experimental cavern. Like internal sources, the gammas come primarily from the uranium and thorium chains and potassium, the radiogenic neutrons come from fission and (α, n) processes induced by the same, and the cosmogenic neutrons are generated by spallation from high energy cosmic ray muons.

The external gamma ray flux in Hall C of LNGS has been measured to be $\sim 1 \text{ cm}^{-2} \text{ s}^{-1}$ [87]. The spectrum of the gamma rays has also been measured, but, in the interest of both simplicity and the effort to be conservative, we assume that all of the flux is at the highest energy gamma in the natural decay chains, the 2.6 MeV gamma of ^{208}Tl . This assumption should be conservative because the majority of the flux is at much lower energies that are more effectively shielded by the water and scintillator tanks.

The fast (up to $\sim 20 \text{ MeV}$) neutron flux has also been measured by several groups, with widely varying results [123]; much of the uncertainty apparently arises from different assumptions used in the necessary Monte Carlo calibrations of detector efficiencies. The most conservative (i.e., largest) measurement of the flux reported in Ref. [123] is $3 \times 10^{-6} \text{ cm}^{-2} \text{ s}^{-2}$. The spectrum for external radiogenic neutrons, like that from detector elements, depends on the composition and levels of the different radioactive contaminations in the cavern rock. The elemental composition of the rock is 47.91% O, 30.29% Ca, 11.88% C, 5.58% Mg, 1.27% Si, 1.03% K and 1.03% Al [123]. The ^{238}U and ^{232}Th rates are reported respectively as 8.1 and 0.27 Bq/kg in Ref. [123] and as 11 and 0.37 Bq/kg in Ref. [124]; we will assume the more conservative (latter) of the two references. The ^{235}U and ^{210}Pb rates are not directly reported, so we will assume natural abundance and secular equilibrium, respectively. The spectrum of radiogenic neutrons produced in rock with the above composition and radioactive abundances, as calculated by SOURCES4A, is shown in Figure 5.7.

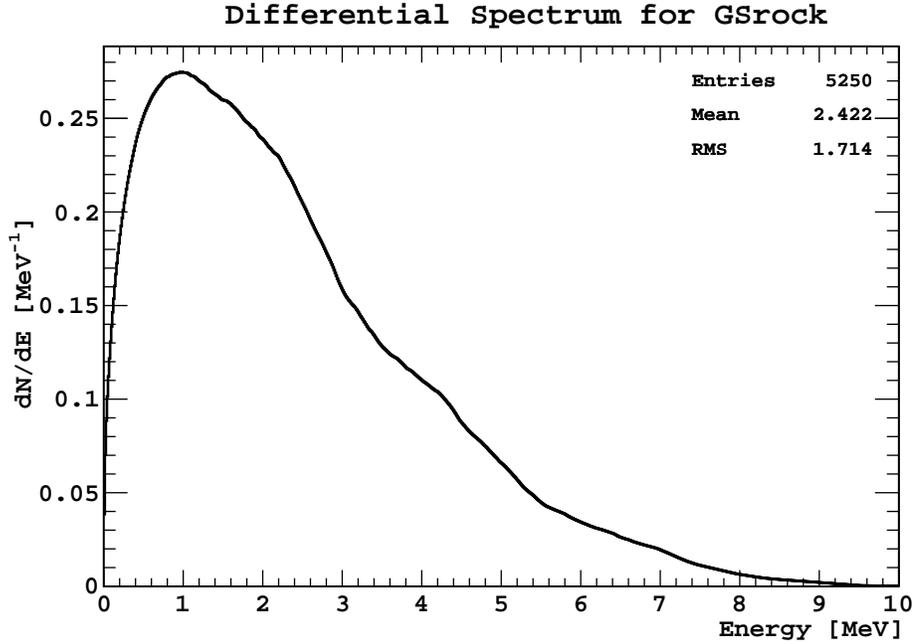


Figure 5.7: Differential spectrum of (α, n) and fission neutrons generated in the rock walls of the LNGS Hall C, calculated with SOURCES4A. The total flux is $3 \times 10^{-6} \text{ cm}^{-2} \text{ s}^{-1}$.

The external cosmogenic flux above 10 MeV from the rock walls at LNGS is $7 \times 10^{-10} \text{ cm}^{-2} \text{ s}^{-1}$, or $\sim 220 \text{ m}^{-2} \text{ year}^{-1}$. The spectrum and angular distribution are shown in Figures 5.5 and 5.6.

Backgrounds Summary

Table 5.9 lists the total gamma and radiogenic and cosmogenic neutron emission rates for each component of DarkSide-50, as well as rates due to external fluxes. Cosmogenic neutrons are only generated outside the inner detector cryostat. Where provided, the number in parentheses is the result for the high ^{210}Pb rate in Table 5.4; the default number assumes ^{210}Pb in secular equilibrium with ^{238}U .

Volume	Gammas primaries/s	Radiogenic Neutrons n/d	Cosmogenic Neutrons n/d
Active (^{39}Ar , not simulated)	<2	–	–
Fused silica windows (combined)	0.08	0.012	–
Inner PMTs (total)	2.8	0.36	–
Metal plates (combined)	0.03 (5.6)	6.8×10^{-5}	–
inner vessel support	0.02 (22)	0.006 (9.0)	–
Dewar	1.7 (75)	0.08	–
Neutron veto PMTs (total)	226	93	–
Neutron veto scintillator	0.4	0.29	18
Neutron veto tank	130	2.6	20
Muon veto water	30	0.76	420
CTF tank	860	17.2	130
External	4.60×10^6	1.20×10^6	300

Table 5.9: Total rate of backgrounds to be simulated for each of the detector components. For a description of the components, see Figure 5.2. Cosmogenic neutrons are only simulated outside the inner Dewar, as muons that pass through the active volume should be almost completely vetoed. In the gammas column, the term “primaries” refers to either the parent isotope of a decay chain (for volumes inside the Dewar, inclusive) or the total rate of gammas to be simulated (for the shield layers outside the Dewar). Where two numbers are given, the first is assuming ^{210}Pb rates in secular equilibrium with ^{238}U the second in parentheses uses the available reported values for ^{210}Pb from Table 5.4.

5.3 Simulation Results

5.3.1 Analysis Procedure

The first step in analysis is to determine the total “visible” or electron-equivalent energy deposited in the active volume. For each particle tracked in the simulation, we record several statistics, such as the particle type, process and parent particle by which it was produced, and the total amount of energy deposited in each volume. Since the light yield for nuclear recoils is quenched relative to electrons, and the amount of quenching is a function of the total energy of the nuclear recoil, we must add to the energy of the parent nuclear recoil the energy deposited by electrons ionized by it. So, we sort through the list of particles and, for electrons ionized by nuclear recoils (i.e., that were tagged by `GEANT4` as created by the “hIoni” process), we take their deposited energy and add it to the parent recoil. Then, we tally all the energy deposited in the active volume, scaling the energy from nuclear recoils by the appropriate quenching factor. We assume that the TPC has an applied drift field of ~ 1 kV/cm, which would make the relative quenching 0.5. (The absolute quenching for nuclear recoils is ~ 0.25 independent of field, and 0.5 for betas with a 1 kV/cm field). We call this tally the total quenched energy deposit (edep). As a check that the ionized electron energies are properly associated to the parent nuclear recoil, we also record the initial kinetic energy of all of the nuclear recoils and verify that the totals are equal.

We assume that the light yield in DarkSide-50 will be 6 p.e./keV_{ee} at zero applied field. This is the stated design goal for DarkSide-50, and can be justified based on the results of DarkSide-10 (Chapter 4). In DarkSide-10, we obtained a zero-field light yield of 4.5 p.e./keV_{ee}, and determined that replacing the bottom, relatively low quantum efficiency PMT with an array of PMTs identical to the top should increase that number by at least $\sim 15\%$. In DarkSide-50, the light yield should be further enhanced by the use of thinner ITO layers and replacing the 3M foil reflector with more efficient newly developed high-

crystallinity PTFE; rough optical simulations by DarkSide collaborators indicate that this should raise the light yield by a further 10–20%.

We apply a WIMP-induced nuclear recoil search window of 30–200 keV_{nr}, or 45–300 photoelectrons, assuming a quenching factor of 0.25 for nuclear recoils, or 15–100 keV_{ee} at a ~ 1 kV/cm drift field. The number of events with total quenched edep in this range is the base background rate that makes it through the shielding. At this stage, we do not simulate the scintillation or any other optical processes, but take all cuts directly on (quenched) edep. We also ignore effects due to finite detector energy resolution.

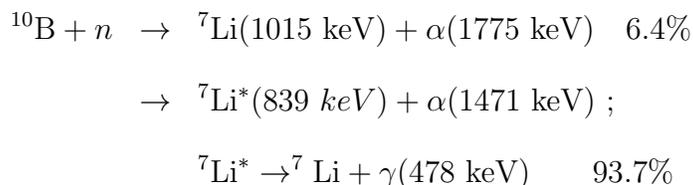
To further reduce the background, we apply a series of analysis cuts. Gamma and neutron backgrounds are considered separately, since gammas have a much higher rate but can be removed by pulse shape and S2/S1 discrimination. Neutron events will sometimes deposit some fraction of their energy through electromagnetic channels, such as inelastic scattering and neutron capture. If 50% or more of the energy in neutron events is deposited via electromagnetic channels, we assume it will be rejected by pulse shape or S2/S1 discrimination.

Surface alpha backgrounds are not simulated, but we assume that they can be effectively removed by a suitable fiducial volume cut. For DarkSide-50, we assume that any event within 1.5 cm of the top or bottom of the active volume or 3 cm of the outside wall will be rejected, leaving a fiducial mass of ~ 33 kg. This fiducial volume cut will also slightly affect the neutron and gamma rates, since event distributions will be higher toward the edge of the fiducial volume.

Neutrons can also be rejected by detecting multiple recoils in the active volume or detecting their capture in the active neutron veto. Multiple recoils, if sufficiently far apart, can be identified by the presence of two separate S2 pulses if the recoils are separated vertically, or the spread of the hit patterns if separated laterally. The minimum separation that can be identified must be determined experimentally, so we make very conservative assumptions for this cut: we assume that neutrons can be rejected when they produce multiple recoils separated by more than 1 cm vertically (which would correspond to ~ 5 μ s S2 separation) or

15 cm laterally (two PMT diameters). For the veto, we assume that an event can be vetoed if it deposits more than 40 keV_{ee} in the scintillator volume. The electron-equivalent energy in the veto is calculated similarly to the active volume, except that the quenching is energy-dependent. The quenching for protons, alphas, and heavier nuclear recoils in pseudocumene assumed for this calculation is shown in Figure 5.8.

Most neutron captures in the boron-loaded scintillator of the neutron veto will occur on ¹⁰B due to its high cross section [83]. The capture process can proceed in two channels:



The quenching of the nuclear recoil products (⁷Li and α) results in an electron-equivalent energy of ~50–60 keV_{ee}, which should be detectable with the 40 keV_{ee} threshold; the excited channel could be significantly higher energy depending on how much of the gamma energy is deposited in the veto.

As can be seen in Table 5.9, the rate of gammas and radiogenic neutrons entering the outermost volumes is extremely large. Directly simulating enough of these particles to obtain a reasonable estimate of the interaction rate in the active volume is not practical. Instead, we use a “bootstrap” technique. In the first stage, we spawn the external backgrounds just outside the outermost shield layer (the CTF tank), and count all the interactions in the neutron veto volume. (The veto volume is used for simplicity, since our analysis tools are already configured to store events that deposit energy in the veto). In the second phase, we repeat the process (using the same spectrum), but with the outer shield layers removed from the simulation and the particles spawned just outside the neutron veto volume.

The total expected background rate in the active volume is then the initial rate from Table 5.9 times the fraction of events in phase 1 that were detected in the veto, times the

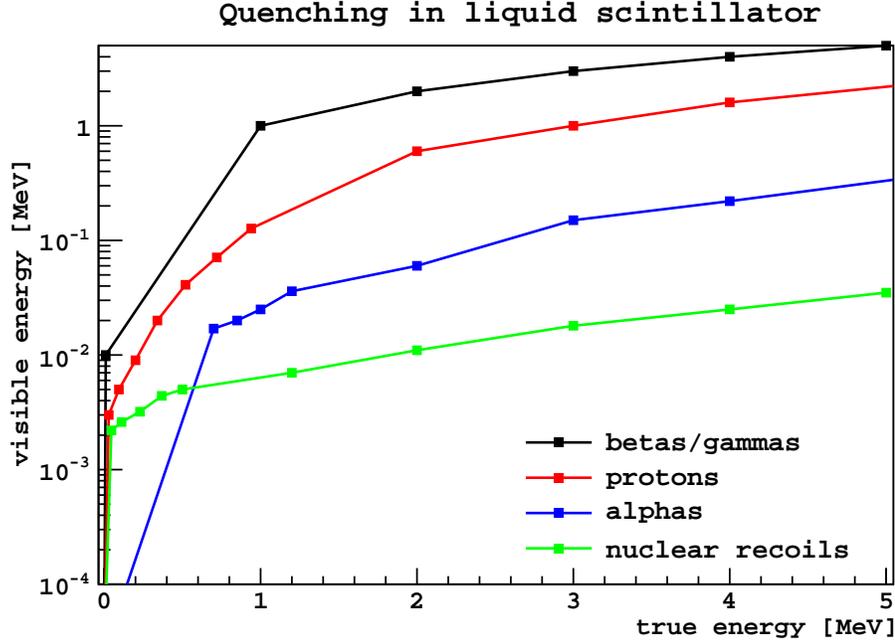


Figure 5.8: Assumed quenching of light yield of heavy particles in liquid scintillator relative to electrons [125, 126].

fraction of events in phase 2 that are detected in the active volume (and pass all cuts). (Since we use an isotropic angular distribution for the backgrounds, there is another factor of 4 to account for the fact that, in each phase, half of the simulated flux is pointing away from the detector.) This estimation should be very conservative, since we use the same energy spectrum as input for phase 1 and 2; realistically, most of the particles will have lost some energy passing through the outer shield layers, and therefore have less range than the higher-energy spectrum we simulate.

For each background type (gammas, radiogenic neutrons, and cosmogenic neutrons), we report the simulated background rate both in total counts over the simulated livetime and as a rate in counts per day per kilogram for gammas and counts per year per kilogram for neutrons, considering in both cases the fiducial mass. Since many of the counts were zero after all cuts, for all the rates we report the 95% upper confidence limit for the rate based on the detected counts assuming Poisson statistics. We do not attempt to evaluate any

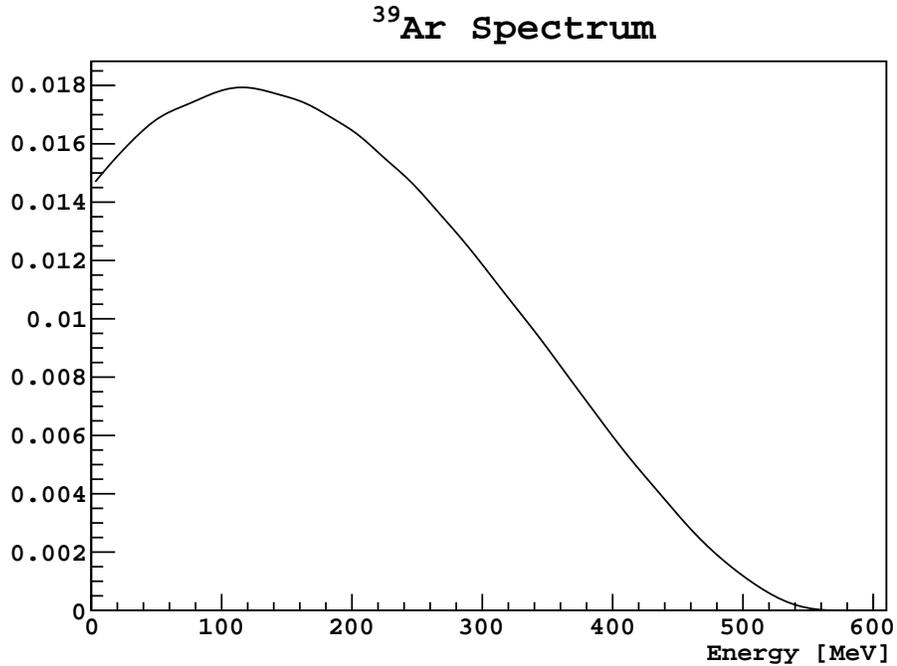


Figure 5.9: Beta decay spectrum of ^{39}Ar .

systematic uncertainties in the input spectra or simulation or analysis procedures, since we have attempted to be conservative at every step of the process.

5.3.2 Gammas

The background due to ^{39}Ar beta decays in the active volume is not reducible via shielding, so the rate of gamma backgrounds reaching the active volume is unimportant so long as it is less than the ^{39}Ar rate in depleted argon. The spectrum of ^{39}Ar , obtained from simulating radioactive decays via GEANT4, is shown in Figure 5.9. Approximately 23% of the ^{39}Ar spectrum falls within the 15–100 keV energy range, which leads to a background rate of $\sim 2 \times 10^5$ counts/kg/day in the nuclear energy window for natural argon, and < 800 counts/kg/day for depleted argon. So, as long as the rate of gamma events in the 15–100 keV energy window is less than a few hundred counts per kilogram per day, it will not have any significant effect on the total background rate.

