

**PRECISION MODELING OF GERMANIUM DETECTOR
WAVEFORMS FOR RARE EVENT SEARCHES**

Samuel Joseph Meijer

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in
partial fulfillment of the requirements for the degree of Doctor of Philosophy in the
Department of Physics and Astronomy.

Chapel Hill
2019

Approved by:

Jonathan Engel

Matthew Green

David C. Radford

Frank Tsui

John F. Wilkerson

©2019
Samuel J. Meijer
ALL RIGHTS RESERVED

ABSTRACT

SAMUEL J. MEIJER: Precision Modeling of Germanium Detector Waveforms For Rare Event Searches
(Under the direction of John F. Wilkerson)

Neutrinoless double-beta decay is a proposed rare nuclear event. Current generation experiments require high sensitivity designs, with the ability to remove background signals. The MAJORANA DEMONSTRATOR is a neutrinoless double-beta decay experiment using high purity p-type point contact germanium detectors. The waveforms produced by these detectors have subtle variation indicating the detailed energy and drift path information for each event. In addition, the waveforms depend sensitively on crystal impurity levels, temperature, and operating voltage. We have developed a machine learning algorithm which, given a set of calibration waveforms, can infer detector parameters. Once these parameters are known, a precision detector model can be used to fit the drift paths of individual waveforms. This method can be used for parameter estimation, and as a sensitive background rejection technique for the DEMONSTRATOR or the proposed future LEGEND experiment. Design, performance, applications, and future plans will be discussed.

ACKNOWLEDGEMENTS

There are many people to thank for helping me get through grad school and for making it a net pleasant experience as well, including many not listed here.

First, I'd like to thank my advisor John for being a strong and faithful proponent of the work of his students at every level, and for his capable ability to manage projects, including finding ample funding for just about everything.

I also appreciate the Department of Energy and by extension, the American taxpayers for their generous funding support to both train me as a scientist and to learn about some fun corners of the universe. It is a beautiful luxury to study the esoterica of neutrinos without always having direct applications in mind, and it is something I'm often conscious of.

I would also like to recognize the assistance of the computing resources at Oak Ridge National Laboratory, with the help of Robert Varner, as well as the Research Computing resources at the University of North Carolina. Additionally, Ben Shanks' previous work on this topic has been incredibly helpful, and his contributions are both recognized and appreciated.

I am also indebted to my undergraduate research advisors, Tom Bensky and Tom Gutierrez, who taught me a lot about getting real experiments to work, and what it means to do research.

I am also grateful for the time I was able to spend working with Stuart Freedman and his group as an undergraduate. The older I get, the more clear it is to me that Stuart was exceedingly patient and generous with me, and I hope to follow his lead as I work with younger researchers myself.

I owe Tom Gilliss a debt of gratitude for being a fun and productive work partner, and helping me figure out a lot of things I should have figured out years before.

I also must thank Rohan Isaac for being a source of many fruitful discussions on a wide range of technical subjects. Your like-minded willingness to think big and try things out in the lab has been greatly appreciated.

I owe Johnny Goett a special debt for introducing me to Majorana, and for asking me to consider applying to UNC to work with John Wilkerson. Who knows what other path I may have found myself on otherwise.

And finally, to my wife, Taylor, for a wild amount of patience and support.

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Award Numbers DE-AC02-05CH11231, DE-AC05-00 OR22725, DE-AC05-76RL0130, DE-AC52-06NA25396, DE-FG02-97ER41020, DE-FG02-97ER41033, DE-FG02-97ER41041, de-sc0010254, de-sc0012612, de-sc0014445, and de-sc0018060. We acknowledge support from the Particle Astrophysics Program and Nuclear Physics Program of the National Science Foundation through grant numbers MRI-0923142, PHY-1003399, PHY-1102292, PHY-1206314, PHY-1614611, PHY-1812409, and PHY-1812356. We gratefully acknowledge the support of the U.S. Department of Energy through the LANL/LDRD Program and through the PNNL/LDRD Program for this work. We acknowledge support from the Russian Foundation for Basic Research, grant No. 15-02-02919. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada, funding reference number SAPIN-2017-00023, and from the Canada Foundation from Innovation John R. Evans Leaders Fund. This research used resources provided by the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory and by the National Energy Research Scientific Computing Center, a U.S. Department of Energy Office of Science User Facility. We thank our hosts and colleagues at the Sanford Underground Research Facility for their support.

TABLE OF CONTENTS

LIST OF TABLES	x
LIST OF FIGURES	xi
LIST OF ABBREVIATIONS	xiv
LIST OF SYMBOLS	xvi
1 Introduction	1
1.1 Overview	1
1.2 Beta decay	2
1.3 Weak interactions	3
1.4 Neutrino properties	5
1.4.1 Baryogenesis via Leptogenesis	7
1.5 Double-beta decay	8
1.6 Sensitivity to double-beta decay	12
2 The MAJORANA DEMONSTRATOR	14
2.0.1 Status of the Demonstrator	16
2.0.1.1 Future of MAJORANA	16
2.1 Outline	17
3 Signal Formation	18
3.1 Semiconductors	18
3.2 Germanium Radiation Detectors	21
3.2.0.1 Detector Processing/Fabrication	21
3.2.0.2 Impurities	22

3.2.0.3	Contacts	23
3.2.0.4	Operation	24
3.2.1	Signal Formation in a Detector	25
3.2.2	Detectors for the MAJORANA DEMONSTRATOR	25
3.2.3	Characterization of MAJORANA PPC Detectors	28
3.3	PPC Germanium Detector Operation	28
3.3.0.1	Drift Velocity	31
3.3.0.2	Charge Cloud	34
3.4	Electronics	36
3.4.1	Readout Electronics for the MAJORANA DEMONSTRATOR	36
3.4.2	Modeling of installed electronics	39
3.4.3	Electronics Effects on Waveforms	39
3.5	Digitization, Readout, and Analysis	40
3.5.1	Readout and Control	42
3.6	Complete Signal Simulation	43
3.6.1	Detector Model	43
3.6.2	Field Simulations	43
3.6.3	Electronics Model	45
3.6.4	Waveform Simulations	46
3.7	Conclusions	47
4	Bayesian Inference, Fitting, and MCMC Methods	48
4.1	Overview	48
4.2	Statistical Methods	49
4.2.1	Bayesian Statistics	49
4.2.2	Markov Chains	50
4.2.2.1	Markov Chain Process	51
4.2.3	Nested Sampling	52

4.2.3.1	Nested Sampling Algorithm	53
4.2.3.2	Diffusive Nested Sampling	55
4.2.3.3	Hyperparameters	56
4.3	Fitting waveforms	56
4.3.1	Data Processing	57
4.3.2	Machine Learning Techniques	58
4.3.3	Bayesian Inference of Waveforms	59
4.3.4	Waffle	60
4.3.5	Training Fits	60
4.3.5.1	Parameters	60
4.3.5.2	Data Selection	61
4.3.5.3	Impurity	62
4.3.5.4	Number of waveforms	62
4.3.6	Waveform Fits	62
4.4	Model Testing	63
4.4.1	Bayes Factor	64
4.4.2	Cross-Validation	65
4.5	Conclusions	65
5	Performance and Results	66
5.1	Convergence Metrics	66
5.2	Training Performance	68
5.3	Waveform Fit Performance	69
5.3.1	Other Detectors and Datasets	73
5.3.2	Future Validation	75
6	Applications	76
6.1	Overview	76
6.2	Inverse Filtering	76

6.3	Future applications	78
6.3.1	Electronics Stability Performance	78
6.3.2	Energy estimation	79
6.3.3	Likelihood Cut	80
6.3.4	Compton Imaging	80
6.3.5	Other future applications	81
7	Conclusions	82
7.0.1	Results	82
7.1	Extensibility	83
7.2	Future Work	83
7.2.1	Near Term Work	83
7.2.2	Further Validation	84
7.2.3	Model Improvement	84
7.2.3.1	Improve the impurity gradient model	84
7.2.3.2	Improve Velocity Model	85
7.2.4	Extend fits	85
7.2.5	Improve Computational Performance	85
7.3	Conclusions	86
Appendix A	IMPURITY CONCENTRATION	87
Appendix B	Z-TRANSFORMS	88
B.1	The Z-Transform	88
Appendix C	DATA ACQUISITION FOR THE MAJORANA DEMONSTRATOR	90
C.1	ORCA	90
C.1.1	Objects in ORCA	90
C.2	The MAJORANA DEMONSTRATOR	90
BIBLIOGRAPHY	92

LIST OF TABLES

3.1	Velocity model parameters included for each charge type.	34
3.2	Electronics filter models considered for this work.	46
3.3	Detector models considered for this work.	46
3.4	Waveform model parameters considered for this work.	47
4.1	Hyperparameters used in DNest4, with the values used in this study indicated.	57
4.2	Example of Bayes Factor comparison between models, for two sets of models, with two sets of waveforms indicating a strong preference for model B.	65

LIST OF FIGURES

1.1	Chart of nuclides, indicating the dominant decay mode by color [1]. The line of stability zig-zags through the middle of the isotopes, near the $Z = N$ line at low mass.	2
1.2	Isobar along odd isobar $A=75$. Data from [2].	4
1.3	Isobar along even isobar $A=76$. ^{76}Ge is unable to single beta decay because both Ga and As would be energetically forbidden endpoints; double-beta decay of ^{76}Ge to ^{76}Se is allowed, but highly suppressed. Data from [2].	4
1.4	Feynman diagrams indicating the two forms of beta decay.	5
1.5	Feynman diagrams indicating double-beta decay	10
2.1	A schematic of the MAJORANA DEMONSTRATOR, indicating the primary experimental components	15
3.1	Valence and conduction bands in a semiconductor, source [47]. The interatomic distance, or <i>lattice constant</i> a is fixed for a given material, and the value it takes on contributes to the exact band structure.	19
3.2	Germanium band structure diagram, adapted from [48].	20
3.3	a) A single crystal being pulled from the melt in a czochralski-style pulling apparatus. b) The resulting boule following a czochralski pull.	22
3.4	The approximate PPC detector cross section as used by MJD. The HV conducting n^+ outer surface is indicated in pink, a nonconductive passivated surface in orange, and the small p^+ point-contact readout electrode in blue. The detector is axially symmetric about the dashed line.	26
3.5	Photograph showing the “bottom” surface of an ORTEC enriched PPC detector. Note the bevel and the small point contact.	27
3.6	Simulated weighting potential, with lines indicating the drift trajectories.	30
3.7	Simulated hole drift time map across a detector cross section. Shown for detector P42575A, at a constant z position of half the detector length. Note that the drift time contours are not azimuthally symmetric, with longer drift times at $\phi = \pi/4$ for a given radius. The map is symmetric about $\phi = \pi/4$. The hole drift time along the white constant-radius curve is shown in Figure 3.8.	32
3.8	Simulated hole drift time along axis angle ϕ . Shown for detector P42575A, at a constant z position of half the detector length and a constant radius of half the detector radius. This corresponds to the drift time along the white dashed constant-radius curve in the middle of Figure 3.7.	33

3.9	Time domain waveforms from charges in a single cloud of gaussian with FWHM of 1 mm from the bulk of an enriched detector. This charge cloud is exaggerated for visualization. The gaussian cloud compares well to a smoothed waveform with a gaussian kernel standard deviation of approximately 10 samples.	35
3.10	Color indicates the maximum difference in time at the 50 percent timepoint between waveforms in a cloud, at each point over the detector cross section. Detector is an enriched ORTEC detector, so the bevel and point contact are visible.	37
3.11	Color indicates the maximum difference in time at the 50 percent timepoint between waveforms in a cloud, at each point over the detector cross section. Detector is a BEGe, so the ditch outline is visible, with some edge effects.	37
3.12	A high-level schematic overview of the MJD signal readout chain.	38
3.13	Regions of signals waveforms discussed in this chapter. The waveform is a voltage signal in the time domain.	40
3.14	The electric field viewed cross-sectionally inside one ORTEC geometry (enriched) PPC detector (P42575A). The highest field location is near the point contact.	45
4.1	The distributions sampled by both Classic Nested Sampling and Diffusive Nested Sampling as the chain proceeds forward. The Diffusive Nested Sampling algorithm allows for continuous mixing at shells of lower likelihood. Adapted from [90].	55
4.2	Data structure for training fits. Other parameters (such as detector geometry) are not allowed to float. When fitting a single waveform, the fit uses the shared detector parameters without floating them and the single waveform parameters are fit to match that waveform. Electronics models are indicated in Table 3.6.3.	60
4.3	Waveform fit comparison for an average waveform. The horizontal scale indicates a time parameter. The vertical scale on the top figure is a voltage, and the bottom figure indicates the residual difference between a data waveform and the proposed simulated value. The data and fit overlap considerably, with a residual feature at the rising edge.	63
5.1	Beginning of level exploration for a fit using 105 particles per each of 4 waveforms over 750 levels. Each level is successively made, compressing the remaining prior mass in each subsequent level, and after all levels have been created, the particles diffuse into exploring through all existing levels.	67
5.2	Top figure shows the log-likelihood as a function of the enclosed prior mass. Bottom figure shows the weights of the saved particles, shown by their enclosed prior mass values; particles to the left of the peak are not expected to contribute to the evidence significantly.	67

5.3	Fit result of impurity at each end of the detector for detector P42575A in dataset 1. The right plots are chains of the posterior values with the mean indicated, plotted as a function of sample number. The left plots are KDEs of the indicated chain, with the median and mean estimates indicated.	68
5.4	Fit result of impurity at each end of the detector for detector P42575A in dataset 1. The three different colors each represent a training fit using 4 different waveforms, and the black line is the average of the posterior maxima.	70
5.5	Comparison of energy parameters between fit data and trapezoidal filter estimation. The two metrics are largely in agreement, although some waveforms differ by as much as 10 keV.....	71
5.6	Comparison of energy parameters between fit data and trapezoidal filter estimation, in a joint KDE plot centered over the 2615 keV photopeak.....	72
5.7	Fitted positions for 112 calibration events in a single detector. Note that radius value is squared to correct for the cylindrical volume element. Events show some higher density near the radial extreme of the detector and some clustering in a small region near the surface. The events which have slower drift times are found to correctly be located farther from the point contact ($r = 0, z = 0$).	73
5.8	Simulated event interaction locations for 630,525 calibration source events in detector P42575A, with each blue point indicating a different event. Note that radius value is squared to correct for the cylindrical volume element. A higher color density indicates a relatively higher number of events in that location. Events show some higher density near the radial extreme of the detector and the density decreases with increasing axial position. The event distribution is otherwise nearly uniform. .	74
6.1	Frequency response of forward and inverse transfer functions. The forward and inverse filters are nicely matched, indicating that much of the frequency-space power can be recovered.	77
6.2	Waveform with electronics effects removed. Note that the falling edge of the waveform is corrected to a flat top, as expected.	78
C.1	The Data Acquisition system for the Majorana Demonstrator.	91

LIST OF ABBREVIATIONS

ADC	Analog-to-Digital Converter
BEGe	Broad Energy Germanium detector
CC	Charged Current
CPU	Central Processing Unit
DL	Discovery Limit
DAQ	Data Acquisition
DNest	Diffusive Nested Sampling
DNS	Diffusive Nested Sampling
FEniCS	Finite Element Computational Software
FWHM	Full width half maximum
GPU	Graphics Processing Unit
HDF	Hierarchical Data Format
HPGe	High-purity germanium
I/O	Input/Output
KDE	Kernel Density Estimator
LBL	Lawrence Berkeley National Laboratory
LNGS	Laboratori Nazionali del Gran Sasso
MCMC	Markov Chain Monte Carlo
MJD	MAJORANA DEMONSTRATOR
MSPS	Mega (10^6) samples per second
MVC	Model-View-Controller
NC	Neutral Current
NIB	NeXT Interface Builder
NS	Nested Sampling
ORCA	Object-Oriented Realtime Control and Acquisition
ORTEC	Oak Ridge Technical Enterprise Corporation
PDE	Partial Differential Equation
PMNS	Pontecorvo-Maki-Nakagawa-Sakata

PPC	P-type Point Contact
QCD	Quantum Chromodynamics
QDC	Charge-to-digital converter
QED	Quantum Electrodynamics
SBC	Single Board Computer
SNO	Sudbury Neutrino Observatory
SNR	Signal-to-noise ratio
SPICE	Simulation Program with Integrated Circuit Emphasis
VME	Versa Module Europa
μ TCA	Micro Telecommunications Computing Architecture

LIST OF SYMBOLS

a	Isotopic Abundance
A	Atomic mass number
B	Background Index
E	Electric Field Magnitude
E_{sat}	Saturation Electric Field Magnitude
$G^{0\nu}$	Neutrinoless double-beta decay phase space factor
g_A	Weak axial coupling constant
$M^{0\nu}$	Neutrinoless double-beta decay nuclear matrix element
$\mathcal{L}(\theta)$	Likelihood of parameter θ
m_e	Electron mass
$\langle m_{\beta\beta} \rangle$	Effective neutrino mass in double-beta decay
n	Neutron
N	Neutron number
N_A	Avogadro constant, 6.02×10^{23}
p	Proton
$T_{\frac{1}{2}}^{0\nu}$	Neutrinoless double-beta decay half-life
$U_{\alpha i}$	The PMNS mixing matrix
v_d	Electron or hole drift velocity
W	molar mass
Z	Atomic number
$Z(\theta)$	Model Evidence
α	Fine-structure constant
β	Equal weight enforcement hyperparameter in DNS
ΔE	Energy resolution
ϵ	Detection Efficiency
Γ	Decay rate
λ	Backtracking scale length hyperparameter in DNS
μ	Charge mobility

μ_{sat}	Saturation charge mobility
$\pi(\theta)$	Prior probability of parameter θ
θ	A parameter or set of parameters, as used in a fit

CHAPTER 1

Introduction

1.1 Overview

Neutrinos are neutrally charged fundamental particles of very low mass, which interact only via the weak nuclear force and gravity. Despite being ubiquitous in the universe, these particles have only minimal interactions with matter, and their presence has proven to be an interesting field of study for generations of scientists.

Experimentally, neutrinos are among the most challenging particles to detect, and the technology to detect them has been an intentional and concerted development since they were first theorized. Due to their low interaction rate, specialized, large volume detectors must generally¹ be used to detect the particles. Neutrinos contribute in important ways to many nuclear processes, however, and their presence can be indirectly noted even without these specialized detectors, as in the case of beta decay.

Theoretically, neutrinos have many interesting and unique characteristics that make them worthwhile to study. They are the only massive fundamental particles to have no electric charge, so their exact properties are somewhat uncharted territory; indeed, charged leptons such as the electron are detected and manipulated primarily through their electromagnetic interactions. Their vanishingly small (but nonzero) masses have forced physicists to consider the mechanisms that grant them mass, which are likely to be different from the mechanisms in other particles. Neutrinos also take part in the unusual phenomena of flavor oscillation and matter effects.

In addition to all this, neutrinos are viewed as a possible mechanism for the observed matter-antimatter asymmetry in the universe. During the big bang, it is believed that matter and antimatter should have been created in equal parts; today, however, we live in a matter dominated universe, devoid of any obvious pockets of antimatter. This baryon asymmetry may be explainable through a lepton asymmetry if neutrinos are their own antiparticles, violating lepton number. This possibility will be discussed further in section 1.5.

¹Coherent neutrino-nucleus scattering can significantly enhance the reaction rate.




Figure 1.1: Chart of nuclides, indicating the dominant decay mode by color [1]. The line of stability zig-zags through the middle of the isotopes, near the $Z = N$ line at low mass.

1.2 Beta decay

There are several common types of radioactive decay, including forms which release electrons, photons, or alpha particles, as well as nuclear fission which often produces lighter nuclei. When neutrinos are created, it is always by nuclear reactions involving the weak force, such as beta decay.

Figure 1.1 shows the table of nuclides, which indicates all the combinations of protons and neutrons that can be combined to form nuclei and, in this version, the dominant decay mode for each isotope. The number of neutrons and protons in a nucleus is approximately equal at low mass (A), but at higher masses all nuclei are neutron-rich. If all the isotopes along a given isobar (constant $A = N + Z$) are viewed together by stability (or mass excess), you get an approximately parabolic trough of nuclei with the most stable nuclides at the bottom. This stable bottom corresponds to the black line of stability in figure 1.1, with beta decays funneling inwards towards the line of stability.

When viewed in such an isobaric “cross-section” the bottom corresponds to the line of stability in the table of nuclides, and isotopes with a higher or lower number of protons will decay inwards towards the line of stability, as seen in Figures 1.2 and 1.3. Figure 1.3 shows the isobar for nuclides with atomic mass of 76.

Even numbered isobars, as 76 is, have an additional effect, rather than lying on a simple parabola. A nucleus is more stable with an even number of protons or neutrons, or even more so with both. Nuclei with an even number of neutrons or protons experience a pairing force which increases stability. For this reason, along the even isobar in Figure 1.3 the stability zig-zags along Z as the isotopes vary from being even-even to odd-odd, but this is not seen in odd isobars like Figure 1.2, where there is always one unpaired nucleon.

Beta decay often occurs in neutron- or proton-rich isotopes, (away from the line of stability), moving them towards greater stability. A beta decay may occur in one of two ways, a positive or negative version:

$$p \rightarrow n + e^+ + \nu_e \quad (1.1)$$

$$n \rightarrow p + e^- + \bar{\nu}_e, \quad (1.2)$$

the former equation being a β^+ decay (in proton-rich isotopes), and the latter a β^- decay (in neutron-rich isotopes). This type of decay is particularly interesting because it is capable of transforming one type of nucleon into another, a process which is otherwise not possible in nature – in a nucleus, this is quite literally transmuting an atom of one element into another. In addition to this change, the emission of an undetected (and therefore, effectively experimentally invisible) neutrino was a puzzle for early experimentalists, who, unable to see all three decay products, struggled to reconcile the decay with energy and momentum conservation. In 1930 Wolfgang Pauli correctly proposed that the existence of an unmeasured light neutral particle – the neutrino – could explain the experimental results [3], but it would be decades later before such a particle was directly measured.