(a) Detected background counts due to gammas in simulations

Volume	Simulated livetime (d)	Counts after cut		
		No Cuts	Energy	Veto
Inner PMTs	53	915415	234050	207305
Metal plates	53 (3.6)	3994 (2156)	853 (912)	782 (905)
IV Support	53 (3.6)	490204 (75092)	111929 (41141)	98228 (40248)
Fused silica windows	53	121230	27580	24933
Dewar	53	573733	154080	111956
Shielding and vetoes	25	1	1	0
Neutron veto PMTs	25	394	83	9
External	52	109	20	3
Total	25 (3.6)	9.93×10^5 (1.92×10^5)	2.49×10^5 (7.30×10^4)	2.09×10^5 (6.73×10^4)
				5.58×10^4 (9.86×10^3)
				55753
				233 (284)
				29205 (2994)
				6955
				26204
				0
				4
				1

(b) 95% upper C.L. estimate for gamma background rates after cuts ($\text{kg}^{-1} \text{day}^{-1}$)

Volume	No Cuts	Energy	Veto	Fiducial
Inner PMTs	340	88	78	33
Metal plates	1.5 (12)	0.33 (5.2)	0.31 (5.7)	0.15 (2.7)
IV Support	180 (410)	42 (225)	37 (220)	17 (26)
Fused silica windows	46	11	9.8	4.32
Dewar	240	72	57	20
Shielding and vetoes	0.004	0.004	0.002	0.002
Neutron veto PMTs	0.34	0.079	0.013	0.011
External	0.094	0.021	0.006	0.006
^{39}Ar , depleted by 0.04	8.6×10^4	800	800	800

Table 5.10: Simulation results for gamma backgrounds in DarkSide-50. By default we assume ^{210}Pb in equilibrium with ^{238}U ; the numbers in parentheses are the results for the high ^{210}Pb concentrations in Table 5.4. See Section 5.3.1 for a description of the cuts. The counts in the "Total" row are obtained by scaling the counts in each component to the smallest livetime simulated. The gamma background rate will be further reduced by a factor of $\sim 10^8 - 10^9$ by pulse shape and scintillation to ionization ratio cuts. Note also that the rate quoted for ^{39}Ar is only an upper limit, and could be much smaller.

The rate of gamma backgrounds in the active volume after each successive cut is presented in Table 5.10. For the external gammas, 6 gammas out of 5×10^7 simulated were detected in the neutron veto volume in the first step of the bootstrap simulation described in Section 5.3.1. With no cuts, $109/1 \times 10^7$ were detected in the active volume in the second step, leading to a total fraction of detected external gammas of $< 5 \times 10^{-12}$, which, given the total incident flux of 4.6×10^6 per second, yields $\sim 2 \times 10^{-5}$ interactions per second in the entire 50 kg volume, or < 0.04 counts/kg/day. This rate, already negligible compared to the ^{39}Ar rate, is further reduced by the energy cut.

5.3.3 Radiogenic Neutrons

The predicted rate of radiogenic neutron backgrounds in the active volume after each cut is shown in Table 5.11. For the external radiogenics, 128 events out of 3×10^6 simulated deposited energy in the neutron veto in the first step of the bootstrap process. In the second step, 62 events out of 5×10^7 simulated deposited energy in the active volume, yielding a maximum rate of $\sim 3 \times 10^{-3}$ background events/year/kg, which is further reduced by the application of cuts.

Table 5.11 highlights two key features of the active neutron veto. First, it reduces the background due to radiogenic neutrons from internal sources by a factor of ~ 100 or more; recall that these are neutrons that are emitted from inside the veto and scatter in the active volume *before* entering the veto. (In fact many are captured on inert elements, but the capture gamma is detected.) The veto power is even stronger for neutrons that must travel through the veto before reaching the active volume.

The second key feature is an *in situ* evaluation of the neutron background rate. For example, the neutron background in the WIMP search window can change by a factor of ~ 1000 if we assume the high reported rate of ^{210}Pb contamination in PTFE. Although all materials will be screened before construction, it is not possible to fully evaluate everything in advance. With the neutron veto, we can measure the rate of background neutron events

(a) Detected background counts due to radiogenic neutrons in simulations

Volume	Simulated livetime (y)	Detected counts after cut					
		No Cuts	Energy	β Fraction	Veto	Mult. Recoils	Fiducial
Inner PMTs	48	2390	694	354	2	0	0
Metal plates	48 (44)	0 (2)	0 (1)	0	0	0	0
IV support	48 (44)	62 (61458)	23 (21090)	8 (10143)	0 (140)	0 (42)	0 (20)
Fused silica windows	48	53	10	8	0	0	0
Dewar	48 (44)	267 (289)	111 (84)	52 (30)	0	0	0
Neutron veto PMTs	48	14	1	0	0	0	0
Neutron veto tank	48	1	1	0	0	0	0
Other shields/vetoers	48	0	0	0	0	0	0
External	670	62	3	0	0	0	0
Total	48 (44)	2786 (64000)	839 (21865)	422 (10563)	2 (142)	0 (42)	0 (20)

(b) 95% upper C.L. estimate for radiogenic neutron background rates after cuts ($\text{kg}^{-1} \text{ year}^{-1}$)

Volume	No Cuts	Energy	β Fraction	Veto	Mult. Recoils	Fiducial
Inner PMTs	1.0	0.32	0.18	0.003	0.001	0.001
Metal plates	0.001 (0.003)	0.001 (0.002)	0.001	0.001	0.001	0.001
Inner vessel support	0.031 (27)	0.013 (9.4)	0.006 (4.5)	0.001 (0.071)	0.001 (0.024)	0.001 (0.020)
Fused Silica windows	0.026	0.009	0.006	0.001	0.001	0.001
Dewar	0.12 (0.14)	0.05 (0.04)	0.027 (0.018)	0.001	0.001	0.001
Neutron veto PMTs	0.009	0.003	0.001	0.001	0.001	0.001
Neutron veto tank	0.002	0.002	0.001	0.001	0.001	0.001
Other shields/vetoers	0.001	0.001	0.001	0.001	0.001	0.001
External	2.6×10^{-3}	2.6×10^{-4}	1.0×10^{-4}	1.0×10^{-4}	1.0×10^{-4}	1.0×10^{-4}

Table 5.11: Simulation results for radiogenic neutron backgrounds in DarkSide-50. By default we assume ^{210}Pb in equilibrium with ^{238}U ; the numbers in parentheses are the results for the high ^{210}Pb concentrations in Table 5.4. See Section 5.3.1 for a description of the cuts. The counts in the "Total" row are obtained by scaling the counts in each component to the smallest livetime simulated.

during the WIMP search campaign without compromising the final “blind” analysis in the WIMP search window. If the measured neutron rates are significantly different than expected, we can reevaluate our background estimates before “opening the box,” so long as we believe we can accurately predict the fraction of unvetoed neutron events.

5.3.4 Cosmogenic Neutrons

In addition to the cuts applied to radiogenic neutrons, cosmogenic neutron backgrounds can also be vetoed by detecting the parent muon if it passes through a sensitive region of the detector or vetoes. Because we simulate the neutrons directly and not the muon or any possible accompanying shower, we assume the muon detection efficiencies obtained by Borexino [85]. Specifically, we assume cosmogenic neutrons that are generated in the CTF tank or water of the muon veto can be tagged and rejected with 99% efficiency, and those generated in the neutron veto and its tank with 99.992% efficiency. We make the very conservative assumption that we do not detect the associated muon or any other shower component for the external cosmogenic neutron flux. The cosmogenic neutron background rates are shown in Table 5.12.

5.4 Predicted Sensitivity

The DarkSide collaboration has proposed an initial three year campaign for DarkSide-50 using the Hamamatsu R11065 PMTs for the inner detector sensors. Depending on the results of this campaign, the inner detector could then be upgraded with QUPIDs. To predict the sensitivity of this campaign to WIMP interactions, we must estimate the total background rate over the three year period. Combining Monte Carlo results in the previous section is non-trivial, since the simulations cover different livetimes. We must also treat the ^{39}Ar rate differently than our simulated gamma and neutron rates, since for ^{39}Ar we have only an upper limit. If we ignore ^{39}Ar for now, then we can combine the neutron simulations

(a) Detected background counts due to cosmogenic neutrons in simulations.

Volume	Simulated livetime (y)	Detected counts after cut					
		No cuts	Energy	β Fraction	Veto	Mult.	Recoils
Neutron veto	15	1803	280	142	16	11	9
Veto tank	15	416	57	30	2	1	0
Muon veto	15	834	122	62	1	1	1
CTF tank	15	25	8	3	0	0	0
External	88	114	18	3	0	0	0
Total	15	3098	470	238	19	13	10

(b) 95% upper C.L. estimate for cosmogenic neutron background rates after cuts ($\text{kg}^{-1} \text{ year}^{-1}$)

Volume	No cuts	Energy	β Fraction	Veto	Mult.	Recoils	Fiducial	Muon Veto
External	0.030	0.0056	0.0017	0.0007	0.0007	0.0007	0.0007	0.0007
CTF tank	0.045	0.019	0.010	0.004	0.004	0.004	0.004	4×10^{-5}
Muon veto	1.1	0.12	0.098	0.006	0.006	0.006	0.006	6×10^{-5}
Veto tank	0.58	0.091	0.052	0.008	0.006	0.006	0.004	5×10^{-7}
Neutron veto	2.4	0.40	0.21	0.031	0.023	0.020	0.020	2×10^{-6}

Table 5.12: Simulation results for background from cosmogenic neutrons in DarkSide-50. See Section 5.3.1 for a description of the cuts. The counts in the "Total" row are obtained by scaling the counts in each component to the smallest livetime simulated. We assume that, for neutrons generated inside the CTF tank and muon veto layers, we can identify the parent muon and veto the event with 99% efficiency. For neutrons generated inside the neutron veto and its containing tank, we assume a 99.992% veto efficiency [85].

to estimate the total background rate due to neutrons in the following way. For a given source i (cosmogenic or radiogenic neutrons from each volume simulated), we can construct the likelihood function for the true rate r_i from the detected number of counts x_i and the simulated livetime t_i as

$$L_i(r_i) = P(x_i|r_i, t_i) = \frac{(r_i t_i)^{x_i} e^{-r_i t_i}}{x_i!} \quad (5.9)$$

where the rightmost expression is the Poisson probability to observe x_i counts with an average expectation of $r_i t_i$. Since this is the likelihood for the *average* rate from this source, for sources that are vetoed with some efficiency (such as cosmogenic neutrons spawned in the shield layers), we can obtain the average rate of events passing the veto cut by scaling the simulated time by $1/(1 - \eta)$, where η is the veto efficiency. The total likelihood that the set of rates from all sources \mathbf{r} is equal to a particular value is the product of the individual likelihoods:

$$L(\mathbf{r}) = \prod_i L_i(r_i). \quad (5.10)$$

Finally, the likelihood that the true total rate is some value r is equal to the maximum likelihood of all the sets \mathbf{r} such that $r = \sum r_i$. Confidence intervals can then be defined as the values of r such that the integral over the normalized total likelihood function is equal to the desired confidence level [127].

Figure 5.10 shows the normalized likelihood for the total rate of background due to neutrons in DarkSide-50 based on the results in Tables 5.10(b) and 5.11(b), assuming the lower rates of ^{210}Pb in PTFE. In this case, the most likely rate is determined by unvetoesd cosmogenics spawned in the muon veto volume, 0.01×9 counts/15 years = 0.006 counts/year. The 90% and 95% upper confidence limit for the average rate are 0.052 and 0.067 counts/year, respectively. If the ^{210}Pb rate in the PTFE is the high value reported in Table 5.4 is correct, then the estimated average total rate is 0.46 counts/year, with 90% and 95% upper C.L. of 0.62 and 0.67 counts/year.

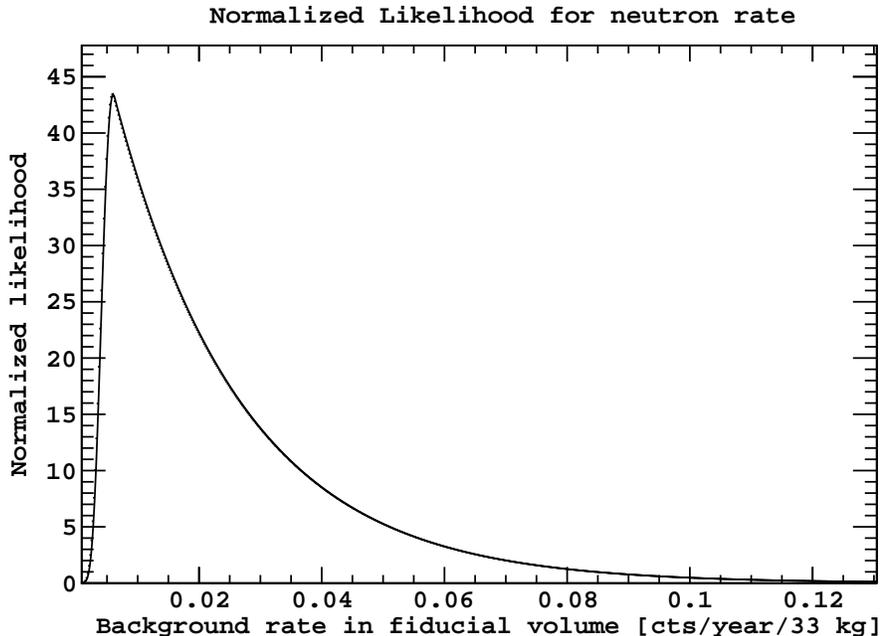


Figure 5.10: Likelihood (normalized to unit integral) for the total rate of background from neutrons in the DarkSide-50 fiducial volume after all cuts.

Since the total rate due to gammas in the detector is $\sim 10\%$ of the limit on the ^{39}Ar rate, we can neglect gammas and consider ^{39}Ar as our only beta-like background. We expect at most ~ 800 counts per day per kg of fiducial mass in the WIMP energy range, or 9.64×10^6 counts/year in the ~ 33 kg fiducial volume, before pulse shape discrimination (PSD) and scintillation to ionization (S2/S1) cuts. If we assume a (conservative) fixed rejection power of 10^8 for the PSD and S2/S1 cuts (with a 50% nuclear recoil acceptance), then the total expected rate of background due to ^{39}Ar is at most ~ 0.1 counts/year. Although this rate is an upper limit, we don't have the necessary information to construct the full likelihood for the expected ^{39}Ar rate, so the best we can do to combine it with the neutron rate already obtained is to add it. Thus, we can state that the 90% upper confidence limit on the expected total rate in DarkSide-50 is at most 0.15 counts/year (again, assuming that the ^{210}Pb rate in PTFE is in equilibrium with ^{238}U).

The sensitivity to WIMP interactions is maximized if there is no background, so the natural conclusion to this background rate estimation is to ask how likely is the three-

year search campaign to be background-free. Unfortunately, we can't answer this question exactly, but we can get the approximate probability if we are willing to abuse the concept of likelihood somewhat. That is, if we are willing to treat the likelihood for a given total rate based on the Monte Carlo results as the probability for that rate, then the probability to receive 1 or more background events is equal to

$$P_{\geq 1} = 1 - P_0 \approx 1 - \int e^{-rt} L(r) dr, \quad (5.11)$$

where, in this case, $t = 3$ years. If we again assume that the rate from gamma-like backgrounds is constant and contains only the ^{39}Ar then the likelihood is given by Figure 5.10 shifted to the right by the ^{39}Ar rate of 0.1 counts/year. I.e.,

$$L(r) = L_n(r - 0.1 \text{ cts/year}), \quad (5.12)$$

where $L_n(r_n)$ is the neutron rate likelihood in Figure 5.10. Following this prescription, the “probability” to receive 1 or more background counts in a 3-year WIMP search with DarkSide-50 is less than 0.31. Again, I specify “less than” because the ^{39}Ar depletion factor of 0.04 for underground argon is an upper limit, and the assumed PSD and S2/S1 rejection power of 1×10^8 is conservative.

All of the preceding discussion (and, in fact, this entire chapter!) can be summarized by stating that we have a better than even chance that the proposed three-year DarkSide-50 WIMP search will have zero background, assuming adequate controls are taken during material selection and construction. If this is the case, the campaign will represent the largest background-free WIMP search yet conducted. If we also see a null signal (i.e., no events passing all WIMP selection cuts throughout the campaign), then the 90% upper C.L. on the expected number of counts is 2.3 (obtained from the likelihood for a Poisson process with 0 detected counts). Taking into account the 50% nuclear recoil acceptance imposed by the PSD and S2/S1 cuts, then the limits DarkSide-50 will be able to place on the WIMP-

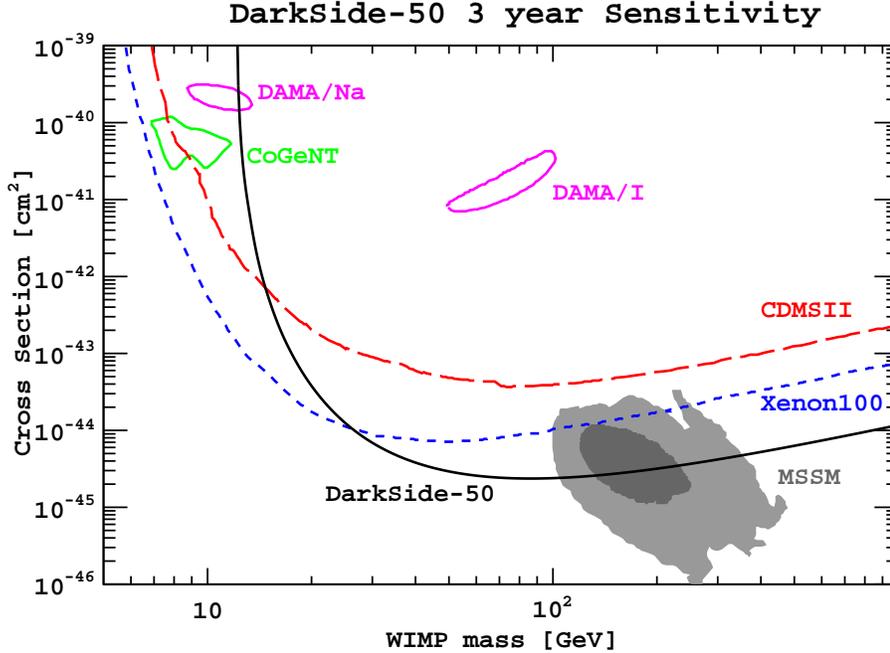


Figure 5.11: Predicted WIMP-nucleon cross section sensitivity for 100 kg-years of livetime in DarkSide-50, assuming 50% nuclear recoil acceptance, and ignoring detector resolution and threshold effects. The 1- and 2- σ preferred regions for the lightest neutralino in the MSSM [128] are shown by the dark and light gray shaded regions, respectively. The dashed exclusion limits for CDMS (red) and Xenon (blue), as well as the acceptance regions for CoGeNT (green) and DAMA (magenta) are taken from Ref. [62], Figure 5.

nucleon cross section are set by the cross section that results in 4.6 expected counts per ~ 100 kg-years when integrating Equation (1.30) over the specified acceptance window of 30–200 keV.

Figure 5.11 plots this cross section sensitivity as a function of WIMP mass. The current limits from Xenon100 and CDMS, as well as the accepted regions from DAMA/Libra and CoGeNT and the predicted region for the neutralino from the MSSM are included for comparison. In the three-year campaign, DarkSide-50 should reach a sensitivity to WIMP-induced nuclear recoils of $\sim 3 \times 10^{-45}$ cm² for a 100 GeV WIMP. Unfortunately, the relatively high threshold needed to reject ³⁹Ar backgrounds and the resulting insensitivity to low-mass WIMPs prevents DarkSide from having any significant input into the ongoing DAMA/CoGeNT modulation signal controversy.

5.5 Conclusions

Our implementation of the `GEANT4` Monte Carlo framework, `s4det`, has been designed to enable users to quickly and easily simulate backgrounds for simple detector geometries consisting primarily of nested cylinders. Using `s4det`, we have shown that the total expected background rate for a 3 year WIMP search campaign in DarkSide-50 is less than 0.5 counts once all cuts have been applied, assuming careful screening of the PTFE for the inner vessel support structure. This result assumes the use of photomultipliers as the inner detector sensors, and will be even lower if QUPIDs are used in future campaigns. Although not calculated rigorously, we have presented arguments that indicate that we can expect this campaign to be completely background free with at least $\sim 70\%$ confidence.

If this estimate is correct, DarkSide-50 will have the largest background-free exposure of any WIMP search experiment to date. This is due in large part to the use of an active neutron veto that surrounds the active volume to remove neutron-induced nuclear recoils, which are the main source of background. The neutron veto also allows for *in situ* monitoring of the neutron background, which is extremely useful information for determining the expected remaining background in the WIMP search window.

As Figure 5.11 shows, DarkSide-50 should have sensitivity to WIMP interactions comparable to the leading detectors of the current generation. In a 3 year run, DarkSide-50 should explore a significant fraction of the remaining parameter space for the neutralino, probably the most favored WIMP candidate. Furthermore, if the initial survey of materials and calibrations confirm the very low estimated background rate, we will be able to claim statistically significant evidence of dark matter interactions in DarkSide-50 with the observation of only 2–3 events.

Appendix A

Measurement of the ${}^7\text{Be}$ solar neutrino flux with Borexino

This chapter describes in brief the analysis of Borexino data that led to the first real-time measurement of the interaction rate of ${}^7\text{Be}$ solar neutrinos [4, 108]. This measurement is an important confirmation both of our understanding of the nuclear processes that power the sun and the mechanisms and parameters of neutrino oscillations in vacuum and matter. The Borexino detector is described in detail in Ref. [129]. The active target is 278 tons of liquid scintillator (pseudocumene mixed with a small amount of PPO waveshifter). The active scintillator is contained within a thin 8.5 m diameter spherical nylon vessel that separates it from two nested buffer volumes of non-scintillating pseudocumene doped with DMP, a quenching agent. The buffer shields the active volume from backgrounds emitted by the 2212 PMTs that view the active volume. The entire inner assembly is contained within a 13.7 m diameter stainless steel sphere, which is itself contained inside an 18 m diameter cylindrical water tank that serves as further shielding from external gamma backgrounds. The water tank is also instrumented with PMTs to detect the Cerenkov light of cosmic muons passing through the water and serve as an active muon veto.

The ${}^7\text{Be}$ rate analysis was carried out by two independent groups working in parallel to serve as mutual cross-checks for each other. I will describe the “mach4” analysis led primarily by Princeton, Virginia Tech, and the U./ Mass./ Amherst students. I was in particular responsible for the low-level hit decoding and calibration, which will be covered in detail in Section A.1. With a few exceptions, this chapter will describe the state of the mach4 analysis at the time of the second ${}^7\text{Be}$ paper (Ref. [4]). Since that time, considerable further development has taken place, including a series of calibrations with sources inserted into the active volume. I will reference some of these further developments, in particular those that serve to confirm assumptions made at earlier stages in the analysis.

A.1 Low level calibration

This section describes the steps used to translate the raw digitized output from the inner detector electronics (the Laben boards) into hit timing and charge (i.e., number of photoelectrons) information for each channel. When the signal for a given PMT rises above threshold (which is set to a value considerably smaller than 1 photoelectron for that channel), the PMT current is integrated for 80 ns, then the channel ID, charge, and timing information are stored in a local buffer. After the 80 ns integration, the channel is dead for a further 70 ns. The current number of hits in each buffer is communicated to the central trigger, the Borexino Trigger Board (BTB). The BTB generates an event trigger if the total number of hits within a given time window is above a set threshold; generally 25 hits in 60 ns. This trigger records the absolute time from a GPS clock and writes to disk all hits stored in buffers and any further hits that arrive over a set trigger. The trigger window for normal inner detector triggers was $7\ \mu\text{s}$ prior to Run 6592 (18 Dec. 2007) and $16.5\ \mu\text{s}$ after that run. After each trigger window there is a fixed dead time of $10\ \mu\text{s}$ while the electronics reset. The BTB also generates special calibration triggers (described later) at a fixed rate and triggers initiated by the outer detector PMTs. After every outer detector trigger (sometimes referred to as

muon or MTB trigger), the BTB immediately generates a neutron trigger with a 1.6 ms window. This trigger was added in Dec. 2007 in order to be able to tag neutrons and some cosmogenic isotopes.

Each hit stored in the raw data file has the following structure which will be explained in more detail later:

- channel ID (16 bit unsigned int)
- ramp1 (8 bit unsigned int)
- ramp2 (8 bit unsigned int)
- gray counter (16 bit unsigned int)
- ADC samp1 (8 bit unsigned int)
- ADC samp2 (8 bit unsigned int)
- hit flag (16 bit unsigned int)
- error flag (8 bit unsigned int)

The timing information is encoded in the “ramp” and “gray counter” variables, and the ADC samples are used to read out the charge collected from the PMT over an 80 ns integration window. The gray counter value is read from a 20 MHz scaler. The ramp values are read from an ADC measuring the amplitude of a 10 MHz triangle wave generated from the same clock signal that drives the gray counter. The ramp1 variable is read at the beginning of the integration time (at the same time the gray counter is read out), and ramp2 is read 80 ns later after integration. The ramp values are used to interpolate between the 50 ns precision of the gray counter. By design, the slope of the triangle wave (ramp) should be rising when the gray counter register is even. Ideally, then, the time for a hit within the $2^{16} \times 50 \text{ ns} \approx 3.28 \text{ ms}$ reset time of the gray counter can be determined by

$$t_{\text{raw}}^0 = 50 \text{ ns} \times \begin{cases} \text{gray_counter} + \text{ramp1}/256 & \text{if gray_counter even} \\ \text{gray_counter} + 1 - \text{ramp1}/256 & \text{if gray_counter odd} \end{cases} \quad (\text{A.1})$$

Unfortunately, the readout is not this simple due to a number of complications:

- The "parity" of the 10 MHz triangle wave and gray counter is sometimes reversed; i.e., ideally the triangle wave is rising when the value of the gray counter register is even, but in some cases it is rising-on-odd.
- The amplitude of the triangle wave signal may not precisely cover the full range of the 8-bit ADC that reads the ramp variables.
- Although there is negligible phase offset between the gray counter and triangles waves, there is a few ns jitter in the precession of the gray counter scalar, which can lead to offsets of 50 ns near the transition point.
- The nominal 80 ns time between the reading of the ramp1 and ramp2 values can vary by channel.
- There can be an offset between the gray counters of different channels.

These issues are made worse because many are set to an essentially random state every time the boards are initialized or power cycled, and so must be determined every run.

To address these issues, before each run 1000 triggers are generated using a pulser connected directly to the front end electronics; the BTB continues to generate pulser triggers throughout the run at a rate of 0.5 Hz. The following section describes the calibrations performed using the pulser events.

A.1.1 Laben Pulser Calibration

Because all channels are in an almost unknown state prior to pulser calibration, we must process all of the pulser events three times. Since there is not a large number of pulser triggers (usually on the order of ~ 2000), we save the pulser triggers in memory during the first pass to avoid having to read through the file on disk again.

First pass

In the first pass through the pulser triggers, we determine the amplitude (in ADC counts) of the triangle wave for each channel (`ramp_width`) and the offset of each channel's gray counter value from the mean. Each channel *should* have exactly one hit in each pulser window, so we also tally the number of times each channel either has no hits or has more than one hit in a single trigger in order to later flag malfunctioning channels.

First we read through each pulser trigger in the entire run, and fill a histogram for each channel with all the values of `ramp1` and `ramp2` (i.e., a single histogram per channel containing both `ramp1` and `ramp2`). Then we calculate the mode of all of the gray counter values in the event, and keep a running tally for each channel of the number of times its gray counter value was equal to the mode.

It can happen that several channels or boards can fail simultaneously during a run without the run itself ending. We attempt to identify these events to exclude events after such an occurrence from the final data set. Because it is very rare that any channels fail in the first 1000 pulser events taken before the run begins (since they take only a few seconds), we use this as our benchmark. After the first 1000 pulser events have been read, we calculate the average number hits during pulser events for that period. Then, if, during any of the pulser events that occur during the run, the number of hits is below 80% of the initial average, that event is flagged as the last good event, and we throw out the rest of the run. Although it may be possible for some of the channels that dropped out to recover at a later point, we try here and elsewhere during channel selection to include only channels and triggers that are known to be good. In other words, all data is suspect until “proven” valid, rather than the more inclusive approach.

After reading through all of the pulser events in the run, we take the following steps:

1. Determine the reference channel, a reliable channel that can be used to determine the relative offset of the gray counters for each other channel. This is determined by the channel whose gray counter value fell most often on the mode during pulser events.

2. Flag for removal all the channels with a high rate of misfires (no hits) or multifires (multiple hits) on pulser events. Channels that misfired or multifired on more than 10% of the pulser events, or that did not fire on at least 1000 triggers are removed.
3. Determine the width of the triangle wave in ADC counts for each channel. A typical histogram compiled in the first pass is shown in Figure A.1. We define the minimum and maximum extent of the ramp as the first and last local maximum in the histogram, respectively (shown by the red boxes in Figure A.1).
4. Based on the ramp determination, the channel can be flagged for removal under several conditions:
 - (a) No local maximum is found within 20 counts of either edge.
 - (b) Either of the bins identified as the maximum or minimum extent contains fewer than 5 counts.
 - (c) The wave extends past the edge of the ADC; i.e. the minimum is at 0 or the maximum is at 255.
5. The final ramp edges are defined to be two bins further from the bins identified in the previous step. This more accurately accounts for the fact that the triangle wave spends less time at the peaks of its pattern.

Second Pass

In the second pass through the pulser triggers, we determine the relative offset between the channels' gray counters, the parity of the slope of the 10 MHz triangle wave relative to the last bit of the gray counter (rising-on-even or rising-on-odd) and the actual length of the nominal 80 ns delay between sampling of ramp1 and ramp2 for each channel. First, we histogram the offset of the gray counter value for each channel relative to that of the reference channel determined at the end of the previous pass.

To determine the parity of the triangle wave, we need a method to guess whether the slope is rising or falling at the time the ramp1 value is read. To do this, we calculate the

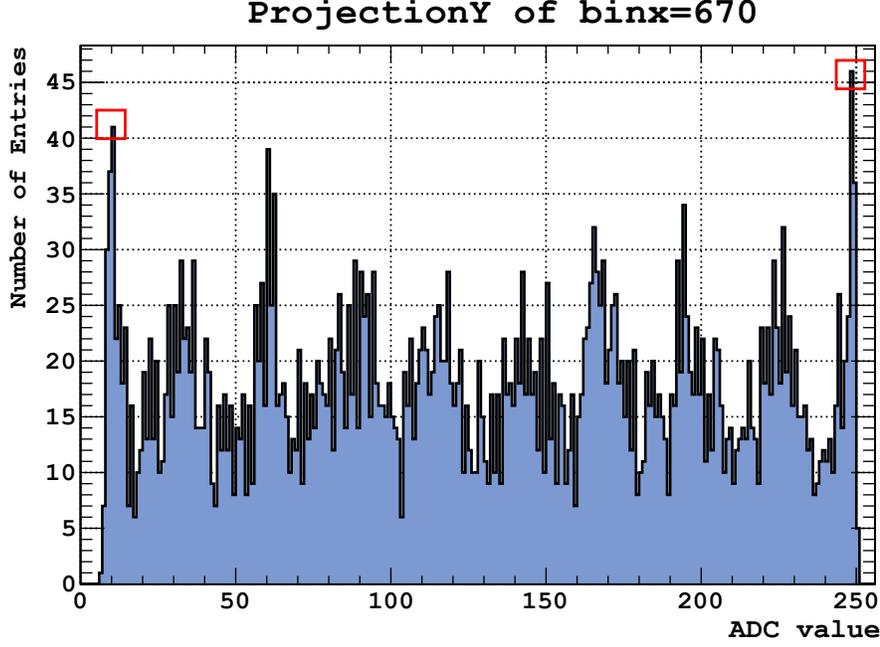


Figure A.1: Typical histogram of the ADC values for the 10 MHz triangle wave generated during precalibration for each channel. The red boxes show the bins identified as the maximum and minimum extent of the ramp.

nominal value for ramp2 under the assumption that the slope is rising or falling (assuming the nominal 80 ns delay between the two samples), then take the difference between the calculated and actual values for ramp2. Whichever assumption yields a smaller difference is assumed to be correct. To calculate the assumed value for ramp2 from ramp1, first we define the slope of the ramp in counts/ns as

$$\text{slope} = \frac{\text{ramp_max} - \text{ramp_min}}{50 \text{ ns}}. \quad (\text{A.2})$$

The time (relative to the gray counter) for ramp1 assuming a rising slope is

$$t_{1,r} = (\text{ramp1} - \text{ramp_min})/\text{slope}. \quad (\text{A.3})$$

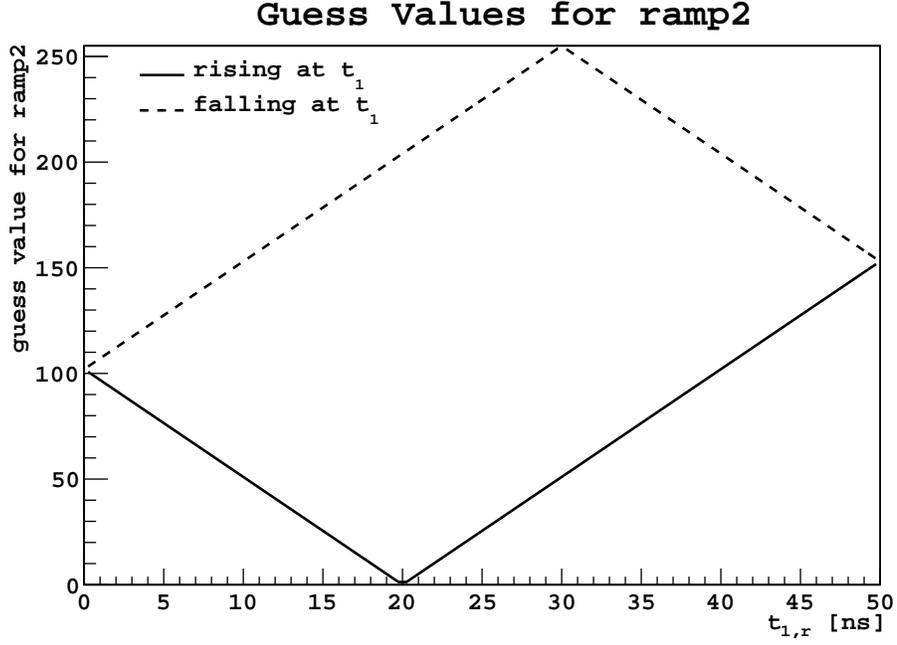


Figure A.2: Predicted value for ramp2 determined from ramp1, under the assumption that the slope is rising (solid line, Equation (A.5)) or falling (dashed line, Equation (A.6)) at t_1 .

If the slope is actually falling, then the time would be

$$t_{1,f} = 50 \text{ ns} - t_{1,r}. \quad (\text{A.4})$$

For an 80 ns delay time, if t_1 is less than 20 ns, then the slope at t_2 is opposite from t_1 , otherwise it is the same. The predicted ramp2 value under the hypothesis that the slope is rising at t_1 is then

$$\text{ramp2}_r = \text{ramp_min} + |\text{ramp_max} - \text{slope} \cdot (t_{1,r} + 30)|. \quad (\text{A.5})$$

Likewise, the prediction under the hypothesis that the slope is falling at t_1 is

$$\text{ramp2}_f = \text{ramp_max} - |\text{ramp_min} - \text{slope} \cdot (t_{1,r} - 30)|. \quad (\text{A.6})$$

A plot of Equations (A.5),(A.6) is shown in Figure A.2.

Then we define the offsets as

$$\delta_r = |\text{ramp2}_r - \text{ramp2}|; \quad \delta_f = |\text{ramp2}_f - \text{ramp2}|. \quad (\text{A.7})$$

From the δ values, we also define an error estimator

$$err = \frac{\min(\delta_r, \delta_f)}{\delta_r + \delta_f}. \quad (\text{A.8})$$

For each channel, we tally the number of pulser triggers where the guessed slope indicates that the parity is rising-on-odd, rising-on-even, or indeterminate, defined by $err > 0.45$.