1.3 Weak interactions

The weak force is mediated by three types of bosons, two charged, and one neutral. There are nine² weak interaction vertices, and they can be categorized as either *charged current* (CC) or *neutral current* (NC) [4]. A neutral current reaction is one in which the exchanged boson is neutral (the Z^0 particle), and therefore has no exchange of electric charge between the vertices. Similarly, a charged current reaction is one in which the

²I count these as (1) CC with a lepton and neutrino, (2) CC with two quarks, (3) Z^0 with 2 of quark or lepton, (4,5,6) 3 vertices connecting only W's and Z^0 's, and (7,8,9) 3 mixed couplings of the photon with W and/or Z^0 . Understandably, others may choose to account for the total number in a slightly different way, particularly with respect to the mixed couplings. Regardless, it is notably more than QED's single vertex, or QCD's three vertices.

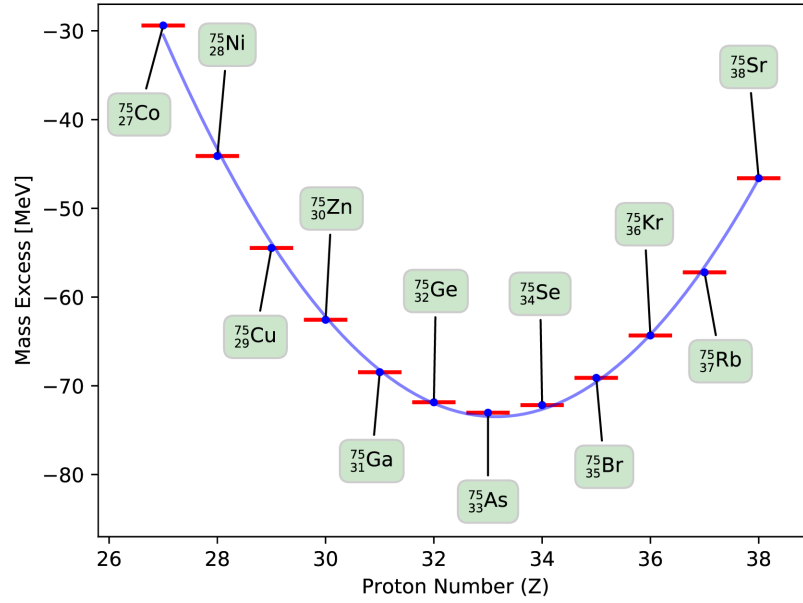


Figure 1.2: Isobar along odd isobar $A=75$. Data from [2].

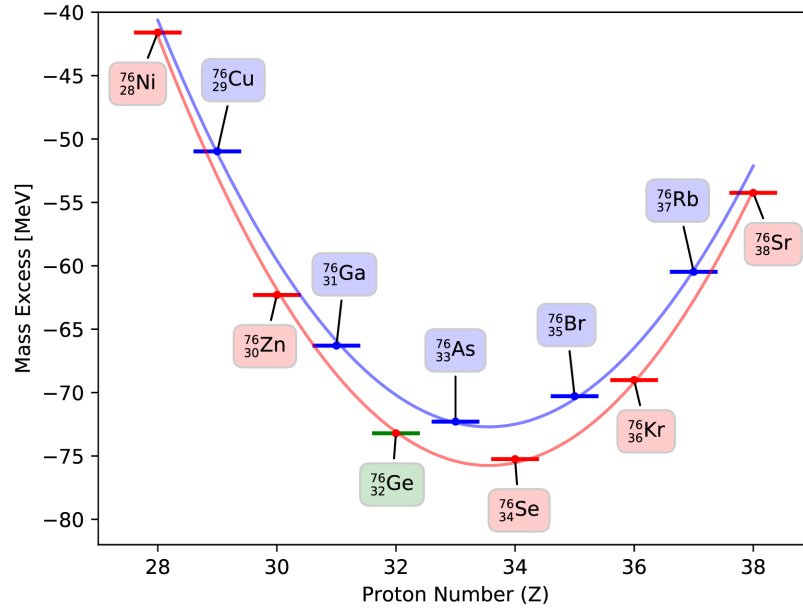


Figure 1.3: Isobar along even isobar $A=76$. ^{76}Ge is unable to single beta decay because both Ga and As would be energetically forbidden endpoints; double-beta decay of ^{76}Ge to ^{76}Se is allowed, but highly suppressed. Data from [2].




Figure 1.4: Feynman diagrams indicating the two forms of beta decay.

exchanged boson is charged (the W^+ or W^-), and does exchange electric charge between the vertices. The difference between these reaction types is easiest to see in the context of some examples.

A beta minus decay, for example, converts a neutron into a proton, an antineutrino, and an electron. In Figure 1.4(a), you can see a neutron decaying into a proton, with emission of an electron and an antineutrino. As the decay starts with an electrically uncharged particle and ends with a charged particle in the same vertex (the neutron decaying to a proton), the connecting boson must carry some charge (the W^-).

Each fundamental particle can be assigned some *lepton number* L , which is $+1$ for lepton particles (electrons, neutrinos, etc), -1 for lepton antiparticles (positrons, antineutrinos, etc), and 0 for particles which are not leptons. Beta decay conserves electric charge as well as total lepton number; both quantities are conserved for all observed decay modes thus far. We will revisit the idea of *lepton number conservation* later in this chapter.

Some interaction modes (such as a neutrino scattering with a charged lepton) can occur as either a NC or CC interaction (and may then have an enhanced total likelihood of occurring).

1.4 Neutrino properties

The neutrino cross section is quite small, making interactions with other matter unlikely. The cross section is primarily small because of kinematics of the decay; the virtual W and Z^0 bosons are very heavy.

After several earlier attempts, the neutrino was first measured by Reines and Cowan in the Savannah River Experiment [5]. Here, three large tanks of liquid scintillator were placed in close proximity to a nuclear reactor (a source of antineutrinos). The experiment looked for evidence of an antineutrino interacting with a

proton in the scintillator, producing a neutron and positron, as indicated by

$$\bar{\nu}_e + p \rightarrow n + e^+. \quad (1.3)$$

The signatures of neutron capture and positron annihilation in near-coincidence provided good evidence that the neutrino had been captured on a proton. Due to its clear signature, this *inverse beta decay* reaction has gone on to become a component of nearly all subsequent neutrino detection measurements.

Following this, the community experienced the Solar Neutrino Problem, first born out of an experiment by John Bahcall and Raymond Davis. Here, Bahcall was responsible for a theoretical model of solar neutrinos, explaining the nuclear reactions in our sun and determining the associated rates and energies of released neutrinos [6]. Davis attended to the challenging experimental details of trying to actually measure solar neutrinos [7]. The experiment came up with only about one third of the expected number: this was the Solar Neutrino Problem. There were many theories as to the cause, varying from experimental mistakes to other theories of the solar model. By 1998, Super-Kamiokande had shown some indications of flavor changing [8], but with some ambiguity to the mechanism³. In addition, the gallium based experiments SAGE and GALLEX/GNO had sensitivity to the low energy proton-proton fusion neutrinos (p-p), and had seen a deficit of solar neutrinos [9]. The solar neutrino problem was finally resolved by the Sudbury Neutrino Observatory (SNO) in 2001, verifying a model theorized by Bruno Pontecorvo in 1967 [10]. This experiment was able to measure all three flavors of neutrinos, and could then see that some of the neutrinos created as electron neutrinos in the sun had *oscillated* into other flavor states by the time of their measurement.

As well as this, we know that the neutrino is a spin-1/2 particle⁴, it weighs much less than an electron but has mass, and it can be produced in beta decays. Cross sections are modestly well known, and we can predict many neutrino-related reaction rates [11].

Despite having decades of study, there are still some fundamental questions about the properties of the neutrino. Perhaps the most important examples are that we do not know the value of the neutrino's mass, nor if it is a Dirac or Majorana particle – a distinction which will be discussed later.

Measurement of the neutrino's mass evades more traditional measurement techniques. The neutrino is the only known fundamental particle having mass but no electric charge. A charged particle can, for example,

³Super-Kamiokande had sensitivity to both the ν_e flavor via CC reactions, and to all flavors via NC/scattering, but could not distinguish between them.

⁴As did Pauli.

be accelerated in an electric field to inform about its mass; this option is not available for the electrically neutral neutrino. Additionally, the neutrino is stable against decay, preventing us from looking at its decay products to infer its energy and momentum.

From neutrino oscillation experiments, we have confidence that the neutrino has mass. Additionally, these experiments can tell us the relative mass difference between the three neutrino mass states. Cosmological measurements from PLANCK currently place lower limits on the sum of all three mass states to $\sum m_\nu < 230 - 540 \text{ meV}$ [12]. While every experiment is somewhat model-dependent, these measurements are arguably more-so, as they rely on a detailed understanding of how the universe evolved and how neutrinos contributed to the formation of structure along the way. Measurements therefore must rely on only the present universe to infer what may have happened during the previous 13 billion years, with (reasonable) models to interpret what we see now.

Despite this knowledge that neutrinos have mass, we do not, however, know the mass ordering, nor the absolute mass scale of neutrinos. That is, we do not know which of the three known mass states is heaviest, and while we know how far apart the mass states are from each other, we don't know how far the lightest neutrino mass is from zero.

As a result, determination of the neutrino mass must come from other sources⁵. There are several proposed (and in progress) experiments to resolve these issues. A direct, "model-independent" way to determine the neutrino mass is by looking at the spectral endpoint in beta decays, which has slight variation for different neutrino mass values [13]; this is the strategy being employed by KATRIN [14] as well as Project 8 [15], although using very different experimental techniques.

1.4.1 Baryogenesis via Leptogenesis

While matter and antimatter are in many ways similar to each other, particles of each have opposite quantum numbers, such as electric and color charges. The antiparticle of the familiar electron is the positron, having identical mass and opposite electromagnetic charge (and lepton number). As far as is known today, the universe consists almost entirely of matter, and not in any substantial way of antimatter [16]. In general

⁵Or possibly a combination of sources.

matter and antimatter are conserved quantities, and in a particle reaction the baryon and lepton numbers remain constant⁶.

It is of interest then to identify mechanism which may be responsible for generating this matter-antimatter asymmetry. One appealing mechanism of baryogenesis is via leptogenesis, where an excess of leptons leads to an excess of baryons, however, this just pushes the problem to a different unexplained excess. The Sakharov conditions are a set of conditions for producing matter from equal amounts of matter and antimatter. These are:

Baryon number violation: Could be via leptons.

C and CP Violation: To distinguish matter and antimatter.

Out-of-equilibrium period: To prevent the asymmetry from reverting immediately back.

One theorized process which results in a net creation of matter is neutrinoless double-beta decay, described in the following section.