Next, for hits with $err < 0.45$, we attempt to determine $\Delta t = t_2 - t_1$, the time between the sampling of the ramp1 and ramp2 values. (Because of the nominal 80 ns value, we refer to this variable as d80.) If the slope is rising at t_1 , then

$$t_2 = \begin{cases} 50 \text{ ns} + (\text{ramp_max} - \text{ramp2})/\text{slope} & \text{if } t_1 < 20 \text{ ns} \\ 100 \text{ ns} + (\text{ramp2} - \text{ramp_min})/\text{slope} & \text{if } t_1 \geq 30 \text{ ns} \end{cases} \quad (\text{A.9})$$

If the slope is falling at t_1 , then

$$t_2 = \begin{cases} 50 \text{ ns} + (\text{ramp2} - \text{ramp_min})/\text{slope} & \text{if } t_1 < 20 \text{ ns} \\ 100 \text{ ns} + (\text{ramp_max} - \text{ramp2})/\text{slope} & \text{if } t_1 \geq 30 \text{ ns} \end{cases} \quad (\text{A.10})$$

The appropriate value for Δt is histogrammed for each channel.

After we have finished filling the tallies and histograms in the second pass, we perform the following steps:

1. Fit the histogram of Δt values for each channel to a gaussian; the mean obtained from the fit is the channel's d80 value. If the value is more than 15 ns away from the nominal value of 80 ns, the channel is flagged.

2. Determine the usual parity of the ramp slope, rising-on-even or rising-on-odd, based on which bin had more entries for that channel. If the bin with fewer counts had greater than 90% of the number of counts in the higher bin, or if twice the number of triggers with $err > 0.45$ is greater than the sum of the counts in the two parity bins (i.e., if more than 1/3 of triggers had indeterminate slope), the channel is flagged.
3. Determine the offset of each channel’s gray counter from the highest bin in the offset histogram. If the offset is more than 2 (positive or negative), the channel is flagged.

Third Pass

In the third pass, we try to identify any remaining systematic offsets between the calculated time for hits in a given channel relative to the reference channel. For each hit in each pulser event, we find the “floating” time, i.e., within the 3.28 ms gray counter reset window, using the following method (the same method used for the remainder of event processing).

1. Guess the slope of the ramp as in the previous section, and calculate t_1 and t_2 from that. If $err > 0.45$, throw out the hit.
2. If the parity from the guessed slope is not equal to that determined in the previous step, the gray counter may have advanced early or late. The last two bits of the “hit flag” from the raw data contain information that can be used to determine whether to correct this effect by adding a ± 1 count correction to the gray counter value. Table A.1 shows the code used to determine the offset.
3. Add to the channel’s gray counter value the gray offset calculated at the end of the second pass and the shift (if any) determined in the previous step. This determines the hit time to within 50 ns.
4. Add t_1 to the time determined in the previous step. Because the time obtained is less precise near the top and bottom edges of the ramp (because of rounding of the edges, etc.), if $ramp1$ is within 10 counts of `ramp_min` or `ramp_max`, instead add $t_2 - d80$, which can be determined more precisely.

```

int gray_shift=0;
unsigned time_bits = hitflag & 3;
if( slope != hw_slope ){
  if( rising_on_even &&
      ( !slope && time_bits == 3) || (slope && time_bits == 0) ) )
    gray_shift = -1;
  else if( !rising_on_even &&
           ( !slope && time_bits == 0) || (slope && time_bits == 3) ) )
    gray_shift = -1;
  else
    gray_shift = 1;
}

```

Table A.1: Code used to determine offset to be applied to gray counter if calculated slope does not match the expected slope based on the parity of the channel and the last bit of the gray counter value. Here `slope` is a boolean with value 1 for rising slope and 0 for falling.

5. Finally, add to this value a fixed offset for each crate, which is a function primarily of signal propagation times between all the crates, and does not change between runs. The crate delays are given in Table A.2.
6. In later evaluation, we will redefine the hit times relative to the GPS clock. There are several “logical” channels not connected to any PMTs that receive an electronic pulse when the GPS clock is read out (near the end of the trigger window), so the final time for all other channels is the floating time of these logic channels minus the channel’s floating time. If the hit time calculated this way is outside the trigger window (defined in the beginning of Section A.1), the hit is thrown out.

For each channel we fill a histogram of the offset of the calculated floating time of its hits from the floating time of the reference channel determined in the first pass. If the absolute value of the mean of this histogram or the RMS of the histogram is greater than 10 ns, the channel is flagged.

Crate	1	2	3	4	5	6	7
Offset (ns)	0	0	0	0	-2.48	-1.85	-4.53
Crate	8	9	10	11	12	13	14
Offset (ns)	-1.41	-0.20	-0.18	-12.65	-2.22	-0.09	-0.86

Table A.2: Fixed offset times applied to each crate in calculating hit times.

A.1.2 Laser and random calibration

The pulser triggers are used to calibrate all of the timing electronics downstream of the PMTs. The BTB also generates two other calibration triggers: laser triggers, in which the PMTs are flashed with a laser pulse through a fiber-optic array, and random triggers, which are simply empty acquisition gates. The laser triggers are used to determine any timing offsets due to the PMTs themselves (slow rise-time, etc) and to calibrate the single photoelectron response for each tube. The random triggers are used to determine the dark rate for each channel. The BTB generates laser and random triggers through all runs at a rate of 0.5 Hz. Because this does not result in enough statistics to accurately determine results for most channels, at least once a week the DAQ shifters acquire a dedicated calibration run with only laser and random triggers at a higher rate (and the usual rate of pulser triggers for electronics calibration). During data analysis, we use the dedicated calibration runs to determine the laser and random parameters, and use the in-run triggers as a consistency check.

Dark rate determination is simple. We tally each valid hit for a given channel during random triggers and divide by the length of the trigger window times the total number of random triggers. If the resulting dark rate is greater than 5 kHz, the channel is rejected.

Similar to the logic channels that are pulsed when the GPS clock is read, there is a separate set of channels that are pulsed when the laser is fired, and we use these channels to determine a hit’s “laser time” in those triggers. As we read through the laser triggers, we consider as a valid hits within a 120 ns window around the nominal laser arrival time of 650 ns. The large time window is used to handle channels that are not directly flashed

by the laser (due to damaged fibers, etc.) but that might see reflections. For hits in that window, we fill two histograms, one with the laser time, and one with the integrated charge collected. Any channel without at least 400 laser hits is flagged.

After processing all of the triggers, we first determine the average delay between the laser firing and hits being detected. To do this, we add the histograms of laser time from each channel and fit to a gaussian in a ± 15 ns window around the maximum bin. This value, minus the nominal 650 ns delay, will be the applied “PMT” offset for the logical channels. (Alternatively, we could define each channel’s PMT offset time relative to this mean.) Then, for each physical channel, we determine the the maximum bin of its laser timing histogram, and fit the range out to the first local minimum on each side to a gaussian. This value, minus the 650 ns nominal laser peak, is the channel’s PMT offset time, and is subtracted from the calculated hit time in all further analysis. If the PMT offset time is more than ± 12 ns, it is flagged.

Next we determine the single photoelectron (spe) response by fitting the charge histogram (over a range 10–150 counts) to a function consisting of two gaussians for one- and two-photon responses:

$$F(x) = A_1 N(x; \mu, \sigma) + A_2 N(x; 2\mu, \sqrt{2}\sigma), \quad (\text{A.11})$$

where $N(x; \mu, \sigma)$ is a normal distribution with mean μ and standard deviation σ . The μ obtained from this fit is assumed to be the spe response. In further analysis, when we want to know the number of photoelectrons (npe) detected by a channel during a hit, we divide the charge collected by μ .

A.1.3 Channel and Hit Rejection

The following summarizes the reasons a channel may be completely rejected from analysis for a given run:

- Fewer than 1000 hits from all pulser triggers.

- Misfire or multifire on more than 10% of laser triggers
- Triangle wave minimum value higher than 20 ADC counts or maximum value less than 235 ADC counts, or height of first or last local maximum of ramp histogram less than 5 counts.
- Triangle wave maximum amplitude outside the range of the ADC (first maximum at 0 or last maximum at 255 counts).
- Unable to reliably determine slope of triangle wave on more than 1/3 of pulser hits.
- Measured d80 value between ramp1 and ramp2 samples outside the range 80 ± 15 ns
- Average gray offset more than ± 2 counts.
- Mean remaining offset in pulser hit times more than ± 10 ns or RMS of offset histogram greater than 10 ns.
- Dark rate greater than 5 kHz.
- PMT offset time greater than ± 12 ns, or the fit to determine the offset did not converge, or less than 400 total hits were collected during laser triggers.

Once all channels failing any of the above requirements have been rejected, we define the variable `live_pmts` as the number of remaining channels.

Channels that fail any of the above cuts are not considered at all in further analysis. We define a separate class of “bad charge” channels that are used for timing but not energy reconstruction if they fail any of the following cuts:

- Mean single photoelectron response outside the range 15–45 counts, or spe fit does not converge.
- RMS of single photoelectron response greater than 25 counts.
- Fewer than 400 hits were acquired during all laser triggers.
- More than 10% of laser hits had charge below 10 ADC counts.

Similar to the `live_pmts` variable, we define a `good_charge_chans` variable that counts the number of channels passing all of these criteria. Figure A.3 shows the number of live PMTs and good charge channels for a subset of runs (from 7000–10000). Unfortunately,

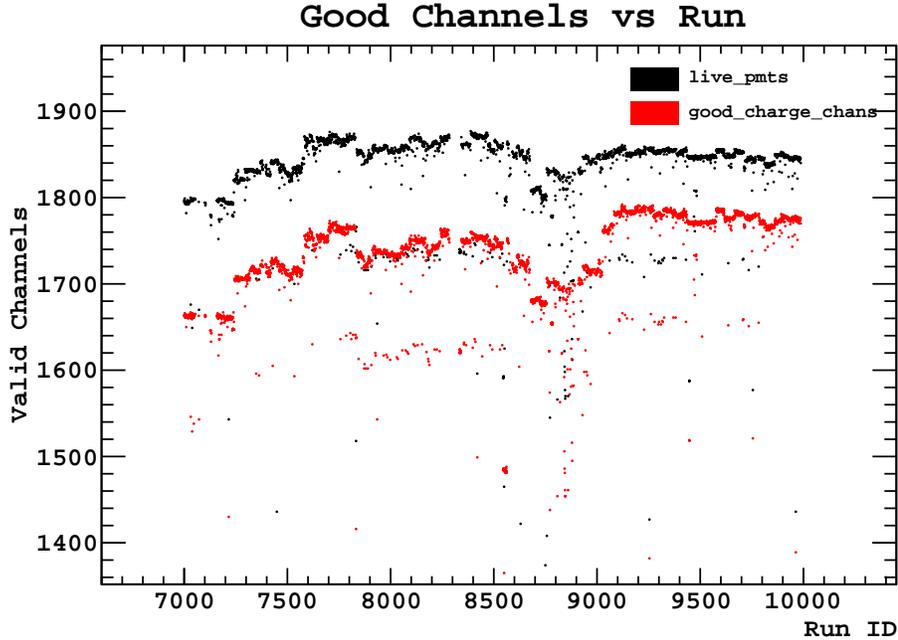


Figure A.3: Number of valid channels (separated into good timing and good charge reconstruction) for runs 7000-10000.

a significant fraction (upwards of 20%) of the channels in Borexino exhibit problems with charge collection, but do have adequate timing performance; we found the loss of precision by simply removing these channels altogether to be unacceptable. Instead, we use these bad charge channels for timing only, but we do not use the measured charge when determining the energy deposited during an event. This will be explained further in Section A.2.2.

Individual hits in a given trigger can be flagged as invalid for any of the following reasons:

- They originate from an invalid channel.
- The reconstructed time of the hit is outside the normal trigger window.
- The hit time is within 180 ns of a previous hit on the same channel. This cut removes spurious “retriggers” that frequently occur immediately after a channel becomes live after the previous hit.
- The hit time is 615 ± 10 ns after the previous hit on the same channel. This cut removes spurious hits due to cable reflections.

Finally, the number of photons detected during a hit on a single channel (npe; integrated charge divided by mean single photoelectron charge) is labeled invalid if it comes from a bad charge channel or if the measured charge is less than 10% of the spe mean.

A.2 Event reconstruction

A.2.1 Clustering

In the mach4 framework, the term “clustering” refers to the process of identifying grouped “clusters” of photons from scintillation emission within each trigger window. This both removes stray dark hits from energy tallies and separates distinct interactions that happen to have occurred close together, either due to pileup or decay coincidences, such as the ^{214}Bi - ^{214}Po coincidence, which is separated by only a few hundreds of nanoseconds. For the neutrino analysis, we reject triggers that contain more than one cluster, as decay coincidences are a source of background and the accidental pileup rate is small, given the few Hz trigger rate. Tagging coincidences is useful, however, for removing some classes of background and estimating the rates of others.

The clustering algorithm is based on that outlined in the thesis of Daniela Manuzio [130]. First, all of the valid hit times are histogrammed into 16 ns bins. For regular inner detector triggers, we define the start of the first cluster as the first group of 3 bins with more than 9 total hits. The end of the pulse is marked by the first window of 6 bins with fewer hits than the average expected from the dark rate plus 3 times the standard deviation in the expected dark rate. The expected dark rate is calculated from the rate measured for all valid channels during the random calibration triggers. With average total dark rates usually on the order of 1 MHz, the expected number of dark hits is ~ 0.1 , so in practice the end is usually the first set of 6 bins with no hits. Then, to ensure that we do not miss late-arriving photons in the tails of scintillation pulses, we add a further 32 bins (512 ns) after the identified end,

unless a second pulse start is found in that region. Only clusters with 20 or more hits are considered in later analysis.

A.2.2 Energy reconstruction

In mach4 we provide two different energy definitions: `nhits`, which simply counts the total number of hits in a cluster, and `npe`, which sums the total charge collected from each channel normalized by that channel's mean single photoelectron response. For events where the average number of photoelectrons detected per channel is significantly less than 1, `nhits` is the best energy variable, since `npe` will include noise from the intrinsic width of the single photoelectron response plus other electronics noise. As the number of photons gets larger, `nhits` more consistently underestimates the true energy of the event, since more channels will have detected multiple photoelectrons but only be counted for a single hit.

We define a separate variable, `nhits_corrected`, to correct statistically for this multiple photoelectron effect. The derivation of this correction is straightforward. Assume we have an event that generates n_0 photoelectrons with N live PMTs. Assuming all channels are equal, the probability for a given channel to detect any given photoelectron is

$$P = \frac{1}{N}. \tag{A.12}$$

If n_0 photoelectrons are detected, then the probability for a channel to receive no photoelectrons is

$$P_0 = \left(1 - \frac{1}{N}\right)^{n_0}, \tag{A.13}$$

and the probability to receive a hit with one or more photoelectrons is

$$P_{\geq 1} = 1 - P_0 = 1 - \left(1 - \frac{1}{N}\right)^{n_0}. \tag{A.14}$$

The expected number of hits is then

$$n = NP_{\geq 1} = N \left(1 - \left(1 - \frac{1}{N} \right)^{n_0} \right). \quad (\text{A.15})$$

If we invert this expectation value, we obtain a formula for estimating the average value of photoelectrons that would result in a given number of hits:

$$n_0 = \frac{\log(1 - n/N)}{\log(1 - 1/N)} \quad (\text{A.16})$$

The final value is rescaled to a nominal 2000 live PMTs to correct for the changing number of live PMTs run per run.

$$\text{nhits_corrected} = \left(\frac{2000}{\text{live_pmts}} \right) \cdot \frac{\log(1 - \text{nhits}/\text{live_pmts})}{\log(1 - 1/\text{live_pmts})} \quad (\text{A.17})$$

This correction makes the `nhits_corrected` variable approximately linear with event energy, but does not account for the large uncertainty due to the method employed of inverting the expectation value. It also relies on the assumption that all of the PMTs are equal, and ignores effects such as differing trigger and quantum efficiencies (generally small factors) and different solid angles (small factors for events near the center of the detector, but can be large for events further outside.)

Because of the large fraction of otherwise valid channels with poor charge reconstruction (see Figure A.3), the uncorrected `npe` variable is also noisy. Simply scaling the `npe` variable to account for the differing values of good charge channels does not take into account the exponential decay of the scintillation pulse, which means that multiple photoelectron hits are much more likely to happen at the beginning of the scintillation pulse, and almost all hits in the tails will be single photoelectrons, largely independent of the total size of the pulse. To account for this, for each hit with invalid charge, we assign to it a number of photoelectrons equal to the average number detected by good charge channels in a 10 ns-wide window

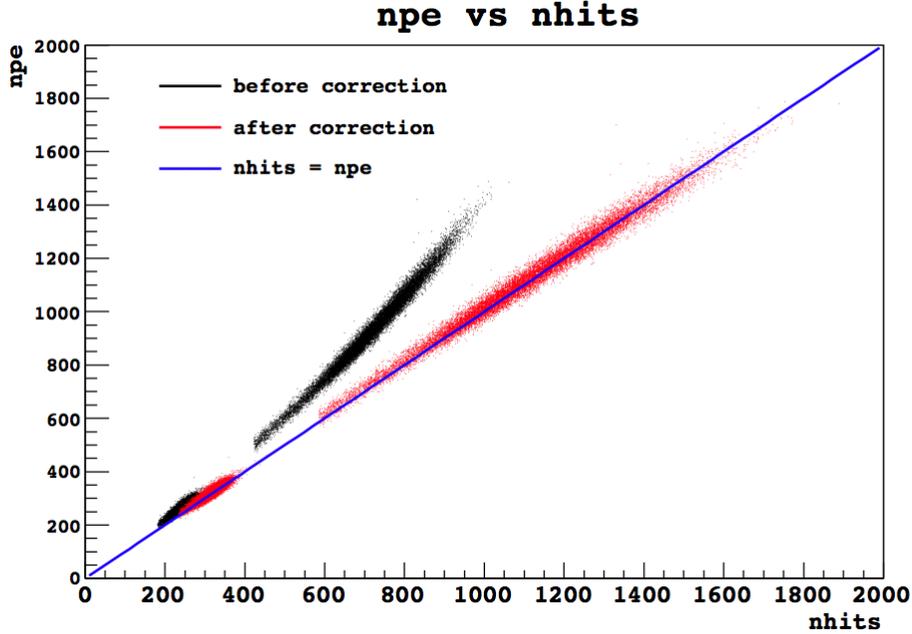


Figure A.4: Comparison of the two energy variables in mach4, npe and nhits, after scaling and statistical correction. Thanks to A. Chavarria.

centered on the invalid hit. If no hits with good charge were acquired in that window, we assign it an npe of 1. Only after we have applied this window averaging correction do we also scale for the number of live PMTs.

Figure A.4 plots these two variables (after correction and scaling) against each other for data from a ^{14}C and radon calibration source at the center of the detector. As the figure shows, the two variables are essentially equivalent after correction. We most often use the npe (`_corrected`) variable. To demonstrate the stability of the energy response, Figure A.5 shows a fit to the ^{210}Po alpha peak, the most prominent feature in the spectrum, as a function of time. (See Section A.4 for more details.) With a few exceptions, the `nhits_corrected` variable fluctuates less than 5% throughout data-taking.

A.2.3 Position Reconstruction

The position and time of each cluster are reconstructed using a maximum likelihood function. In essence, we pick a test point in the scintillator, compute the time of flight for photons to

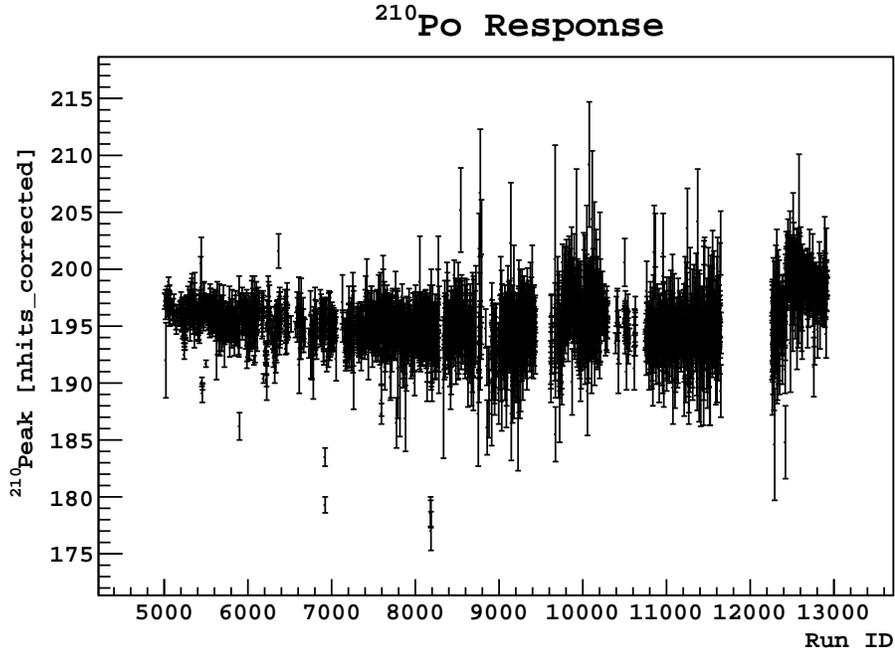


Figure A.5: Fit of the ^{210}Po alpha peak over time, showing the stability of the nhits_corrected variable.

reach the PMTs that registered hits from the test point, subtract off this time of flight from the measured hit times, and compare the resulting distribution to the scintillator emission PDF. Then we repeat with new test points until we find the point for which the time-of-flight subtracted hit times most closely match the scintillation PDF.

If we neglect hit pileup (i.e., assume we record a hit for each photoelectron), the effects of scattering in the scintillator, and the finite size of the PMT detector surface, then the likelihood function is

$$\mathcal{L}(\mathbf{x}, t) = \prod_{i=0}^{nhits} p\left(t_i - t - \frac{n}{c}|\mathbf{x}_i - \mathbf{x}|\right), \quad (\text{A.18})$$

where \mathbf{x} and t are the position and time of the scintillation emission, $p(t)$ is the scintillation PDF, \mathbf{x}_i is the position of the PMT that detected the i th hit at time t_i , c is the speed of light, and n is the index of refraction of the scintillator (and buffer). The reconstructed position and time of the event are the coordinates \mathbf{x} and t at which \mathcal{L} is a maximum. For most efficient computation, in practice we actually find the minimum of $-\log \mathcal{L}$ using MINUIT.

The scintillation PDF was measured using dedicated test cells, and fit to an analytical function consisting of two exponentials. When computing the time-of-flight-subtracted pulse, the timing resolution of the PMTs and electronics remains convolved in, so the final emission PDF is

$$p(t) = N(t; 0, \sigma) * \left(\frac{A_1}{\tau_1} e^{-t/\tau_1} + \frac{1 - A_1}{\tau_2} e^{-t/\tau_2} \right), \quad (\text{A.19})$$

where $N(t; 0, \sigma)$ is a normal distribution with mean 0 and standard deviation σ representing the electronics timing resolution, $*$ denotes convolution, τ_i is the time constant of exponential component i , and the A_1 is the relative weight of the first exponential. The measured values for τ_1 and τ_2 are 0.7415 ns and 13.8752 ns, respectively, with $A_1 = 0.5509$.

When performing reconstruction using Equations (A.18) and (A.19), we found that events were generally reconstructed $\sim 10\%$ further from the origin than the actual position. This was first shown from the apparent vessel radius from reconstructed ^{14}C and ^{210}Po events and later confirmed with calibration sources. Two modifications were necessary to the fitting procedure. First, we use an effective index of refraction of $n=1.68$ compared to the measured scintillator index of ~ 1.53 . Initially, the effective index was used as a simple empirical correction to obtain roughly the correct position. A more recent analysis suggests that, since the index of refraction actually varies slowly over the scintillator emission wavelengths, the higher effective index is actually the correct one for the photon group velocity, while the smaller measured index is appropriate for the phase velocity [131].

The second correction is a modification of the emission PDF to account for photon pileup. Specifically, if multiple photoelectrons are detected by a single PMT, we only record a single “hit” with $\text{npe} > 1$. So, we are less likely to record separate hits for photons that arrive later in the pulse, and the time-of-flight-subtracted pulse is biased toward earlier times. To correct for this, we separate PDFs for hits with 1, 2 or ≥ 3 npe and modify Equation (A.18) accordingly. Figure A.6 shows the three different PDFs currently used in mach4. The standard deviation of the reconstructed positions after these corrections, as measured by

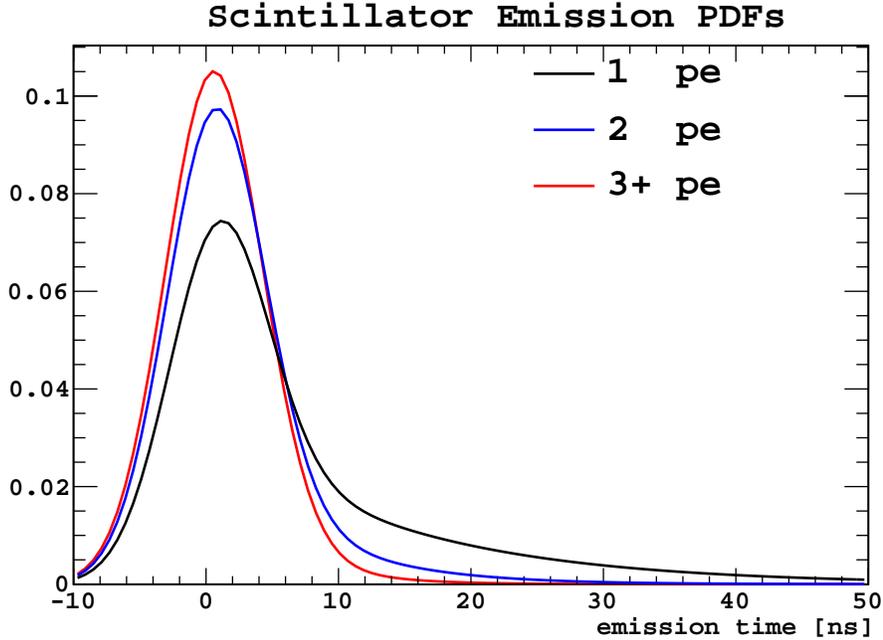


Figure A.6: Effective scintillator emission PDFs for 1, 2 and ≥ 3 detected photoelectrons, including the smearing from the detector response function.

calibration sources, is ~ 12 cm at 1 MeV near the center of the vessel, and slightly larger at lower energies and higher radii.

A.3 Cuts and Event Selection

Unlike counting experiments, there is no signal in Borexino that uniquely identifies the ${}^7\text{Be}$ or other neutrinos. Rather, the ${}^7\text{Be}$ rate is determined by fitting the expected spectrum of ${}^7\text{Be}$ interactions in the scintillator plus the known backgrounds to the total measured spectrum. The ability of the fit to accurately determine the ${}^7\text{Be}$ interaction rate is therefore dependent on the signal to background ratio over the ${}^7\text{Be}$ energy range. The ${}^7\text{Be}$ spectrum in Borexino, as well as the contributions due to the other low-energy solar neutrinos, is shown in Figure A.7. Although the other neutrinos must be accounted for in the total spectral fit, the low rates do not present a significant background to ${}^7\text{Be}$.

The most prominent backgrounds to the ${}^7\text{Be}$ neutrino in Borexino are

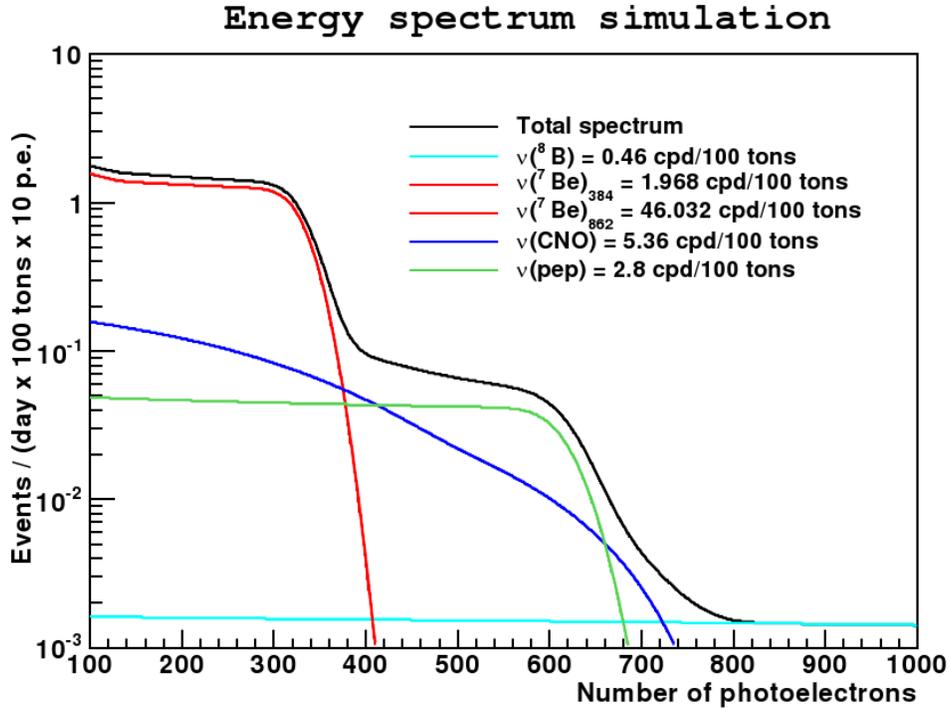


Figure A.7: Expected spectrum from solar neutrinos in Borexino [132]. The relative amplitudes are from Ref. [133], and the spectra account for the energy resolution of Borexino.

- ^{14}C beta decays ($Q=156 \text{ keV}$), intrinsic to the organic liquid scintillator. ^{14}C decays account for the majority of the trigger rate in Borexino. The low endpoint energy is below most of the ^7Be spectrum.
- ^{85}Kr beta decays ($Q=687 \text{ keV}$). ^{85}Kr is naturally present in the atmosphere and found in the scintillator at a rate of $\sim 25 \text{ cpd}/100 \text{ tons}$, introduced by a small leak that occurred during initial scintillator processing. ^{85}Kr is the most serious background, as the spectrum is very similar to ^7Be , and the decay rate is of the same order as the ^7Be interaction rate. With sufficient statistics, the ^7Be spectrum can be differentiated due to the sharp “shoulder” at the end-point energy. Limits for the total ^{85}Kr rate can be estimated by identifying coincidences due to rare decays to the metastable $^{85\text{m}}\text{Rb}$ coincidences, discussed in more detail in Section A.3.4
- ^{11}C $\beta+$ decays. ^{11}C is created cosmogenically from ^{12}C , either directly by photo-nuclear evaporation interactions with cosmic ray muons or indirectly through muon shower

products and (n,2n) reactions. The ~ 20 minute half-life prevents simple vetoing of the parent muon. The two 511 keV annihilation gammas from the positron emission bring the ^{11}C spectrum above that of the ^7Be neutrino, but, due to the finite energy resolution of the detector, the lowest energy decays do affect the high energy endpoint of the ^7Be spectrum.

- ^{10}C β^+ decays. ^{10}C is also cosmogenically created; it has a half-life of 19 s, and always decays to an excited state of ^{10}B , which then emits a gamma (718 keV 98.5% of the time, 1021 keV 1.5%). The minimum energy for a ^{10}C decay in the detector is therefore 1.74 MeV, so it only needs to be considered when studying the spectrum at higher energies.
- Other cosmogenic sources. These are effectively reduced by vetoing any event within 2 ms following a muon detection in the water tank.
- External gamma rays, primarily from the PMTs and stainless steel sphere. These are effectively removed by a fiducial cut that accepts only events within 3 m of the vessel center and with $|z| < 1.7$ m. The z cut is to remove events emanating from the nylon vessel endcaps, which are one of the major background sources.
- Daughters of ^{222}Rn decay. Radon, like krypton, is present naturally in the atmosphere. Radon daughters, especially polonium, also frequently contaminate surfaces such as the nylon vessel, and diffuse into the scintillator from there. The alpha sources, ^{214}Po and ^{210}Po can be removed by application of the Gatti parameter, discussed in Section A.3.3. Some of the remaining backgrounds can be effectively removed by detecting the ^{214}Bi - ^{214}Po coincidence, discussed in Section A.3.2. Some part of the scintillator distillation and purification procedures breaks the secular equilibrium of the various decay products; in particular, the rate of alpha decays from ^{210}Po is much larger than that of its parent, ^{210}Bi . In fact ^{210}Po is the second most dominant source of decays in Borexino after ^{14}C .

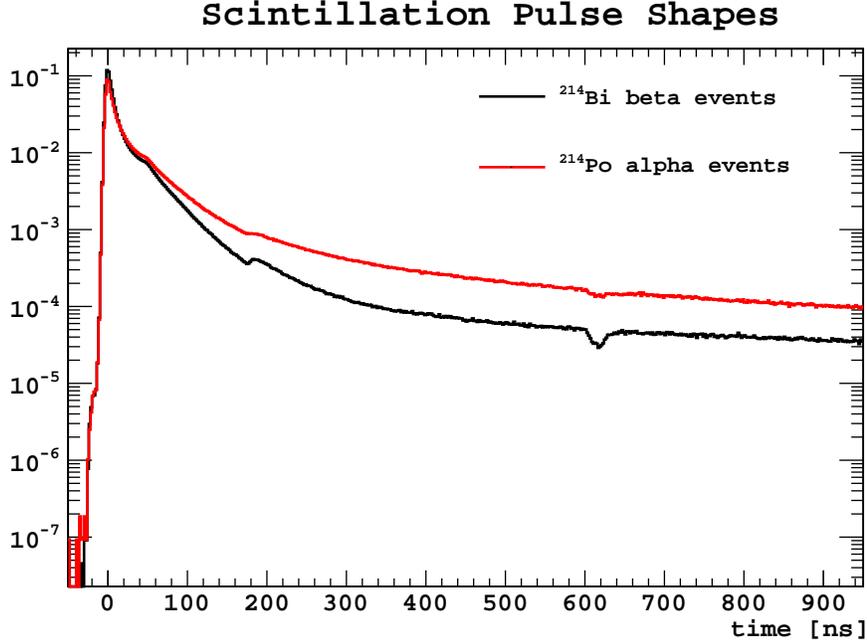


Figure A.8: Average scintillation pulse shapes from ^{214}Bi beta and ^{214}Po alpha decays tagged by their coincidence. The dips at ~ 150 ns and ~ 650 ns are due to the retriggering and cable reflection cuts, respectively.

A.3.1 Alpha-beta discrimination and the Gatti parameter

The time profiles of the scintillation response in the Borexino liquid scintillator are slightly different for alpha and beta decays, as shown in Figure A.8. We can use this difference to identify and reject alpha backgrounds in the spectrum. Because the difference between the pulse shapes is small, we cannot generally identify alphas on an event-by-event basis, but must rely on statistical methods to determine the average numbers.

We use two parameters to characterize the alpha-like quality of a pulse: the tail-to-total ratio (ttr) and the gatti parameter. The ttr is defined as the fraction of hits in a cluster that arrive after some time t_0 ; early investigations determined that $t_0 = 70$ ns gave the best separation.

Even better separation is obtained using Gatti's optimum filter method [134]. Gatti's filter requires reference PDFs for the average response for alpha and beta decays, $\bar{\alpha}_i$ and $\bar{\beta}_i$, where the pulse shapes are binned into discrete times Δt_i . Based on these reference shapes,

we define a weight for each bin

$$P_i = \frac{\bar{\alpha}_i - \bar{\beta}_i}{\bar{\alpha}_i + \bar{\beta}_i} \quad (\text{A.20})$$

The weights P_i are positive where the alpha PDF is higher than the beta, and negative otherwise. The amplitude of the weights is largest where the PDFs are most significantly different. Given these weights, the gatti parameter for a given event is given by

$$G \equiv \frac{1}{N} \sum_i P_i S_i \quad (\text{A.21})$$

where N is the total number of hits and S_i is the number of hits in bin i . Pulses from beta events will have a negative gatti parameter, while alpha decays will have a positive gatti parameter.