1.5 Double-beta decay

An isotope which is forbidden from beta decaying may instead, in some cases, double-beta decay. This decay is usually described as one of

$$2n \rightarrow 2p + 2e^- + 2\bar{\nu}_e, \quad (1.4)$$

$$(A, Z) \rightarrow (A, Z + 2) + 2e^- + 2\bar{\nu}_e, \quad (1.5)$$

and is a single nuclear event consisting of two simultaneous beta decays. In the even A=76 isobar shown in Figure 1.3, ⁷⁶Ge is unable to single beta decay because both Ga and As would be energetically forbidden endpoints; two-neutrino double-beta decay of ⁷⁶Ge to ⁷⁶Se is allowed by the standard model, though highly suppressed. This decay is theoretically possible in 35 isotopes [17], and has been directly measured in 9 isotopes, with half-lives ranging from 10¹⁸ to 10²¹ years [18].

A related process which is interesting for several reasons is a rare decay mode known as *neutrinoless double-beta decay*. This process is theorized to be possible in some models, but has never been observed.

⁶This is not true in kaon decays, however, it is an indirect CP violation, and of insufficient magnitude to explain the observed matter asymmetry.

The decay converts two neutrons within a single nucleus into two protons and two electrons, without the release of any neutrinos, as in

$$2n \rightarrow 2p + 2e^{-}. \quad (1.6)$$

As the protons are bound in the nucleus, the two electrons receive the majority of the energy in the decay, and their summed energy is nearly constant; this means that electrons in this decay would contribute to a monoenergetic peak when measured together.

For this neutrinoless decay to occur, *total lepton number* would be violated, as it would prevent emission of two electron antineutrinos (both have $L = -1$, and their nonexistence in the final state would violate lepton number by $\Delta L = 2$). This would indicate that neutrinos are their own antiparticles, a class of particle referred to as Majorana particles [19], as well as demonstrating that lepton number is not a conserved quantity. This is in contrast to the more standard Dirac particle, which has a distinct particle and antiparticle. Since the neutrino is fundamental and lacks electrical charge, it is conceivable that the particle and antiparticle are the same object. Indeed there are no comparable fundamental particles from which we might infer otherwise.

As neutrally charged fundamental particles, the difference between neutrinos and antineutrinos appears to be just an opposite chirality; this is an intrinsic property, like total spin. In light of the fact that neutrinos have mass, this statement is effectively just noting that neutrinos and antineutrinos have opposite helicity (a state property, like the observable spin projection). As a massive particle is necessarily traveling slower than the speed of light, it is in principle possible for another particle to gain more speed, and pass it; from this boosted reference frame, the neutrino momentum would be going the opposite direction, and the helicity would be flipped. So, if the only difference between a particle and antiparticle state of a neutrino is helicity, then the two states should be (at times) equivalent and interchangeable.

Efforts to measure the postulated process of neutrinoless double-beta decay have placed only limits rather than measuring the process, even after a number of efforts over several decades. The process is highly suppressed, and based on current experiments has an expected half-life of more than 10^{25} years in most isotopes, and greater than 10^{26} years in ^{136}Xe [20].

A decay rate can be estimated using Fermi's Golden Rule, which, for this decay is usually stated as

$$\Gamma = \left(T_{\frac{1}{2}}^{0\nu}\right)^{-1} = G^{0\nu}(Q_{\beta\beta}, Z) |M^{0\nu}|^2 \langle m_{\beta\beta} \rangle^2, \quad (1.7)$$




Figure 1.5: Feynman diagrams indicating double-beta decay

which assumes the decay proceeds via light neutrino exchange. Here, G is a phase space factor, $M^{0\nu}$ term is a *nuclear matrix element* (or “transition amplitude”), and $m_{\beta\beta}$ is the effective neutrino mass.

The phase-space factor $G^{0\nu}$ modifies this transition amplitude with information about the final energy and momentum. This factor is in principle exactly calculable for a given isotope [21]. A decay from a heavy primary into a light final product has many different ways to arrange the energy and momentum (so has a large phase-space factor), but a decay from a light primary into a light final product does not give as many options for how to arrange the energy and momentum, and so is phase-space suppressed. Additionally, increasing the number of particles in the final state decreases the available phase-space (and therefore the magnitude of this factor) due to the increased sharing of energy and momentum.

The nuclear matrix element $M^{0\nu}$ is a statement about the likelihood of the transition from one quantum state into another. It is effectively a measure of how much the wavefunctions for each state overlap. Calculation of this value is extremely nontrivial (especially for heavier nuclei off of closed shells), and estimates of the dimensionless quantity vary between about 3 and 6 for ^{76}Ge , due to differences in techniques and assumptions; this difference is large and the nuclear matrix element enters the rate squared, amplifying the effect. In order to perform the calculation, the initial and final state nuclear wavefunctions must be known or estimated, which is a many-body problem. There are a variety of computational methods used

for calculating nuclear matrix elements, including nuclear shell model, energy-density functional theory (EDF), quasiparticle random phase approximation (QRPA), and the interacting boson model (IBM). In order to reduce some of the uncertainties associated with this value, other related nuclear structure measurements are being made [22]. For a current description of the challenges associated with double-beta decay nuclear matrix elements, see [23] and [24].

While not shown in Equation 1.7, the weak axial coupling constant g_A is often included such that the nuclear matrix element is explicitly $M^{0\nu} = g_A^2 M$. There are phenomenologically interesting reasons to consider why the axial coupling constant may be renormalized inside the nucleus to one of several possible values [25]. This intrinsically produces some uncertainty in the scale of the value, amplified by g_A appearing to the fourth power in Γ . If g_A is excessively quenched, the decay half-life could end up excessively large [26], leaving the reaction even more difficult to measure.

The effective neutrino mass accessible in a double-beta decay, $\langle m_{\beta\beta} \rangle$, is the mass-weighted sum over the Leptonic Mixing Matrix⁷ $U_{\alpha i}$, expressed as

$$\langle m_{\beta\beta} \rangle = \left| \sum_{i=1}^3 U_{ei}^2 m_i \right|. \quad (1.8)$$

The $\langle m_{\beta\beta} \rangle$ term in the rate contains the leptonic part of the full matrix element, leaving the remaining (hadronic) processes to the $M^{0\nu}$ term [27, pg. 173], [28, pg. 164]. As the rate is proportional to the (small but nonzero) $\langle m_{\beta\beta} \rangle$, the rate is then “automatically” encoded with the requirement of nonzero neutrino mass directly. The full value of the term is

$$\langle m_{\beta\beta} \rangle = \left| c_{12}^2 c_{13}^2 m_1 + s_{12}^2 c_{13}^2 m_2 e^{i\alpha} + s_{13} m_3 e^{i\beta} \right| \quad (1.9)$$

where s_{ij} indicates $\sin(\theta_{ij})$ and c_{ij} indicates $\cos(\theta_{ij})$, with α and β the complex Majorana phases. Neutrino mass eigenstates are not simultaneous with neutrino flavor eigenstates. The three neutrino flavor eigenstates correspond to the e , μ , and τ particles, with each corresponding to a mixture of mass eigenstates 1, 2, and 3, as

$$|\nu_\alpha\rangle = \sum_{i=1}^3 U_{\alpha i}^* |\nu_i\rangle \quad (1.10)$$

The angles θ_{ij} are then the mixing angles that describe

⁷Also known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix

The components of this $\langle m_{\beta\beta} \rangle$ term are mostly modestly well known from oscillation experiments, and limits can be placed from cosmological observations as well [29]. There is no experimental knowledge of the Majorana phase terms, and limits on $\langle m_{\beta\beta} \rangle$ are usually placed by allowing the phase angle values to float on the range $[0, 2\pi]$ and using the resultant range of outputs. The m_i terms are not known, but limits can be stated.

In addition to this light-neutrino exchange mechanism, there are other possible ways for the decay to proceed, such as by a heavy particle exchange or a mixture of heavy and light particle exchange. Measurement of neutrinoless double-beta decay is generally agnostic to the mechanism, although some techniques which are sensitive to the angular correlation or energy distribution of the final state electrons may be informative, as in the SuperNEMO experiment [30]. The “standard” light neutrino exchange is favored, as higher order effects may washout, leaving it an ineffective baryogenesis mechanism [31]. Regardless of which mechanism is responsible, if a neutrinoless double-beta decay is observed, the neutrino must have Majorana mass and is then a Majorana fermion [32, 33].

The observed quantity for a neutrinoless double-beta decay experiment is either a decay rate or a limit on the minimum value of this rate. After a measurement, Equation 1.7 can be rearranged to give the effective mass (assuming a particle lepton-violating mechanism). As different isotopes have different matrix elements (and event rates), the uncertainty and model dependence can be improved by combining results from multiple isotopes.

1.6 Sensitivity to double-beta decay

In the design of a neutrinoless double-beta decay experiment, it is important to consider what aspects will improve the sensitivity to measuring double-beta decay. Sensitivity is a measure of the ability to recognize a signal with some certainty in the possible presence of background signals.

The number of counts expected in the region of interest (ROI) for a counting experiment is given by

$$N = \ln(2) \underbrace{\left(\frac{aMN_A}{W} \right)}_{\text{available nuclei}} \left(\frac{\epsilon t}{T_{1/2}^{0\nu}} \right) \quad (1.11)$$

for isotopic abundance a , molar mass W , and elemental mass M , with detection efficiency ϵ , for a half life of $T_{1/2}^{0\nu}$, in an experiment running for time t .

The sensitivity of an experiment is related to the number of counts, and is usually defined in one of two ways, depending on whether backgrounds contribute to the measurement or not, as

$$T_{1/2}^{0\nu} \propto \begin{cases} a\epsilon Mt, & \text{background-free.} \\ a\epsilon \sqrt{\frac{Mt}{B\Delta E}}, & \text{with backgrounds} \end{cases} \quad (1.12)$$

with background index B and energy resolution ΔE [27, pg 197],[28, pg. 171][24, 34, 35]. The distinction between these is somewhat nontrivial, as they are not limiting cases of one another. The transition between cases with and without backgrounds depends on the expected background rate. A low-background experiment is defined by the number of background counts it receives; a background with a lower rate will require more total exposure to accumulate enough counts to worsen the result. See [36, 37] for more details.

Furthermore *Discovery Limit* (DL) or *Discovery Potential* is a quantity which specifies the ability to not just set a limit, but to credibly make a measurement. After all, we would like to eventually make measurements of physics, not just continue to set limits. A DL has two basic required components, a level of confidence you require ($N\sigma$), and a nonzero number of counts that you could credibly be expected to measure [37]. The DL will be a lower limit than the sensitivity, and is a more realistic picture of the usefulness of a particular experiment.

When faced with the possible isotopes which may undergo double-beta decay, there are several metrics to use to evaluate which is most ideal to actually measure. There are basic physics metrics, such as the phase space factor and the $Q_{\beta\beta}$ endpoint energy value, which can give an estimate of the expected count rate for a given isotope and spectral feasibility. Also important is the expected nuclear matrix element magnitude [38]. Perhaps the most compelling arguments are the experimental realities of trying to build scalable, efficient detectors using or containing the isotope of interest.