To define the reference shapes, we search for pairs of events consistent with the ^{214}Bi - ^{214}Po coincidence. ^{214}Bi beta-decays ($Q=3.3$ MeV) to ^{214}Po , which then quickly undergoes alpha decay ($Q=7.8$ MeV) with a half-life of $164 \mu\text{s}$. As we will show in Section A.4, the light yield is ~ 500 pe/MeV for beta events and ~ 50 pe/MeV for alpha decays. The ^{214}Po alpha spectrum is a narrow peak at ~ 400 pe while the ^{214}Bi is the usual broad beta spectrum. So we tag as a bipo coincidence any two clusters with intervening time 20 – $500 \mu\text{s}$ and reconstructed within 60 cm of each other in which the first cluster has energy 100 – 1750 pe and the second cluster has energy 300 – 500 pe. These tagged clusters provide convenient alpha and beta reference shapes.

A.3.2 Radon daughter removal

Since the ^{214}Bi - ^{214}Po coincidence is such a distinguishable feature in the data, and given the relatively short (~ 27 minutes) half-life of ^{214}Pb , the parent of ^{214}Bi , we can use the ^{214}Bi - ^{214}Po events to remove most ^{214}Pb decays as well as the ^{214}Bi , and ^{214}Po events. In practice, because the ^{214}Po rate is small and well separated due to the gatti parameter, we do not utilize the coincidence for the radon removal. Instead, we define any event with energy 300 –

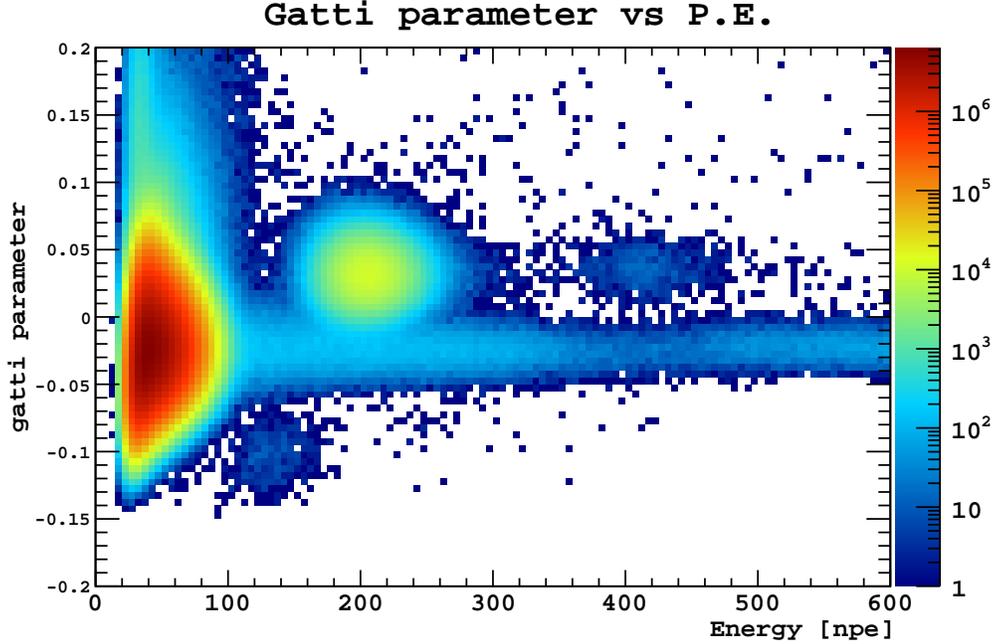


Figure A.9: Spectrum of gatti parameter vs energy for events in the fiducial volume.

500 pe and with $\text{gatti} > 0$ as a ^{214}Po decay, then remove any event over the previous 3 hours with reconstructed position within 80 cm of the ^{214}Po . This cut should remove essentially all of the ^{214}Pb and ^{214}Bi , except in cases where the ^{214}Po decay occurs during the dead time in between triggers or in the gaps between runs. The effect of the cut on the total live mass-time for a given spectrum is estimated by taking the ratio of events in the fiducial volume in a region of the ^{14}C spectrum before and after the cut, since ^{14}C has a much higher rate than other sources and is uniformly distributed throughout the vessel.

A.3.3 Statistical alpha subtraction

Figure A.9 shows the spectrum of events in gatti-energy space after fiducial volume cuts, but before radon daughter removal. In the figure, the small group of events with $\text{gatti} \sim -0.1$ between 100–160 pe are caused by electronic noise and are effectively removed by the cut accepting only triggers with a single cluster. The two regions with $\text{gatti} \sim 0.04$ around 200 and 420 pe are due to ^{210}Po and ^{214}Po , respectively. Finally, the large group below 100 pe is

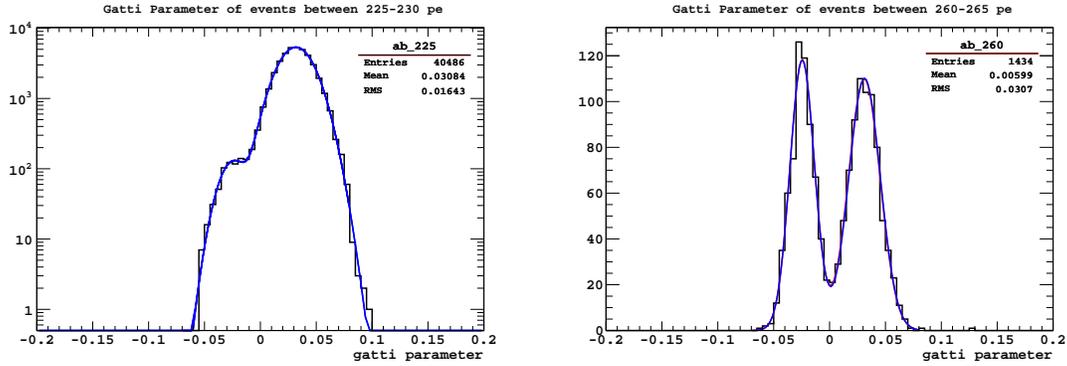


Figure A.10: Example fits of gatti parameter distributions to two gaussians representing the population of alpha and beta events in that energy bin. Left: 225–230 pe, a region dominated by ^{210}Po alphas. Right: 260–265 pe, a region with approximately equal alpha and beta populations. Note the scale on the independent axis is linear in the right figure and logarithmic on the left.

mostly ^{14}C . Because of the high rate of ^{210}Po decays, it is not possible to remove them all using a hard cut on the gatti parameter without removing most of the beta-like events as well. We therefore have adopted a statistical subtraction approach, which proceeds in three steps.

First, we separate all the events into 5 pe-wide bins, and, for each bin, histogram the gatti parameter spectrum for the events in that bin. Then we fit the gatti spectrum in each bin to two gaussians, one for betas and one for alphas. As examples, the results of two of these fits, one bin that contains primarily alphas and one that contains roughly equal amounts of alphas and betas, are shown in Figure A.10.

The mean for the alpha and beta distributions obtained by the fit for each bin is shown in Figure A.11. Ideally the mean should not be a function of energy, but some small variation is not unexpected due to effects like multiple hits on a single PMT, etc. In regions where one species (alphas or betas) is very dominant, the fit can be biased, as shown by the green points in the beta distribution under the ^{210}Po peak and the magenta points on the edges of the alpha distribution in Figure A.11. To correct for this effect, we fit both distributions to a linear function and repeat the fits of step one with the means fixed to the linear fit. We fit for the means only in regions where we expect “good” behavior from the step 1 results:

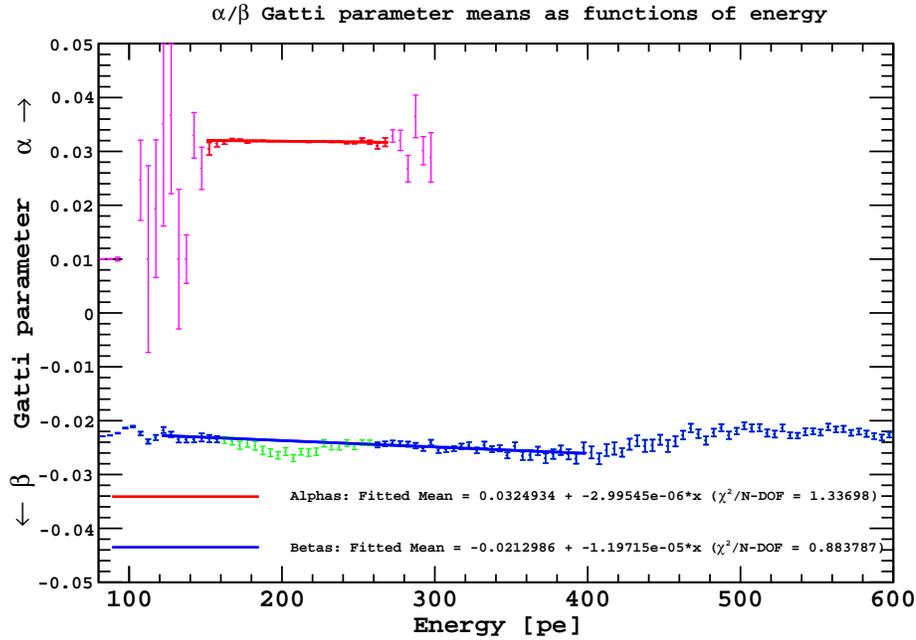


Figure A.11: Mean gatti parameter for alpha and beta events in each energy bin obtained from fit to two gaussians. The red and blue points are fit to linear functions to fix the mean of the gaussians during the second stage of statistical alpha subtraction (see Section A.3.3).

only under the ^{210}Po peak for the alphas (red points), and below the ^{11}C energy, excluding the ^{210}Po range (green points) for the betas.

Recently a third step has been added that attempts to correct for biases introduced due to the fact that the true shape of the Gatti distributions assuming Poisson fluctuations in the scintillation are not perfectly gaussian. This correction is largest in bins where the two populations are roughly equal, but is less than 2% maximum and generally less than 1%.

Figure A.12 shows the spectrum after each of the major cuts: fiducial volume, radon daughter removal, and statistical alpha subtraction.

A.3.4 ^{85}Kr Tagging

Most of the time, ^{85}Kr decays by beta emission with endpoint energy 687 keV directly to the ground state of ^{85}Rb . This type of decay is indistinguishable from elastic scattering from ^7Be neutrinos and cannot be vetoed. However, in 0.43% of decays, the krypton decays to a

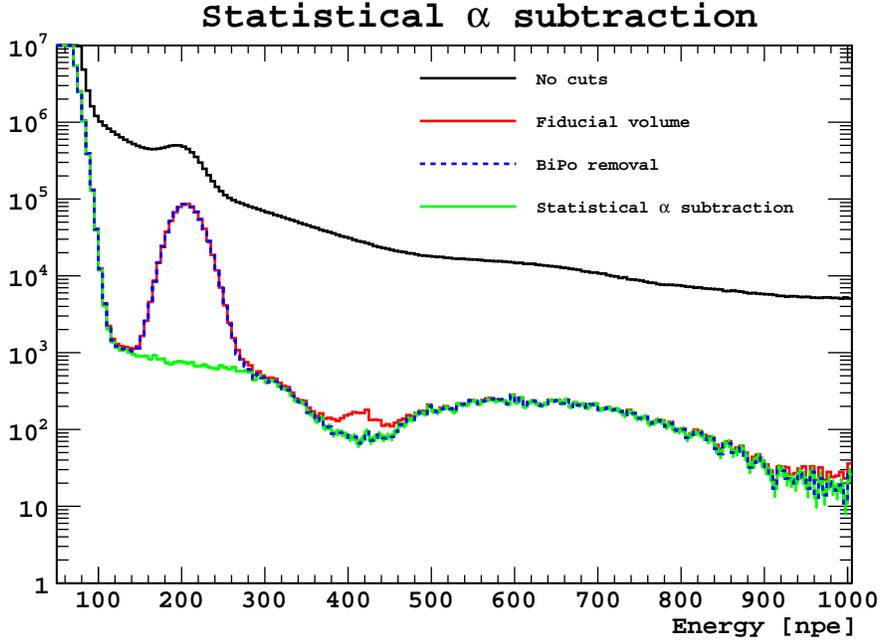


Figure A.12: The spectrum in Borexino after each major event selection cut: No cuts applied (black), reconstructed position with 3 m of the center of the detector and $|z| < 1.7$ m (red), removal of ^{214}Pb , ^{214}Bi , and ^{214}Po by vetoing a small volume prior to each ^{214}Po event (blue dashed), and statistical subtraction of ^{210}Po alphas using the gatti parameter (green).

metastable state $^{85\text{m}}\text{Rb}$, which then decays to the ground state by emitting a monoenergetic 514 keV gamma with a half life of 1.46 μs . We tag these decays by identifying two clusters within 6 μs and reconstructed within 1.5 m where the second cluster (the rubidium gamma) has energy 140–280 pe and $\text{gatti} < 0.02$ and the first has energy less than 130 pe. Once the total tagging efficiency of the cuts ($\sim 20\%$) and the rate of accidental coincidences (primarily ^{14}C followed by ^{210}Po) is evaluated ($\sim 10^{-2}$ cpd/100 tons), the total rate for all ^{85}Kr decays in the detector is estimated to be $\sim 30 \pm 5$ cpd/100 tons [135]. This rate is in good agreement with that obtained by fit to the total spectrum, discussed in Section A.4.

A.4 Spectral Fit

The final step in determining the ^7Be neutrino interaction rate is to fit the spectrum after all cuts in Figure A.12 to the expected ^7Be spectrum plus the spectra of all remaining

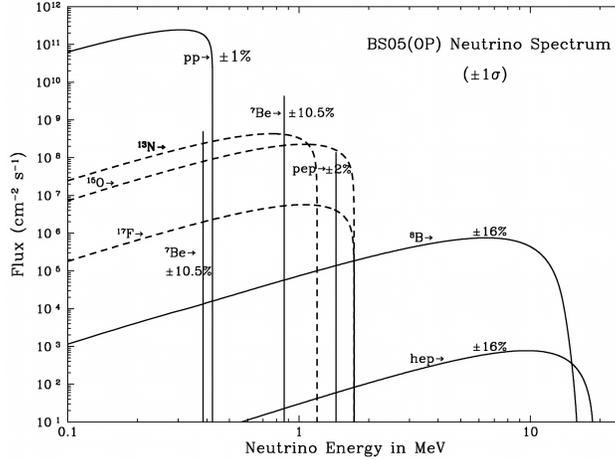


Figure A.13: Predicted flux and spectrum from solar neutrinos according to the BS05(Op) model [13].

backgrounds. The light yield of the scintillator (and the resolution as a function of energy) and the relative weights of all of the different input spectra are free parameters in the fit.

A.4.1 ${}^7\text{Be}$ Interaction Spectrum

Figure A.13 shows the predicted flux and spectrum for the various solar neutrino species from the BS05(OP) solar model [13]. Below a few MeV, the region of interest for Borexino, only the ${}^7\text{Be}$, pp, pep, and CNO cycle (dashed lines in Figure A.13) make a significant contribution to the spectrum. The ${}^7\text{Be}$ electron capture process has two branches, emitting an 862 keV neutrino in 89.6% of reactions and a 384 keV neutrino the remaining 10.4% [136]. We detect the neutrinos in Borexino by elastic scattering off of electrons in the scintillator, $\nu + e^- \rightarrow \nu + e^-$. The Feynman diagrams for this reaction are shown in Figure A.14. All three flavors of neutrino can scatter via the neutral-current interactions; the cross-section for ν_e scattering is higher due to the additional possibility of charged-current interactions. Capture of electron neutrinos on a neutron requires a threshold energy of 17 MeV for ${}^{12}\text{C}$ and 2.2 MeV for ${}^{13}\text{C}$, and so is not possible for these low-energy neutrinos [136].

Figure A.14: Neutrino-electron elastic scattering diagrams. All three neutrino flavors (e , μ , and τ) have equal amplitudes for the neutral-current process (left), but, due to flavor conservation, only electron-neutrinos participate in the charged-current diagram (right).

Kinematically, a neutrino with energy E_ν can produce recoiling electrons with a maximum energy (in units where $c=1$)

$$E_{\max} = \frac{E_\nu}{1 + m_e/(2E_\nu)}. \quad (\text{A.22})$$

For solar neutrino energies, which are well below the W and Z masses, the differential cross section as a function of energy for the diagrams in Figure A.14 is given by

$$\frac{d\sigma}{dE} = \frac{\sigma_0}{m_e} \left[g_L^2 + g_R^2 \left(1 - \frac{E}{E_\nu} \right)^2 - g_L g_R \frac{m_e E}{E_\nu^2} \right], \quad (\text{A.23})$$

where $\sigma_0 = 2G_F^2 m_e^2 / (\pi \hbar^4) = 8.81\text{E-}45 \text{ cm}^2$. The g 's are related to the W and Z coupling constants to left- and right-handed chirality particles; $g_R = \sin^2 \theta_w \approx 0.222$ for all neutrinos; $g_L = \sin^2 \theta_w + 1/2 \approx 0.722$ for electron neutrinos and $\sin^2 \theta_w - 1/2 \approx -0.278$ for μ and τ neutrinos [136]. The hard cut-off at E_{\max} results in the distinct Compton-like edge in the ${}^7\text{Be}$ recoil spectrum at 665 keV.

For the mono-energetic neutrinos, Equation (A.23) is directly proportional to the interaction spectrum. For the neutrino sources whose fluxes have a spectral distribution, we convolve this distribution with Equation (A.23) to obtain the final spectrum

$$\frac{dR}{dE} = n_e V \int \frac{d\Phi}{dE_\nu} \frac{d\sigma}{dE} dE_\nu \quad (\text{A.24})$$

where n_e is the electron number density and V the total volume of the target.

All of the neutrinos generated in the sun are emitted as electron neutrinos, but can oscillate to the other flavors as they travel from the sun to the earth. Although the oscillations

involve all three neutrino flavors, current experimental limits on the upper limit of the mixing angle θ_{13} allow us to approximate the process as involving only two neutrino states. In this case, the probability for an initially electron-flavor neutrino with energy E_ν after it has traveled a distance L through vacuum to interact as an electron neutrino is [136]

$$P_{ee} = 1 - \sin^2(2\theta_{12}) \sin^2 \left(1.27 \Delta m_{12}^2 \frac{L}{E_\nu} \frac{\text{GeV}}{\text{eV}^2 \text{km}} \right). \quad (\text{A.25})$$

In the case of a distributed neutrino source, such as solar neutrinos produced roughly evenly throughout the core of the sun, if $\delta L \gg E/\Delta m^2$ (i.e., several oscillation wavelengths), then the second term in Equation (A.25) averages to 1/2, and $P_{ee} \rightarrow 1 - \sin^2(2\theta_{12})/2 \approx 0.57$ for $\sin^2(2\theta_{12}) = 0.87$ [136, 137].

Solar neutrinos do not travel all the way to Borexino in vacuum, but must first escape the sun. As they pass through the dense matter (in particular high electron density) of the sun, electron neutrinos interact more with the electrons due to the charged current interaction. This process, the MSW effect, changes the Hamiltonian for propagation and hence the oscillation parameters [138]. At energies below ~ 2 MeV, the MSW effect is small for solar neutrinos, and the oscillations are near that of vacuum propagation. At much higher energies, neutrinos exit the sun as almost pure ν_2 (the second mass eigenstate), so that $P_{ee} \rightarrow \sin^2 \theta_{12} \approx 0.32$.

For the 862 keV ${}^7\text{Be}$ neutrino, the survival probability after accounting for the MSW effect becomes $P_{ee} = 0.54$ [4]. Therefore, when calculating the expected interaction spectrum using Equation (A.23), we obtain for the average value for the coefficient g_L

$$\langle g_L \rangle = 0.262; \quad \langle g_L^2 \rangle = 0.317$$

A.4.2 Energy Response Function

The expected spectrum in units of energy (generally MeV) for the solar neutrino spectra is derived according to the previous section. The spectra for the remaining cosmogenic backgrounds are obtained from the Table of Isotopes and the National Nuclear Data Center [112]. These spectra are in units of energy, whereas the spectrum measured by Borexino is in units of photoelectrons (npe). To convert between the two, we use Birk's law [139]

$$\Delta L = \frac{Y_0 \Delta E}{1 + k_B \cdot (dE/dx)}, \quad (\text{A.26})$$

where L is the signal in npe, Y_0 is the photon yield (photons emitted per unit energy deposited) for a minimum-ionizing particle in the limit $dE/dx \rightarrow 0$, and k_B , the Birks constant, is an empirical constant that is a property of the scintillator. Y_0 for this scintillator has been measured to be $\sim 11,500$ photons/MeV [136], and k_B is 0.0115 ± 0.007 . The dE/dx curve for pseudocumene is shown in Figure A.15. To find the actual number of photons emitted by a particle with total energy E , we take the continuum limit of Equation (A.26) and integrate from E until the particle has lost all energy:

$$L(E) = Y_0 \int_0^\ell \frac{Y_0}{1 + k_B (dE/dx)} (dE/dx) dx \quad (\text{A.27})$$

where ℓ is the range of the particle and is given by

$$\ell = \int_0^E \frac{1}{(dE'/dx)} dE' \quad (\text{A.28})$$

Finally, the total number of photoelectrons detected is $QL(E)$, where Q is the probability for a given photon to hit a PMT photocathode, be converted into a photoelectron, and produce a (digitized) current pulse above threshold. Based on the location of the ${}^7\text{Be}$ shoulder in Figure A.12 at ~ 300 pe, we know that the total yield $Y \equiv QY_0 \sim 500$ photoelectron/MeV. This value is determined more precisely by the spectral fit (discussed in the following section).

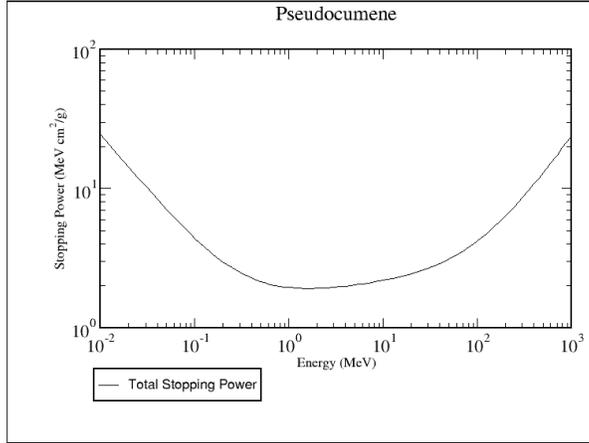


Figure A.15: dE/dx curve for electrons in pseudocumene, the Borexino liquid scintillator.

The preceding analysis informs us only about the average number of detected photoelectrons. The actual number of detected photoelectrons in a given event will vary based on the Poisson statistics of photon counting and on the resolution of the electronics chain. We model these resolution effects as a gaussian distribution with standard deviation

$$\sigma_n = a + b\sqrt{n}, \tag{A.29}$$

where n is the average number of detected photoelectrons and a and b are free parameters of the spectral fit.

Therefore, the spectrum in npe for a given source with energy spectrum dR/dE is given by

$$\frac{dR(n)}{dn} = \int N(n; QL(E), \sigma_{\text{npe}}) \frac{dR}{dE} dE, \tag{A.30}$$

where n is the number of photoelectrons and $N(x; \mu, \sigma)$ is a normal distribution with mean μ and standard deviation σ .

If we label each different species included in the fit with an index i , then the total number of counts expected in a bin with range n to $n + \Delta n$ npe is

$$N = MT \sum_i \int_n^{n+\Delta n} \frac{dR_i(n')}{dn'} dn', \quad (\text{A.31})$$

where M is the mass of the fiducial volume (78.5 tons) and T is the live time over which the spectrum to be fit was acquired. So, for each step of the fit, we calculate N for each bin of the final spectrum in Figure A.12; MINUIT then seeks the set of parameters that maximize the agreement between the two based on a likelihood or χ^2 function. The free parameters in the fit are:

- Total decay rate for each species in the spectrum. The species considered are
 - ${}^7\text{Be}$ neutrinos
 - ${}^{14}\text{C}$
 - ${}^{85}\text{Kr}$
 - ${}^{210}\text{Bi}$
 - ${}^{11}\text{C}$
 - ${}^{10}\text{C}$ if the upper limit of the fit range is high enough
- Other species are included in the generated spectrum, but their rates are fixed:
 - ${}^{14}\text{C}$ pileup (two events at approximately the same time that are reconstructed as a single event; rate calculated from the ${}^{14}\text{C}$ rate.
 - pp neutrinos (fixed to standard solar model (SSM) prediction)
 - pep neutrinos (fixed to SSM prediction)
 - ${}^{214}\text{Pb}$ (fixed by estimating the unvetoes rate from ${}^{214}\text{Po}$ tagging)
- Light collection parameters:
 - Y , the limit of photoelectron yield for a minimum-ionizing particle.
 - Y_γ , the limit of light yield for gammas (sets the start point of the ${}^{11}\text{C}$ spectrum).
 - the a and b resolution parameters from Equation (A.29).

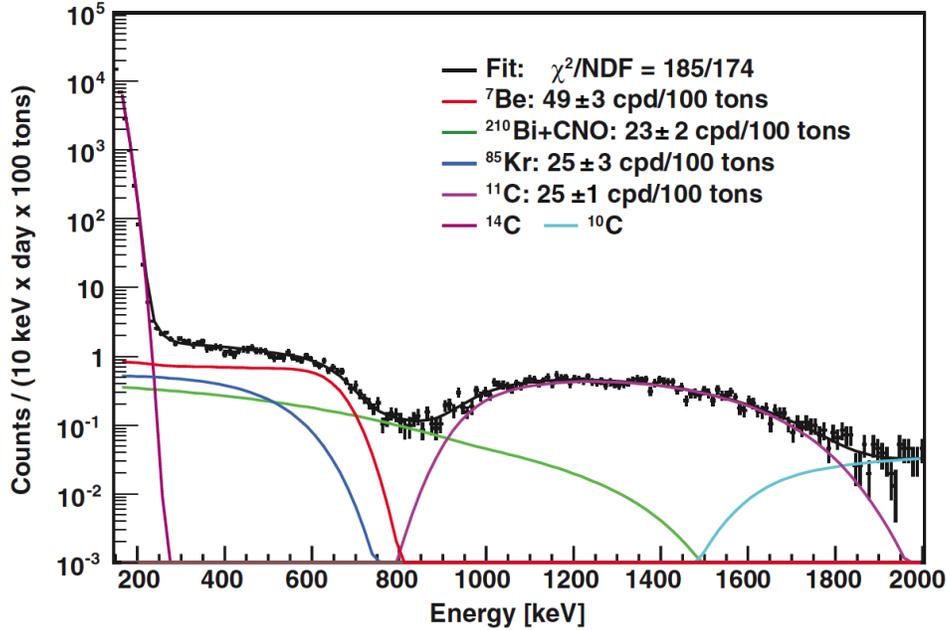


Figure A.16: Results of fit to spectrum of 41.3 ton-years of Borexino data [4].

- For consistency, we will sometimes fit the spectrum before statistical alpha subtraction; in this case there are two further parameters for the rate and light yield of ^{210}Po decays. The CNO neutrino interactions also have a measurable effect on the total spectrum, but their spectrum is almost indistinguishable from that of ^{210}Bi . Therefore, we treat the fitted amplitude for the ^{210}Bi spectrum as the total rate for $^{210}\text{Bi} + \text{CNO}$ interactions.

A.4.3 Fit results

Figure A.16 reproduces the results of the spectral fit from Ref. [4], which covers 41.3 ton-years of data. The ^{14}C pileup, ^{214}Pb and pp and pep neutrinos are not shown for clarity. The rate returned for ^7Be neutrino interactions is 49 ± 3 counts per day per 100 tons of scintillator. The ^{85}Kr rate from the fit is 25 ± 3 cpd/100 tons, in good agreement with the rate predicted from the $^{85}\text{Kr} \rightarrow ^{85\text{m}}\text{Rb}$ coincidence measurements (30 ± 5 cpd/100 tons). The value of χ^2 (185 with 174 degrees of freedom) is higher than optimal; the excess is dominated

by the fit to the ^{14}C spectrum, which is not modeled in the table of isotopes with sufficient resolution for the very high rate in Borexino.

A systematic uncertainty of 8.5% is estimated for this fit. The systematic uncertainty is dominated by the fiducial volume determination and the detector response function. The fiducial volume is determined by finding the fraction of some class of event that is distributed evenly throughout the inner vessel that passes the fiducial cut; then we multiply this fraction by the total mass of scintillator. We use ^{14}C and 2.2 MeV gamma-rays from cosmogenic neutrons capturing on hydrogen as our distributed sources. The fiducial volume estimate when using ^{14}C tends to be low due to external gamma rays that begin to populate the ^{14}C toward the edge of the inner vessel. On the other hand, the fiducial volume estimation tends to be high when using the 2.2 MeV capture gammas because of the long range (10–20 cm) of gammas of this energy; some fraction of them escape toward the vessel edge. The uncertainty due to this method of fiducial volume determination is estimated at 6% [4]. More recently, the accuracy of the position reconstruction has been determined over the course of dedicated calibration campaigns, and the uncertainty in the fiducial volume has been reduced to $\sim 1\%$. The energy-response uncertainty is dominated by the near-degeneracy of the ^7Be and ^{85}Kr spectra; any small change in the energy response can cause a large trade-off in the ^7Be and ^{85}Kr rates. This uncertainty is also estimated at 6%.

Including the systematic uncertainty, the ^7Be neutrino rate in Borexino for the 41.3 ton-year dataset considered in Ref. [4] was determined to be $49 \pm 3(\text{stat}) \pm 4(\text{sys})$ counts per day per 100 tons of scintillator. The predicted rate, using the neutrino oscillation parameters in Ref. [137] and the solar model from Ref. [13], is 48 ± 4 cpd/100 tons, in very good agreement.

Appendix B

daqman File Format

Two different formats have been used for the raw binary data output of `daqman`. All of the data taken in the first two campaigns of DarkSide-10 (the analysis discussed in Chapter 4) is saved in the older, “legacy” format. Data acquired after May 2011 is saved in the newer format.

B.1 Legacy format

The legacy format has no global header, and all events are in a single file. All of the writing for the legacy events is compressed on-the-fly using the `zlib` compression library. Each event has a header consisting of 3 32-bit words, followed by the raw data unchanged from the CAEN digitizer(s). The header data are

- Event size in bytes
- Event ID (index iterated for each event saved)
- Timestamp (standard unix `time_t`, the number of seconds elapsed since January 1, 1970, UTC).

The data from each CAEN board consists of a header of 4 32-bit words, followed by all the recorded samples for each channel in order, as shown in Figure B.1 [14].

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	1	0	EVENT SIZE																								HEADER			
BOARD-ID				RES	0	PATTERN												CHANNEL MASK													
reserved				EVENT COUNTER																											
TRIGGER TIME TAG																															
0	0	0	0	SAMPLE [1] - CH[0]												0	0	0	0	SAMPLE [0] - CH[0]									DATA CH0		
0	0	0	0	SAMPLE [3] - CH[0]												0	0	0	0	SAMPLE [2] - CH[0]											
...																															
0	0	0	0	SAMPLE [N-1] - CH[0]												0	0	0	0	SAMPLE [N-2] - CH[0]											
0	0	0	0	SAMPLE [1] - CH[1]												0	0	0	0	SAMPLE [0] - CH[1]									DATA CH1		
0	0	0	0	SAMPLE [3] - CH[1]												0	0	0	0	SAMPLE [2] - CH[1]											
...																															
0	0	0	0	SAMPLE [N-1] - CH[1]												0	0	0	0	SAMPLE [N-2] - CH[1]											
...																															
0	0	0	0	SAMPLE [1] - CH[7]												0	0	0	0	SAMPLE [0] - CH[7]								DATA CH7			
0	0	0	0	SAMPLE [3] - CH[7]												0	0	0	0	SAMPLE [2] - CH[7]											
...																															
0	0	0	0	SAMPLE [N-1] - CH[7]												0	0	0	0	SAMPLE [N-2] - CH[7]											

Figure B.1: Raw data format for a trigger from a single CAEN V1720 digitizer [14].

B.2 New format

The new format is designed to be flexible and updateable in the future, and so contains information about the size and version number of global and event headers. Events in a single data acquisition run can be split into multiple files. The header information (both global and per-event headers) are not compressed. The new file format begins with a global header with 13 32-bit words:

- Magic number: 0xdecoded1. Identifies the file as being the new format.
- Size of global header in bytes.
- Global header version number (currently 1).
- Size of per-event header in bytes.

- Per-event header version number (currently 1).
- Total size of the file in bytes.
- Start time: unix timestamp for when the file was first created.
- End time: timestamp for when the acquisition period finished.
- Run ID: unique identifier for the run this file belongs to.
- File index: serial number for the place of this file in a run split into multiple files. The first file in a run has index 0.
- Number of events stored in this file (NOT in this run).
- Event ID min: ID of the first event stored in this file.
- Event ID max: ID of the last event stored in this file.

Each event in the new format has an event header followed by several “datablocks,” which are designed to hold different types of data that can be defined dynamically. The event header is 4 32-bit words:

- Event size in bytes.
- Event ID.
- Timestamp (Unix timestamp as in the legacy format).
- Number of datablocks in the event.

Each datablock has its own sub-header of 3 32-bit words, followed by the actual data for the event, which is compressed using the zlib library. The header consists of:

- Total blocksize on disk: the actual size in bytes of the compressed data plus the header.
- Data size: the size of the UNcompressed data segment.
- Type: a numerical identifier for what kind of data is in the block. Currently defined types are CAEN_V172X (0) and MONTECARLO (1). (These are defined in the header for the RawEvent class).

Bibliography

- [1] D. Acosta-Kane, et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **587**, 46 (2008), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/S0168900207024692>.
- [2] B. Loer, Journal of Physics: Conference Series **136**, 042084 (2008), URL <http://stacks.iop.org/1742-6596/136/i=4/a=042084>.
- [3] B. Loer, *The darkside-10 prototype dark matter detector*, Contributed talk at APS April Meeting 2011. (2011), URL <http://meetings.aps.org/link/BAPS.2011.APR.H8.4>.
- [4] C. Arpesella, et al. (Borexino Collaboration), Phys. Rev. Lett. **101**, 091302 (2008).
- [5] D.-M. Mei and A. Hime, Phys. Rev. D **73**, 053004 (2006).
- [6] C. Arpesella, et al., Astroparticle Physics **18**, 1 (2002), ISSN 0927-6505, URL <http://www.sciencedirect.com/science/article/pii/S0927650501001797>.
- [7] M. Chadwick, et al., Nuclear Data Sheets **107**, 2931 (2006), URL <http://www.sciencedirect.com/science/journal/00903752>.
- [8] J. Sebastian, S. A. Peter, and R. V. Jasra, Langmuir **21**, 11220 (2005), URL <http://pubs.acs.org/doi/abs/10.1021/la051460e>.
- [9] J. Kayser and K. Knaebel, Chemical Engineering Science **41**, 2931 (1986), cited By (since 1996) 9, URL <http://www.scopus.com/inward/record.url?eid=2-s2.0-0022862531&partnerID=40&md5=410fa57ae6c588df04cfbd49513d1194>.
- [10] M. Berger, et al., *Estar, pstar, and astar: Computer programs for calculating stopping-power and range tables for electrons, protons, and helium ions*, [Online] Available: <http://physics.nist.gov/Star> [2011, May 21]. National Institute of Standards and Technology, Gaithersburg, MD. (2005).
- [11] P. Cennini, et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **345**, 230 (1994), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/0168900294909962>.