The isotope chosen by the MAJORANA collaboration is ^{76}Ge . As will be discussed in the following chapters, germanium is also an ideal material for the construction of high resolution detectors; germanium detectors are a standard nuclear physics spectroscopy tool for this reason. For a recent overview of other major experimental searches, see [39].

CHAPTER 2

The MAJORANA DEMONSTRATOR

The MAJORANA DEMONSTRATOR is a germanium detector array experiment with the goal of evaluating methods to be used for a large-scale neutrinoless double-beta decay experiment. The approach is intended to be modular and scalable. It consists of 58 detectors together weighing 44 kg distributed between two different cryostat modules, referred to as Module 1 and Module 2.

There are many constraints when designing detector systems for a double-beta decay experiment. The guiding principle is that radioactive backgrounds must be reduced as much as possible while increasing the exposure and improving the energy resolution as well. Backgrounds are kept low through a careful and deliberate program of radioassaying all components to be used near the detectors [40], and simulating the potential impact each would have on the final result [41]. The MAJORANA collaboration underwent a significant and successful effort to electroform and machine high-purity copper underground to reduce backgrounds and cosmogenic activation. Copper parts are used extensively to support the detectors and build all close structural components, as copper is machinable and strong, yet has no natural radioactive isotopes, and can be made very radiopure.

To shield from radioactive background sources near the detector array, the experiment uses a dense, compact shield consisting of 54 tonnes of lead bricks surrounding inner layers of clean copper; the *inner copper shield* was electroformed and machined underground [40]. The lead shielding is surrounded by borated polyethylene plastic sheets, which can shield against neutrons. The inner volume is then flushed with liquid nitrogen boiloff gas, which contains lower levels of radon than air.

The background constraint limits the number and type of electrical cables which can be used near the detectors, and requires that most of the preamplifier circuit is outside the vacuum system. Signal and high voltage cables are chosen to have low mass, an unusual constraint¹. It is also desirable to have the most active

¹Though specialized industries such as aerospace have some overlapping interest.




Figure 2.1: A schematic of the MAJORANA DEMONSTRATOR, indicating the primary experimental components

detector mass with the smallest number of readout channels, as additional channels encumber additional cables; this leads to development of large mass, single electrode detectors.

Many materials which are springy yet conductive contain low levels of radioactivity. As most electrical connectors rely in some way on the elasticity of pins and sockets to force contact when meshing together, it is difficult to form reliable, low-background connectors, especially over a broad temperature range. This process was a major design effort for the collaboration, perhaps surprising to those outside the low-background community.

An underground location is ideal as a way of shielding the experiment from cosmic rays (and cosmic ray secondaries, such as muons). For this reason, the DEMONSTRATOR is located 4850 feet underground at the Sanford Underground Research Facility, in the former Homestake Gold Mine, in Lead, SD. This is the site of the well-known “Homestake Experiment” performed in the 1960s by Raymond Davis to measure solar neutrinos [7]. The surface muon flux is $2.0 \pm 0.2 \mu/s/cm^2$, and a measured flux at the site of the DEMONSTRATOR of $(5.31 \pm 0.17) \times 10^{-9} \mu/s/cm^2$ [42], corresponding to a reduction by nearly 9 orders of magnitude.

The detectors used in the DEMONSTRATOR are made of germanium crystals (which are further discussed in section 3.2). Our detectors are able to detect with high resolution the amount of energy deposited in them by ionizing radiation.

To make the desired measurement, we have fabricated the detectors using germanium enriched in the isotope ^{76}Ge , and are looking to observe the decay of the detectors themselves, with no external sources present. This presents the advantage of capturing nearly all energy from events occurring within the detector, and maximizing the detection geometry (nearly 4π -coverage). It does have the curious feature of lacking the ability to take “true” background data *without* the double-beta decaying source present, however, the expected rate is so small, this is in practice never a problem for the DEMONSTRATOR. An experiment with more modest resolution would struggle to differentiate between the endpoint $2\nu\beta\beta$ electrons and any possible $0\nu\beta\beta$ electrons.

2.0.1 Status of the Demonstrator

The MAJORANA DEMONSTRATOR began commissioning on the first module in June 2015, followed by production data taking in January 2016; the experiment is currently taking data on both modules. In 2018, the collaboration released its initial results from the full array, with blind data [43]. Here, we demonstrated an energy resolution of 2.5 keV FWHM at the 2039 keV region of interest, and placed a lower limit on the half life of 1.9×10^{25} years. More recently, we released another result, improving this limit to 2.7×10^{25} years [44].

The DEMONSTRATOR was conceived of as an experiment to validate the feasibility of building a larger-scale detector; the techniques being proposed had never been tried before and it is valuable to develop experience with an intermediate-scale project before committing to a large-scale project. One issue that was identified in the construction and operation of the experiment was that the cables and connectors to be used can provide a number of reliability problems which reduced the yield of operational detectors.

2.0.1.1 Future of MAJORANA

In order to make a competitive measurement of double-beta decay, it is clear that a larger experiment is necessary, and preferably one with further refined implementation. Although the MAJORANA and GERDA collaborations are in principle competing to measure double-beta decay, the two groups have common interests and experiences, and have worked closely together since their formation, with the goal of eventually joining together into a single collaboration.

In 2017, the LEGEND collaboration was formed from the GERDA and MAJORANA collaborations with the goal of reaching a double-beta decay discovery sensitivity of 10^{28} years [45]. This experiment is a scaled

approach, which will demonstrate a medium-scale experiment of around 200 kg, followed by a full-scale experiment at tonne scale. The design will be a hybrid of the technologies and strengths of each collaboration.

While detectors are currently being fabricated for this project, using the existing detector mass of the DEMONSTRATOR and GERDA projects can bootstrap the construction process and get the project started much more quickly. As a result, the DEMONSTRATOR will stop taking data and transition detectors into the LEGEND-200 cryostat located at Laboratori Nazionali del Gran Sasso (LNGS), in Assergi, Italy. As there are naturally competing interests between prolonging the DEMONSTRATOR and constructing the next generation LEGEND-200, the detailed timeline of this process is currently in flux.

2.1 Outline

The signals produced by the germanium detectors of the DEMONSTRATOR are described in the following chapter. Using numerical simulations, we can simulate the expected fields and signals, given their dependence on the detector geometry and characteristics.

Using a bayesian machine learning algorithm, we can use these simulations to estimate the actual parameters of our detectors, as well as the individual waveform parameters (such as deposited energy and hit position). In a training stage, we analyze a selection of non-simulated (data) waveforms from a single detector to characterize the detector properties. Using a continuous feedback by simulating new fields and signals, the waveforms are all simultaneously fit to give a best estimate of the common detector parameters. After finding these values, they can be “frozen”, and individual waveforms can use this information to do a similar fit, using the known fields and *siggen*. The details of this modeling, simulation, and fitting will occupy the following chapters.

CHAPTER 3

Signal Formation

Radiation detectors are devices which can convert ionizing radiation into a signal that is readily measurable; there are many types of such detectors, each with various specific benefits and limitations. In scintillators, radiation is converted to optical light, which can be measured with a light sensor. A bolometer converts the deposited energy into thermal energy, which is measured with a temperature sensor such as a thermistor. In a solid-state semiconductor detector, the radiation is converted directly into an electrical current signal, measurable by the appropriate current-sensing instrumentation.

A germanium detector is a semiconducting single crystal often used to detect radiation. Ionizing radiation incident on a crystal of germanium liberates electron-hole pairs, creating a small but measurable current signal. The details of this process and the associated readout electronics necessary are subjects of this chapter. Furthermore, we will discuss the techniques of modeling and simulating the signals from detectors.

3.1 Semiconductors

In 1931, Wolfgang Pauli famously proclaimed “One shouldn’t work on semiconductors, that is a filthy mess; who knows whether any semiconductors exist.” [46] With the benefit of nearly a century of additional study, we can now see that semiconductors have been one of the most fruitful avenues of scientific investigation, and are responsible in some way for nearly all modern technology. They are also responsible for a valuable technique in radiation detection, the semiconductor or solid-state detector.

A single atom has electronic structure, giving rise to specific allowed energy levels corresponding to each electron’s quantized energy. When atoms are bound together at short interatomic distances in a crystal lattice, the long range order of the lattice changes these levels to be bands, energy ranges in which the electrons may occupy one of a more broad selection of energies, as shown in Figure 3.1. In reality, the bands are composed of $O(N_A)$ discrete levels, but the close spacing of levels with finite widths makes the collection of levels into continuum bands.




Figure 3.1: Valence and conduction bands in a semiconductor, source [47]. The interatomic distance, or *lattice constant* a is fixed for a given material, and the value it takes on contributes to the exact band structure.

Of primary interest is the behavior of the electrons in the *conduction* and the *valence* bands. The conduction band electrons are the highest energy electrons in the material, and are loosely bound, not limited to any particular lattice site. The valence band electrons are slightly lower in energy, and are bound to particular lattice sites. The valence band is the highest energy filled band, and is completely filled in a semiconducting material.

Germanium, like silicon and carbon, has 4 valence electrons. These 4 electrons are covalently shared with adjacent atoms, most commonly forming a diamond cubic lattice, the same structures as formed by silicon and carbon-based diamond.

In an electrical conductor such as a metal, the valence and conduction bands overlap in energy (there is no gap), and the valence electrons can very easily be moved into the conduction band; by default, the conduction “band” will already contain some electrons. In an insulator, there is a large gap between these bands, and valence electrons require considerable energy to be elevated into the conduction band. For a semiconductor, the gap is of a “moderate” size, and electrons may be pushed into this higher energy band relatively easily. It is this mechanism that we will eventually describe as being useful for making radiation detectors.

Electrons and holes feel some effective mass¹ while in the lattice, which will depend nontrivially on their momenta and energy. An electron band structure diagram, as in Figure 3.2, shows the relationship between the charge carriers’ energy and momenta (wavevectors), usually in a 2-dimensional representation of the

¹Their effective mass being the constant relating their acceleration to a force ($m = F/a$), as from an electric field ($\vec{F} = q\vec{E}$).




Figure 3.2: Germanium band structure diagram, adapted from [48].

4-dimensional (E, k_x, k_y, k_z) space. This figure is also referred to as a dispersion relation. Such a diagram contains a significant amount of information about a material's electronic properties such as the density of states and band gap energy. The band gap E_g of germanium occurs at the L point, along the $\langle 111 \rangle$ direction. As there are 8 equivalent $\langle 111 \rangle$ directions, there are 8 conduction band minima. Other conduction band local minima include the Γ_1 and Γ_2 minima, though they are less important for this study. Other properties can visually be deduced, such as the momentum necessary for the indirect transition between the band gap minimum and the zero-momentum gamma point. Furthermore, the effective mass is inversely proportional to $\frac{d^2 E}{dk^2}$, the second derivative of the dispersion curve at a given point. As the electron and hole will be in different bands, the two will in general have different dynamics. This is interesting, because, without any knowledge of the electron or hole wavefunctions, we can still deduce the transport properties such as instantaneous velocity and acceleration of an electron or hole in an external field [49]. This will be valuable when calculating the signals from electrons and holes inside a detector.