- [12] Argonne National Laboratory Environmental Science Division, *Natural decay series: Uranium, radium, and thorium*, [Online] Available: <http://www.ead.anl.gov/pub/doc/natural-decay-series.pdf> [2011 June 21] (2005).
- [13] J. N. Bahcall, A. M. Serenelli, and S. Basu, *The Astrophysical Journal Letters* **621**, L85 (2005), URL <http://stacks.iop.org/1538-4357/621/i=1/a=L85>.
- [14] *Technical Information Manual. Revision n. 14. Mod. V1720 8 Channel 12 Bit 250 MS/s digitizer.*, CAEN (2010).
- [15] V. Trimble, *ANNUAL REVIEW OF ASTRONOMY AND ASTROPHYSICS* **25**, 425 (1987), ISSN 0066-4146.
- [16] G. Bertone, D. Hooper, and J. Silk, *Physics Reports* **405**, 279 (2005).
- [17] M. Roos, arXiv:1001.0316 (2010).
- [18] S. Dodelson, *Modern Cosmology* (Academic Press, 2003).
- [19] N. Jarosik, et al., *The Astrophysical Journal Supplement Series* **192**, 14 (2011), URL <http://stacks.iop.org/0067-0049/192/i=2/a=14>.
- [20] K. G. Begeman, A. Broeils, and R. Sanders, *Mon. Not. R. Astr. Soc.* **249**, 523 (1991).
- [21] D. J. Fixsen, et al., *The Astrophysical Journal* **473**, 576 (1996), URL <http://stacks.iop.org/0004-637X/473/i=2/a=576>.
- [22] D. Larson, et al., *The Astrophysical Journal Supplement Series* **192**, 16 (2011), URL <http://stacks.iop.org/0067-0049/192/i=2/a=16>.
- [23] D. J. Eisenstein, et al., *The Astrophysical Journal* **633**, 560 (2005), URL <http://stacks.iop.org/0004-637X/633/i=2/a=560>.
- [24] S. Burles, K. M. Nollett, and M. S. Turner, arXiv:astro-ph/9903300 (1999).
- [25] A. Friedmann, *General Relativity and Gravitation* **31**, 1991 (1999).
- [26] A. Friedmann, *General Relativity and Gravitation* **31**, 2001 (1999).
- [27] S. Burles and D. Tytler, *The Astrophysical Journal* **499**, 699 (1998), URL <http://stacks.iop.org/0004-637X/499/i=2/a=699>.
- [28] S. Burles, K. M. Nollett, and M. S. Turner, *Phys. Rev. D* **63**, 063512 (2001).
- [29] D. Perkins, *Particle Astrophysics, Second Edition (Oxford Master Series in Physics)* (Oxford University Press, USA, 2009), ISBN 0199545464.
- [30] K. Nakamura et al. (Particle Data Group), *Journal of Physics G* **37**, 075021 (2010).
- [31] G. Blanger, A. Pukhov, and G. Servant, *Journal of Cosmology and Astroparticle Physics* **2008**, 009 (2008), URL <http://stacks.iop.org/1475-7516/2008/i=01/a=009>.

- [32] R. D. Peccei and H. R. Quinn, Phys. Rev. Lett. **38**, 1440 (1977).
- [33] L. J. Rosenberg and K. A. van Bibber, Physics Reports **325**, 1 (2000), ISSN 0370-1573, URL <http://www.sciencedirect.com/science/article/pii/S0370157399000459>.
- [34] E. Arik, et al., Journal of Cosmology and Astroparticle Physics **2009**, 008 (2009), URL <http://stacks.iop.org/1475-7516/2009/i=02/a=008>.
- [35] S. J. Asztalos, et al. (The ADMX collaboration), Phys. Rev. Lett. **104**, 041301 (2010).
- [36] I. Aitchison, *Supersymmetry in Particle Physics: An Elementary Introduction* (Cambridge University Press, 2007), ISBN 0521880238.
- [37] G. Bertone, ed., *Particle Dark Matter: Observations, Models and Searches* (Cambridge University Press, 2010).
- [38] J. L. Feng, Annual Review of Astronomy and Astrophysics **48**, 495 (2010), <http://www.annualreviews.org/doi/pdf/10.1146/annurev-astro-082708-101659>.
- [39] R. H. Helm, Phys. Rev. **104**, 1466 (1956).
- [40] J. D. Lewin and P. F. Smith, Astroparticle Physics **6**, 87 (1996), ISSN 0927-6505, URL <http://www.sciencedirect.com/science/article/B6TJ1-3W9102K-9/2/2bc9c6f41f509a2ff99e7e8b56ecb07c>.
- [41] J. Diemand, et al., Nature **454**, 735 (2008), URL <http://dx.doi.org/10.1038/nature07153>.
- [42] M. C. Smith, et al., Monthly Notices of the Royal Astronomical Society **379**, 755 (2007), ISSN 1365-2966, URL <http://dx.doi.org/10.1111/j.1365-2966.2007.11964.x>.
- [43] C. Savage, et al., Journal of Cosmology and Astroparticle Physics **2009**, 036 (2009), URL <http://stacks.iop.org/1475-7516/2009/i=09/a=036>.
- [44] J. Feng, *Dark particles*, Presentation at APS April meeting. [Online] Available: <http://hep.ps.uci.edu/~jlf/research/presentations/1105aps.pdf> [7 July 2011] (2011).
- [45] F. Halzen and S. Klein, Review of Scientific Instruments **81**, 081101 (2010).
- [46] C. de los Heros (IceCube Collaboration), arXiv:1012.0184 (2010).
- [47] *Atic collaboration*, <http://atic.phys.lsu.edu/index.html>.
- [48] *Fermi collaboration*, <http://www-glast.stanford.edu/>.
- [49] *Hess collaboration*, <http://www.mpi-hd.mpg.de/hfm/HESS/pages/collaboration/>.
- [50] *Pamela collaboration*, <http://pamela.roma2.infn.it/>.
- [51] *Ams collaboration*, <http://ams.nasa.gov/index.html>.

- [52] J. Chang, et al., *Nature* **456**, 362 (2008), ISSN 0028-0836, URL <http://dx.doi.org/10.1038/nature07477>.
- [53] O. Adriani, et al., *Nature* **458**, 607 (2009), ISSN 0028-0836, URL <http://dx.doi.org/10.1038/nature07942>.
- [54] O. Adriani, et al. (PAMELA Collaboration), *Phys. Rev. Lett.* **102**, 051101 (2009).
- [55] D. Grasso, et al., *Astroparticle Physics* **32**, 140 (2009), ISSN 0927-6505, URL <http://www.sciencedirect.com/science/article/pii/S0927650509001078>.
- [56] N. Arkani-Hamed, et al., *Phys. Rev. D* **79**, 015014 (2009).
- [57] J. L. Feng, M. Kaplinghat, and H.-B. Yu, *Phys. Rev. D* **82**, 083525 (2010).
- [58] D. P. Finkbeiner, et al., *Journal of Cosmology and Astroparticle Physics* **2011**, 002 (2011), URL <http://stacks.iop.org/1475-7516/2011/i=05/a=002>.
- [59] V. Barger, et al., *Physics Letters B* **678**, 283 (2009), ISSN 0370-2693, URL <http://www.sciencedirect.com/science/article/pii/S0370269309007035>.
- [60] E. Behnke, et al. (COUPP Collaboration), *Phys. Rev. Lett.* **106**, 021303 (2011).
- [61] S. Archambault, et al., *Physics Letters B* **682**, 185 (2009), ISSN 0370-2693, URL <http://www.sciencedirect.com/science/article/pii/S0370269309013525>.
- [62] E. Aprile, et al. (XENON100 Collaboration), arXiv:1104.2549 (2011), 1104.2549, URL <http://arxiv.org/abs/1104.2549>.
- [63] The CDMS II Collaboration, *Science* **327**, 1619 (2010), URL <http://www.sciencemag.org/content/327/5973/1619.abstract>.
- [64] R. Bernabei, et al. (DAMA Collaboration), *The European Physical Journal C - Particles and Fields* **56**, 333 (2008), URL <http://dx.doi.org/10.1140/epjc/s10052-008-0662-y>.
- [65] D. Hooper, et al., *Phys. Rev. D* **82**, 123509 (2010).
- [66] C. E. Aalseth, et al. (CoGeNT Collaboration), *Phys. Rev. Lett.* **106**, 131301 (2011).
- [67] C. E. Aalseth, et al., arXiv:1106.0650 (2011).
- [68] L. Grandi, Ph.D. thesis, University of Pavia (????).
- [69] T. Doke, K. Masuda, and E. Shibamura, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **291**, 617 (1990), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/016890029090011T>.
- [70] A. Hitachi, et al., *Phys. Rev. B* **23**, 4779 (1981).

- [71] D. Gastler, et al., arXiv:1004.0373 (2010), 1004.0373.
- [72] A. Hitachi, et al., Phys. Rev. B **27**, 5279 (1983).
- [73] A. Hitachi, T. Doke, and A. Mozumder, Phys. Rev. B **46**, 11463 (1992).
- [74] P. Benetti, et al. (WARP Collaboration), Astroparticle Physics **28**, 495 (2008), ISSN 0927-6505, URL <http://www.sciencedirect.com/science/article/pii/S0927650507001016>.
- [75] W. H. Lippincott, et al., Phys.Rev.C **78**, 035801 (2008), 0801.1531.
- [76] M. G. Boulay, et al., arXiv:0904.2930 (2009), 0904.2930.
- [77] P. Benetti, et al. (WARP Collaboration), Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **574**, 83 (2007), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/S0168900207001672>.
- [78] R. Acciarri, et al., Nuclear Physics B - Proceedings Supplements **197**, 70 (2009), ISSN 0920-5632, 11th Topical Seminar on Innovative Particle and Radiation Detectors (IPRD08), URL <http://www.sciencedirect.com/science/article/pii/S0920563209007580>.
- [79] R. Acciarri, et al., Journal of Instrumentation **5**, P05003 (2010), URL <http://stacks.iop.org/1748-0221/5/i=05/a=P05003>.
- [80] W. Walkowiak, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **449**, 288 (2000), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/S0168900299013017>.
- [81] T. Pollmann, M. Boulay, and M. Kuniak, arXiv:1011.1012 (2010).
- [82] P. Mosteiro and C. Galbiati, Tech. Rep., Princeton University (2008).
- [83] A. Wright, et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **644**, 18 (2011), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/S0168900211007364>.
- [84] V. Culler and R. W. Waniek, Phys. Rev. **95**, 585 (1954).
- [85] G. Bellini, et al. (Borexino Collaboration), Journal of Instrumentation **6**, P05005 (2011), URL <http://stacks.iop.org/1748-0221/6/i=05/a=P05005>.
- [86] M. Berger, et al., *Xcom: Photon cross section database*, [Online] Available: <http://www.nist.gov/pml/data/xcom/index.cfm> [2011, June 3]. National Institute of Standards and Technology, Gaithersburg, MD. (2010), URL <http://www.nist.gov/pml/data/xcom/index.cfm>.

- [87] C. Arpesella, Nuclear Physics B - Proceedings Supplements **28**, 420 (1992), ISSN 0920-5632, URL <http://www.sciencedirect.com/science/article/pii/0920563292902079>.
- [88] B. Bonnevier, Plasma Physics **13**, 763 (1971), URL <http://stacks.iop.org/0032-1028/13/i=9/a=007>.
- [89] R. C. Jones and W. H. Furry, Rev. Mod. Phys. **18**, 151 (1946).
- [90] T. Tokuda, T. Ando, and K. Fukui, Journal of Applied Physics **41**, 2854 (1970).
- [91] H. Loosli, Earth and Planetary Science Letters **63**, 51 (1983), ISSN 0012-821X, URL <http://www.sciencedirect.com/science/article/pii/0012821X83900213>.
- [92] D.-M. Mei, et al., Phys. Rev. C **81**, 055802 (2010).
- [93] H. Loosli, B. Lehmann, and W. Balderer, Geochimica et Cosmochimica Acta **53**, 1825 (1989), ISSN 0016-7037, URL <http://www.sciencedirect.com/science/article/pii/0016703789903037>.
- [94] D. M. Ruthven, S. Farooq, and K. S. Knaebel, *Pressure Swing Adsorption* (Wiley-VCH, 1993), ISBN 0471188182.
- [95] K. Knaebel and F. Hill, Chemical Engineering Science **40**, 2351 (1985), cited By (since 1996) 41, URL <http://www.scopus.com/inward/record.url?eid=2-s2.0-0022213047&partnerID=40&md5=56c8bef19a8433281f3b839c34787a00>.
- [96] G. W. Miller, K. S. Knaebel, and K. G. Ikels, AIChE Journal **33**, 194 (1987), ISSN 1547-5905, URL <http://dx.doi.org/10.1002/aic.690330204>.
- [97] <http://www.airsep.com/>, URL <http://www.airsep.com/>.
- [98] G. C. Ray and E. O. Box, Industrial & Engineering Chemistry **42**, 1315 (1950), URL <http://pubs.acs.org/doi/abs/10.1021/ie50487a021>.
- [99] E. L. Pace and A. R. Siebert, The Journal of Physical Chemistry **64**, 961 (1960), URL <http://pubs.acs.org/doi/abs/10.1021/j100837a001>.
- [100] C. H. Lally, et al., Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms **117**, 421 (1996), ISSN 0168-583X, URL <http://www.sciencedirect.com/science/article/pii/0168583X96003187>.
- [101] J. Brodsky, Tech. Rep., Princeton University (2011), URL <http://darkside-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=97>.
- [102] H. R. Vega-Carrillo, et al., Applied Radiation and Isotopes **57**, 167 (2002), ISSN 0969-8043, URL <http://www.sciencedirect.com/science/article/pii/S0969804302000830>.

- [103] J. Brodsky, Tech. Rep., Princeton University (2011), URL <http://darkside-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=86>.
- [104] R. Dossi, et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **451**, 623 (2000), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/S0168900200003375>.
- [105] S. Kubota, et al., Phys. Rev. B **17**, 2762 (1978).
- [106] D. Alton et al. (DarkSide collaboration), *Darkside-50: A direct search for dark matter with new techniques for reducing background*, Proposal submitted to NSF (2011), URL http://darkside-docdb.fnal.gov:8080/cgi-bin/RetrieveFile?docid=82&version=7&filename=DS50_DOE_2011.pdf.
- [107] A. Teymourian, et al., arXiv:1103.3698 (2011), 1103.3689.
- [108] C. Arpesella, et al., Physics Letters B **658**, 101 (2008), ISSN 0370-2693, URL <http://www.sciencedirect.com/science/article/pii/S0370269307011732>.
- [109] G. Alimonti, et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **406**, 411 (1998), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/S0168900298000187>.
- [110] S. Agostinelli, et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **506**, 250 (2003), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/S0168900203013688>.
- [111] W. Wilson, et al., *SOURCES 4A: A Code for Calculating (α, n), Spontaneous Fission, and Delayed Neutron Sources and Spectra* (LANL, 1999), LA-13639-MS.
- [112] *National nuclear data center*, NuDat2 database. [Online] Available: <http://www.nndc.bnl.gov/nudat2> [2011, June 20].
- [113] D. Leonard et al., NIMA **591**, 490 (2008).
- [114] *ILIAS database on radiopurity of materials*, [Online] Available: <http://radiopurity.in2p3.fr> [2011, June 20].
- [115] F. Calaprice and E. de Haas, *Gran Sasso γ measurements*, Presentation slides.
- [116] R. Brunetti, F. Carbonara, et al. (2004), experimental Proposal.
- [117] G. Bellini, et al. (1991), borexino collaboration.
- [118] G. Bellini, et al. (Borexino Collaboration), Phys. Rev. D **82**, 033006 (2010).
- [119] G. Bellini, Nuclear Physics B (Proc. Suppl.) **48**, 363 (1996).

- [120] D. Alton et al. (DarkSide collaboration) (2009), DarkSide collaboration.
- [121] Hamamatsu Corporation (2008), private communication.
- [122] P. Mosteiro (2008), experimental Project.
- [123] H. Wulandari, et al., *Astroparticle Physics* **22**, 313 (2004), ISSN 0927-6505, URL <http://www.sciencedirect.com/science/article/pii/S0927650504001276>.
- [124] C. Arpesella, *Applied Radiation and Isotopes* **47**, 991 (1996), ISSN 0969-8043, proceedings of the International Committee for Radionuclide Metrology Conference on Low-level Measurement Techniques, URL <http://www.sciencedirect.com/science/article/pii/S0969804396000978>.
- [125] F. Arneodo, et al., *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **418**, 285 (1998), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/S0168900298006792>.
- [126] J. Hong, et al., *Astroparticle Physics* **16**, 333 (2002), ISSN 0927-6505, URL <http://www.sciencedirect.com/science/article/pii/S0927650501001141>.
- [127] M. Hannam and W. Thompson, *Nucl. Instr. Meth. A* **431**, 239 (1999).
- [128] O. Buchmueller, et al., arXiv:1102.4585 (2011).
- [129] G. Alimonti, et al., *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **600**, 568 (2009), ISSN 0168-9002, URL <http://www.sciencedirect.com/science/article/pii/S016890020801601X>.
- [130] D. Manuzio, Ph.D. thesis, Universita degli Studi di Genova (????).
- [131] B. Caccianiga, Y. Koshio, and S. Manecki, *Estimation of the fv systematic error*, Internal collaboration memo (2011).
- [132] Borexino pep working group, Tech. Rep. (2011), URL https://borex.lngs.infn.it/Doc/general_meetings/2011-07/pep%20analysis%20report.pdf.
- [133] A. M. Serenelli, W. C. Haxton, and C. Pena-Garay, arXiv:1104.1639 (2011).
- [134] E. Gatti and F. D. Martini, *Nuclear Electronics* **2**, 265 (1962).
- [135] S. Hardy and S. Zavatarelli, *The ^{85}kr contamination through the delayed coincidence analysis*, Internal collaboration presentation (2010).
- [136] K. McCarty, Ph.D. thesis, Princeton University (2006).
- [137] S. Abe, et al. (The KamLAND Collaboration), *Phys. Rev. Lett.* **100**, 221803 (2008).
- [138] S. P. Mikheev and A. Y. Smirnov, *Sov. J. Nucl. Phys.* **42**, 913 (1985).
- [139] J. B. Birks, *The Theory and Practice of Scintillation Counting* (Pergamon Press, 1964).