3.2 Germanium Radiation Detectors

Germanium radiation detectors are a class of semiconductor detectors, all which operate on a similar principle. They are ubiquitous, well-established tools within nuclear radiation spectroscopy, and have been in use for spectroscopy in various forms since the 1960s. They have several advantages over other technologies, most prominently including their excellent energy resolution and timing characteristics [50].

The earliest solid-state detectors were simple, thin planar detectors and were primarily used for charged particle detection. The primary use for germanium detectors in the 1960s and onwards was in gamma ray spectroscopy. These detectors were mostly coaxial and semi-coaxial geometry detectors. In the 1980s, it was realized that semiconductor patterning techniques could be used to make highly segmented detectors, with many channels to be read out in parallel on the same single crystal. These segmented detectors were usually silicon and had an increased complexity in operation, but with a high degree of spatial resolution, allowing for different types of analysis of high value for some experiments. Additionally, by segmenting the detector, the total signal rate in each segment is lower, which may be practically necessary for high rate applications [51, pg. 318] Though their spatial resolution may help eliminate some background signals, their need for more readout channels would introduce too much background near the detectors, so they are not currently used for neutrinoless double-beta decay experiments. Additionally, segmented tracking detectors are often designed to be lower mass reduce scattering, counter to the exposure needs of rare event searches.

3.2.0.1 Detector Processing/Fabrication

The basic germanium detector is a single monocrystal of germanium, often shaped in a right circular cylinder. The crystals are of very high purity, and are usually grown by the Czochralski method. In this method, a rotating seed crystal is slowly pulled from a melt to form a boule, as indicated in Figure 3.3. Usually detectors are pulled with the $\langle 100 \rangle$ crystal axis along the pull-axis [52, pg. 19], though $\langle 111 \rangle$ may also be used in some applications². In general, cylindrical detectors have an unknown azimuthal axis orientation in operation, as most spectroscopy applications are not able to see any effects from this.

In order to purify the material that is used for Czochralski growth, several important techniques are used. After isotopic enrichment, the detectors are in a $^{76}\text{GeO}_2$ form, which is then chemically reduced to form the metal. This material is zone refined until it reaches a resistivity of at least 47 ohm-cm, equivalent to

²The $\langle 100 \rangle$ direction is preferred for detectors because it grows with fewer charge trapping dislocations.

10^{13} electrically active impurities per cubic centimeter [53]. With this material, the detector manufacturer further zone refines the material to at most 10^{11} impurities per cubic centimeter, then it is grown into a single crystal by the Czochralski method which further eliminates impurities in the material; further discussion of the consequence of these impurities will be found in the next section. Often, the final zone-refining step must be performed after a first crystal growth, as this prevents impurities from collecting at the grain boundaries inside the material [54]. After crystal growth, the boule is sectioned and adequate parts may be used for detector fabrication, while others may be re-grown to further purify as needed. This crystal section will then be machined into a detector geometry and patterned with contacts.

3.2.0.2 Impurities

The crystal may be doped with an “impurity” of some sort in order to give it the appropriate semiconducting properties. While the term impurity may colloquially imply that it is undesirable, it should instead be interpreted to mean that it is a difference from being a perfect lattice of germanium atoms; the impurity is one of the most important properties in the manufacture of any semiconductor. These impurities may come from simple lattice defects, such as a dislocation or vacancy, or they may come from an atom of an element with a larger or smaller number of valence electrons than germanium taking a germanium atom’s place in the lattice.




Figure 3.3: a) A single crystal being pulled from the melt in a czochralski-style pulling apparatus. b) The resulting boule following a czochralski pull.

Regardless of how it is achieved, the result is a relative lack or excess of electrons in the area localized near the impurity. The intentional addition of a material into a semiconductor in this way is known as doping.

A doped semiconductor has an unequal concentration of free electrons and holes. A p-type semiconductor has a majority of its free charge carriers as positive holes, whereas an n-type semiconductor has a majority of its free charge carriers as electrons. Most commonly, group IV semiconductors like germanium are doped with a group III element (usually boron) to make them more p-type or a group V element (such as phosphorus or arsenic) to make them more n-type. High-purity germanium (HPGe) has very low concentrations of impurities, with typical values being less than 10^{10} electrically active impurities per cubic centimeter, equivalent to parts per trillion concentration (see Appendix A).

Normally the impurity levels at the tail and seed ends of the boule will not be the same, for a given impurity type. This is because the solubility of impurities is different in the liquid than solid phase; the ratio of the concentrations in the solid and liquid phases is known as the segregation coefficient for that impurity type. Aluminum (a p-type impurity) has a segregation coefficient near 1, causing it to be nearly uniformly distributed in the crystal. Several n-type impurities such as phosphorous and oxygen have segregation coefficients less than 1, and will then be preferentially found in the liquid phase, corresponding to a higher concentration in the tail end of the crystal (as they become concentrated while the crystal is pulled from the melt). For this reason, HPGe crystals are usually p-type at the seed end and n-type at the tail [52]. While in general this gives a predictable gradient of impurity concentration, material may also be introduced into the melt while being pulled to change this behavior, as well as diffused into the material after growth. Usually the details of what is done here are considered proprietary trade secrets. The impurity concentration gradient is of great importance in germanium detector operation, and will be discussed further in later sections.

As only a very small impurity concentration determines the majority carrier type, it is possible when growing a high-purity germanium detector to inadvertently produce the wrong type crystal. It is only by detailed process control that high-purity detector production is made possible.

3.2.0.3 Contacts

Additionally, the detector must have contacts at the anode and cathode. This allows for the detector to have a uniform potential over the relevant surfaces, as well as giving a location for readout. The presence of these contacts also enable the detector to be biased without conducting excessive current.

Contacts come in two basic types, *ohmic* (or, *non-rectifying*) and *blocking* (also called *rectifying* or *non-injecting*). Normally, a detector has (at least) one of each type.

The rectifying contact in a p-type detector is the n^+ contact, which is dissimilar to the bulk p material [50]. These contacts form a P-N junction, which is effectively low resistance under forward bias and high resistance under reverse bias. A positive voltage is applied to the n^+ contact, in the reverse bias configuration. In this way, a voltage can be applied to the detector (and in turn an electric field), which depletes the bulk to form the active detector volume.

The non-rectifying contact in a p-type detector is the p^+ contact, which is held at 0 V, and is where the signal is read out from. This contact is effectively high resistance, and reduces the leakage current through the detector. As both the bulk and the contact are p-type, this is not a semiconductor junction, but instead current is blocked by the minority carrier motion[50]; that is, there simply are not enough minority carriers (free electrons, in this case) in the material to propagate negative charge through the junction. The current under bias of several thousand volts would then be unacceptably high without a blocking contact.

The area between the contacts must not be conductive, so as to avoid forming surface currents across the detector which could mask the intended signal current through the detector. In silicon, the surface naturally forms a robust oxide which is nonconductive. While germanium also readily forms a surface oxide film when exposed to oxygen or water, the naturally formed germanium oxide is chemically inhomogeneous, porous, and low-density, which can lead to charge accumulation along the surface of the passivation layer [55]. Therefore, germanium must be carefully treated to form such a passivation layer, and there are many (often proprietary) techniques for accomplishing this.

3.2.0.4 Operation

Germanium has a band gap of approximately 0.6 eV and an electron-hole creation energy of around 3 eV. The difference between these values can primarily be understood as some energy being transferred to the crystal lattice as phonons [56]. The germanium bandgap is relatively small among semiconductors, at roughly half the value of silicon. The result of this small bandgap is that thermal energy at room temperature is sufficient to excite electrons across the bandgap, creating a high leakage current. In order to circumvent this, the detectors are operated at low temperature, usually using liquid nitrogen or mechanical cooling methods to keep detectors around 80 K. Subsequently, the detectors must also be kept under vacuum, otherwise water vapor will condense onto surfaces, potentially increasing surface conductivity or introducing backgrounds.

An alternative that has been investigated in the past is operating detectors directly immersed in a liquid cryogen such as liquid argon; this is the technique used by the GERDA collaboration.

3.2.1 Signal Formation in a Detector

Ionizing radiation incident on a crystal of germanium will liberate electrons and holes in pairs. Having opposite electric charge, these electron-hole pairs would normally just diffuse randomly in the crystal, trapping on crystal defects or mutually attracting and recombining immediately. Instead, if a voltage is applied across the detector, the electrons and holes (with opposite charges) will experience forces in opposite directions, preventing recombination. These charges will drift in the applied electric field, inducing a charge in a readout electrode, where the signal is read out.

It is possible to have several readout electrodes on a single detector. Additional electrodes provide additional information about the position of the charges as they move through the detector volume.

The raw detector signal that is formed has a detailed dependence on the electric fields present, and therefore also the detector geometry and impurity profile.

A useful quantity when considering signal formation by charges moving through a detector is the *weighting potential*. The weighting potential is the solution to Laplace's equation for a given detector geometry, with unit potential on the electrode of interest and zero potential on all other electrodes [57]. In a multi-electrode detector, you would have several different weighting potentials corresponding to the signal observed for each electrode; in this scenario there is increased utility in being able to calculate the expected signal for each contact. In a single electrode configuration, where the readout electrode is held at ground, the result is effectively a normalized electric potential.

3.2.2 Detectors for the MAJORANA DEMONSTRATOR

The germanium detectors used by the MAJORANA DEMONSTRATOR are of a type known as P-type Point Contact (PPC) detectors. The point-contact germanium detectors were first developed by Luke et. al. [58] as a way of fabricating large volume detectors capable of low energy threshold operation, for rare event searches such as dark matter³. These detectors have a geometry which is approximately indicated in the cross-sectional line drawing in Figure 3.4, and a photograph in Figure 3.5.

³The detectors they made were n-type, but otherwise very similar




Figure 3.4: The approximate PPC detector cross section as used by MJD. The HV conducting n^+ outer surface is indicated in pink, a nonconductive passivated surface in orange, and the small p^+ point-contact readout electrode in blue. The detector is axially symmetric about the dashed line.

The detectors used by the MAJORANA collaboration consist of a p-type semiconductor in the bulk of the crystal, with a small point-like readout electrode. This point-contact geometry is desirable for three primary reasons:

- Lower capacitance
- Long drift times
- Large volume.

These optimizations are somewhat unique to double-beta decay.

The relatively low capacitance of the detector is clear if you picture it as a familiar parallel plate capacitor, with one plate wrapped around the outside surface, and another shrunk to the size of the point contact. Having a low detector capacitance is ideal to reduce the series noise, which therefore allows for higher signal-to-noise ratio (SNR) and lower detector energy thresholds to be achieved [51, pg. 124]. In addition to these requirements, there is a more basic constraint for double-beta decay experiments to have detectors at the largest reasonable masses possible. Doing so allows for the highest exposure, with the fewest number of cables and parts near to the detectors.

Due to the specific residual impurities that are present after fabricating high-purity germanium, holes are less sensitive to charge-trapping. Therefore, in order to maximize the detector performance (energy resolution) p-type material is used for fabricating our detectors. N-type material, on the other hand, is preferred for many high-rate or space-based applications, due to its relatively higher radiation hardness; this is because the point defects caused by most particles will trap holes. For a double-beta decay experiment with very low rates and backgrounds, a p-type type detector is then a clear choice.



Figure 3.5: Photograph showing the “bottom” surface of an ORTEC enriched PPC detector. Note the bevel and the small point contact.

3.2.3 Characterization of MAJORANA PPC Detectors

After the MAJORANA detectors were fabricated it was important to verify the characteristics and operation of them. This characterization covered all aspects of the detector, including detailed geometric measurements, and impurity measurements. Most will not be mentioned here, but several of these measurements are of particular interest to this study.

The manufacturer performed Hall effect measurements via the van der Pauw method to estimate the impurity concentration at each end of the detector crystal. In order to perform the measurement, the manufacturer cut wafers from the boule above and below the detector segment, and measured the Hall mobility in each [59]. These measurements are considered to be low-accuracy, but they are the only experimental estimate we have of the impurity gradient in the detectors.

Additionally, we used a Starrett Manual Visual System optical measuring device to precisely measure the dimensions of each detector, including the point contact and bevel. These dimensions are crucial for calculating the fields in each detector.

3.3 PPC Germanium Detector Operation

P-type germanium may be doped with an acceptor such as boron to increase the carrier concentration of holes in the bulk material (making holes the dominant carrier). A P-N junction acts as a diode, and it is common to dope an outside layer of opposite type from the bulk to make a P-N junction on the surface; this acts as a *rectifying contact*, in which the direction of current flow is limited [50]. In order to make a detector which can be readily handled without damaging the surface, manufacturing limitations require the outer surface to be a thick, lithiated n^+ layer. This layer is doped to such a high degree that it is nearly as conductive as a metal, allowing the surface to form an equipotential High Voltage (HV) contact. The thick surface layer also acts to form a “dead layer” of low electric field in which liberated charge contributes minimally to the measured signal (or not at all), as the charges mostly recombine before they can diffuse out of the low-field region. This is convenient as it prevents many surface alpha events from contributing as potential background events. In order to make the surface have an n^+ rectifying contact on the outside surface, the detector bulk must then be p-type. The (small) point contact⁴ surface is of p^+ type [50].

⁴Or, *inner surface* for a p-type coaxial detector.

A voltage is applied across the crystal, which causes free charge in the lattice to migrate out of the bulk, forming a region in the bulk of the crystal depleted of free charges ⁵. The voltage is increased until the full bulk of the crystal is entirely depleted. This depleted region is an active region, which is capable of detecting ionizing radiation. Ionizing radiation interacting in this depleted region will liberate electrons and holes from the crystal lattice, which, under the influence of the established electric field, will drift to either the HV contact or the ground.

Due to its method of signal formation, a germanium detector can be thought of as a solid-state drift chamber⁶. When ionizing radiation is incident onto a detector, it liberates electron-hole pairs, which then drift to their respective electrodes. As the charge drifts through the detector, it induces a charge at the point contact according to the Shockley-Ramo theorem,

$$i = q\vec{v} \cdot \vec{E}_0(x) \quad (3.1)$$

which states that the instantaneous current is given by the product of the charge, the instantaneous velocity (which may depend on the electric field) and the value of the weighting field at the position of the charge. Equivalently, this theorem may be stated in terms of a weighting potential to give the integrated charge, as in

$$\Delta Q = q(\Phi(x_1) - \Phi(x_0)) \quad (3.2)$$

$$Q(t) = q\Delta\Phi(x(t)). \quad (3.3)$$

One way to conceptualize this is that the value of the weighting potential at a point is the fraction of the total induced charge resulting from the minority carrier motion from its starting point until full collection. Therefore, in the low weighting potential bulk of these detectors (as visible in Figure 3.6), the electrons will contribute only minimally to the total signal.

Inside of our PPC detectors, the field is designed such that the weighting potential is very low in most of the detector, and sharply increases in magnitude near the point contact [58, 60], as is visible in Figure 3.6. As the n^+ contact is at high voltage and the point contact is at 0V, the holes move towards the point contact and

⁵While not done in practice, one way to think about this is that the total integrated current while biasing is the amount of charge stripped out of the depletion region (with the leakage current subtracted off). If you knew either the active volume or the impurity concentration, you should then be able to relate them in this way.

⁶With the advantages of better timing and energy resolution, in a much more compact and higher density detector

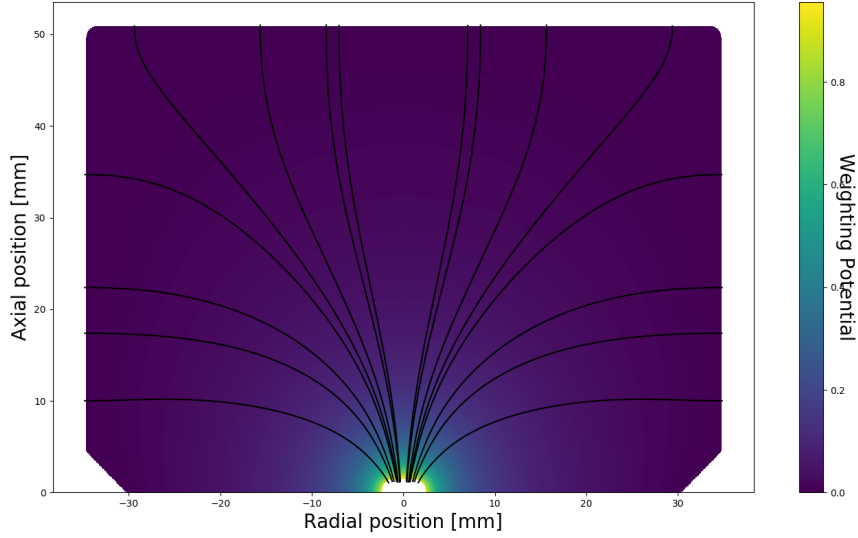


Figure 3.6: Simulated weighting potential, with lines indicating the drift trajectories.

the electrons move away from it. From the Shockley-Ramo theorem (Equation 3.2), we can see that the holes which move through the high weighting potential region will contribute the majority of the current read out, and the electrons will only make large contributions if they start near the point contact and move out of the high weighting potential region.

This weighting potential is well suited for performing pulse-shape analysis of signals [58, 59, 61, 62]. The long drift times provide temporal resolution to allow discrimination between events which deposit energy at a single location within the detector from those which deposit energy in more than one location. Beta-decay electrons, such as those from a double-beta decay event, should travel less than a millimeter within the crystal, effectively making them single site events [61]. Background events from sources such as gamma rays that interact more than once in a detector have waveforms which are categorically different from single site events. This *single-site* and *multi-site* event discrimination is a key way to identify and remove background events. Additionally, the small point contact electrode provides a low-capacitance, which reduces the noise and therefore also improves the energy threshold.

3.3.0.1 Drift Velocity

In low magnitude electric fields, the charge mobility μ of a material is the drift velocity per electric field $\mu = v_d/E$ in a linear relationship. The velocity then is

$$v_h = \mu_h E \quad (3.4)$$

$$v_e = \mu_e E \quad (3.5)$$

for holes and electrons, respectively. This relationship breaks down at higher field values, as the relationship is not actually linear [50, 63, 64]. At high field values, the drift velocity will saturate to a maximum value which no longer increases with additional field. This is parameterized empirically [63] as

$$v_d(E) = \mu_{sat} \frac{E}{[1 + (E/E_{sat})^\beta]^{1/\beta}}. \quad (3.6)$$

This parameterization was developed for silicon, but is widely used for germanium. In addition, the original paper only verified the model over the temperature range from 300 K to 450 K, however, it is applicable even at low temperature [65]. Indeed, materials may have more complicated mobility, as in the negative differential mobility of CdTe or GaAs, where increasing the electric field eventually decreases the drift velocity; this effect is also seen in Si or Ge [66] at low temperatures, around 30 or 40 K. The empirical model of Equation 3.6 does not contain any way to account for these unusual effects, as they do not contribute at realistic temperatures⁷.

In addition to this field-dependent mobility, germanium also exhibits anisotropy, with axis-dependent effects. This accounts for a several-percent effect, depending on the location within the detector. Shown in Figure 3.7 is a map of the drift time as seen in a constant-z cross section of the detector, showing the small but notable effect with changing azimuthal angle. Additionally, Figure 3.8 shows the same map for a single radius value. In order to account for this, the mobility has been measured in multiple axis directions, and can be rotationally transformed with respect to the local electric field direction to given a local mobility value [64].

⁷Experiments such as EDELWEISS or CDMS using germanium detectors at much lower temperatures may be able to see effects from this.

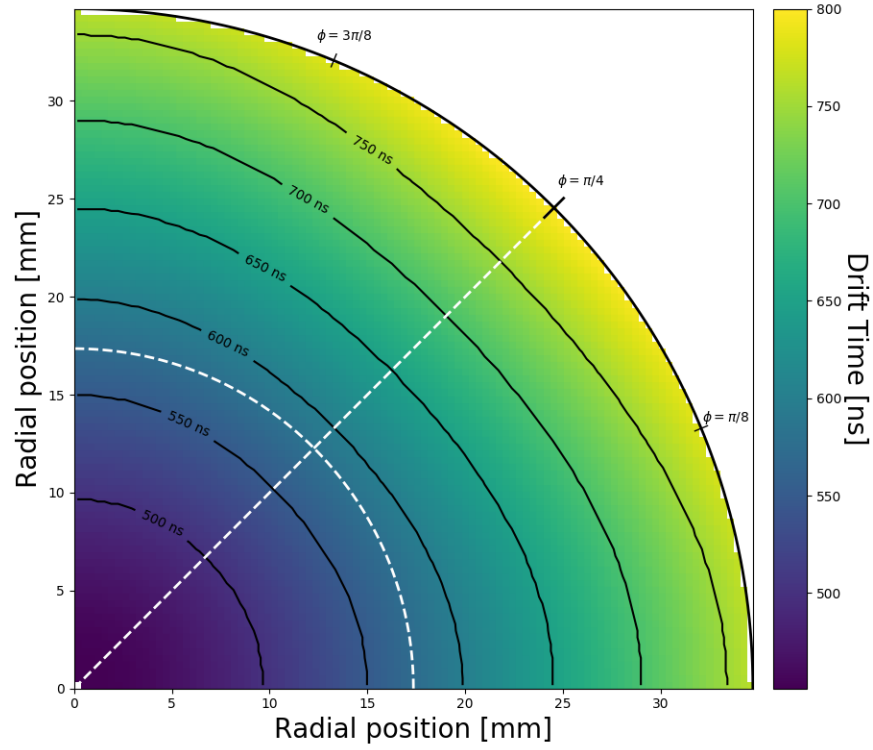


Figure 3.7: Simulated hole drift time map across a detector cross section. Shown for detector P42575A, at a constant z position of half the detector length. Note that the drift time contours are not azimuthally symmetric, with longer drift times at $\phi = \pi/4$ for a given radius. The map is symmetric about $\phi = \pi/4$. The hole drift time along the white constant-radius curve is shown in Figure 3.8.




Figure 3.8: Simulated hole drift time along axis angle ϕ . Shown for detector P42575A, at a constant z position of half the detector length and a constant radius of half the detector radius. This corresponds to the drift time along the white dashed constant-radius curve in the middle of Figure 3.7.

carrier	axis	μ_0	β	E_0	μ_n
hole	100	✓	✓	✓	✗
	111	✓	✓	✓	✗
electron	100	✓	✓	✓	✓
	111	✓	✓	✓	✓

Table 3.1: Velocity model parameters included for each charge type.

As mentioned in Section 3.1, in general, electrons and holes will have different mobilities in a material. In each of the axis orientations indicated above, the mobility has been measured for both electrons and holes [67]. The experimental technique used is a time-of-flight measurement wherein a planar sample is exposed to a fast pulse of ionization radiation, and the drift time of charge carriers across the device is measured [68].

The hole mobility is distinct in each of the three principal crystal axis directions $\langle 100 \rangle$, $\langle 111 \rangle$, and $\langle 110 \rangle$, with the fastest drift velocity in the $\langle 100 \rangle$ direction, and the slowest in the $\langle 111 \rangle$ direction. Due to the diamond cubic lattice symmetry, the mobility for any position can be stated in terms of only the $\langle 100 \rangle$ and $\langle 111 \rangle$ axis mobilities [64, 69]. Using this model, we can then parameterize the complete hole mobility using each of the two single-axis mobility values, according to Equation 3.6.

In the simulation models used in this study, the velocity alone is parameterized using 14 terms, as indicated in Table 3.1. Despite this large number of parameters, the fit only allows the 6 hole parameters to float, and uses the experimental values for the 8 electron parameters as found by [67]. This is justified because the holes contribute the dominant component of the signal for most waveforms. In principle, the hole values could also be held constant to their known values while fitting, however, the use of isotopically-enriched high-purity material leaves the exact values unknown. In general, isotopic enrichment of semiconductors can change the material properties significantly [70, 71, 72, 73]. To our knowledge, there have been no published charge carrier mobility measurements specifically applicable to isotopically enriched ^{76}Ge at liquid nitrogen temperatures; such a measurement would be valuable to validate and refine models of germanium detectors.

3.3.0.2 Charge Cloud

As charges drift through the volume of the detector, they have some charge cloud extending beyond just a point charge. The size of the charge cloud will in general be different for different interaction types (e.g. x-ray photoelectric absorption vs. minimum ionizing muon track). There are diffusion effects which can spread out the charge cloud in space both at the interaction point and during the drifting of charge [50,




Figure 3.9: Time domain waveforms from charges in a single cloud of gaussian with FWHM of 1 mm from the bulk of an enriched detector. This charge cloud is exaggerated for visualization. The gaussian cloud compares well to a smoothed waveform with a gaussian kernel standard deviation of approximately 10 samples.

pg. 369]. This corresponds to some spread in time for the arriving signal, and some small variation in the waveform shape for waveforms along the same isochrone.

In order to simulate the effect of a charge cloud it would then be most appropriate to calculate a large number of charges in a charge cloud, and step them all through the detector volume. Calculating a large number n of points in a true charge cloud is computationally expensive as $\mathcal{O}(n)$, so it is ideal to find a more simple approximate solution. One way which has been used is to smooth a single position waveform in the time domain with a gaussian kernel. This is often a valid approximation due to the fact that the physical extent of the charge cloud will result in some events originating along faster or slower isochrones, with most being close to the center value.

Indeed, visually inspecting simulated waveforms from events occurring in a gaussian spherical spatial distribution about some point in the detector, there is clearly a spread introduced in the rise time. The result of this can be seen in Figure 3.9, where a larger cloud has been used to exaggerate the effect. This effect occurs because a cloud will spatially distribute charge over a region, some of which will drift slightly faster or slower than the center of the charge cloud. In addition, there will be some distribution along the same isochrone, with small waveform shape changes due to the azimuthal angle.

The use of a gaussian kernel to smooth the waveform in the time domain is an approximation of a true cloud. This was validated by simulating a large number of waveforms in a cloud around many points across a detector volume, and looking at the difference at that point among waveforms.

While it is straightforward to find the error between the smoothed waveform and the average of a cloud of waveforms, this numerical value is not particularly meaningful. Instead, we align all waveforms at their start time, and find the difference (in ns) between when the earliest and latest waveforms reach their 50 percent risetime. This corresponds then to the spread at a y-value of 50 in Figure 3.9. By finding this spread for each point in a grid across the detector cross section, we can estimate the validity of this assumption. Figures 3.10 and 3.11 show the results for an enriched and natural detector, respectively. For enriched detectors, the bulk is relatively uniform, though BEGes are less so. The result is that, while the charge cloud size parameter will be found for any position using the gaussian kernel, there is a nonuniform mapping between the position in the detector and the spread that a given size charge cloud will produce. In practice this is not a problem, but any use of the fitted charge cloud size parameter must recognize this and correct for it.

3.4 Electronics

In order to read out the signal from a germanium detector, it is necessary to have a few basic components in place. Generically, the raw current signal from a detector must first be amplified by a *preamplifier*, then perhaps further amplified by an *amplifier*. Then this amplified signal may be read out by digitizing the signal, resulting in a waveform for later, offline analysis. When waveforms are to be digitized they often use only a preamplifier, as the amplifier usually provides some shaping which is useful for compressing the signal into a single recorded point value, such as amplitude.

3.4.1 Readout Electronics for the MAJORANA DEMONSTRATOR

The electronics readout chain for the MAJORANA DEMONSTRATOR consists of a low-mass front end (LMFE), a two-stage preamplifier, and a digitizer. The exact implementation used is somewhat unusual, due to the specific low-radioactivity constraints in place. A schematic diagram of the principal components is shown in figure 3.12.

A current signal is formed inside the detector, which is directly injected onto the gate of a JFET on the LMFE. The LMFE is physically located inside the detector unit, approximately 1 cm from the detector itself;

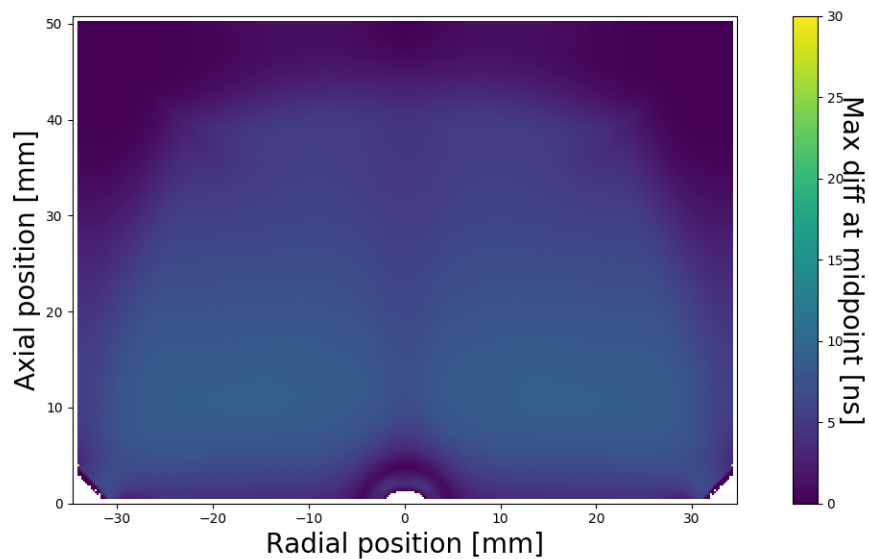


Figure 3.10: Color indicates the maximum difference in time at the 50 percent timepoint between waveforms in a cloud, at each point over the detector cross section. Detector is an enriched ORTEC detector, so the bevel and point contact are visible.




Figure 3.11: Color indicates the maximum difference in time at the 50 percent timepoint between waveforms in a cloud, at each point over the detector cross section. Detector is a BEGe, so the ditch outline is visible, with some edge effects.

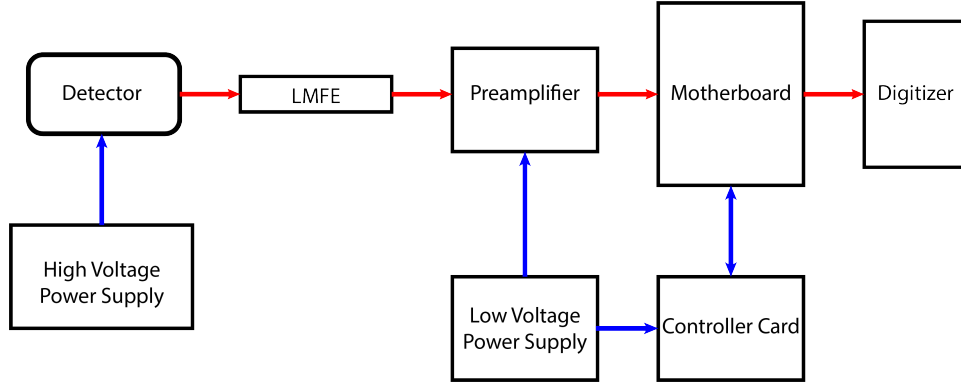


Figure 3.12: A high-level schematic overview of the MJD signal readout chain.

it is near the detector’s base temperature, slightly above liquid nitrogen temperature, and under vacuum. The LMFE is electrically connected to four coaxial signal cables, the *source*, *drain*, *feedback*, and *pulser*. These all connect to the warm electronics outside through several meters of cable, a vacuum flange, and multiple connector pairs. Immediately on the warm, non-vacuum side of the flange, the signals are passed to the “preamplifier card”, mounted on a motherboard which provides power.

The LMFE is electronically a component of the feedback loop of the preamplifier, but is physically distinct from the rest of the preamplifier due to background constraints, so the components are often considered as separate features.

The preamplifier design was performed at LBL, and is based on a folded cascode arrangement. The folded cascode design is ideal for producing a good output voltage swing and common-mode range with high output impedance [74]. The primary challenges of the MJD design are that it must be low noise and low radioactive background. The low background aspect is achieved by reducing the potentially radioactive mass of components near the detector, giving rise to a feedback loop of several meters in length. As the signal frequency increases, the time required for a signal to be transmitted along the cable eventually exceeds the time for the signal to go out of phase. This long feedback then limits the maximum risetime and bandwidth of the signals which are amplified.

3.4.2 Modeling of installed electronics

In order to develop a realistic model of the electronics, we performed numerical simulations of the full amplifier circuit between the detector and the digitizer. Although as-built schematics of the electronics were not available, we were able to infer the full design from the physical circuits.

Perhaps the most common software for modeling analog and mixed-signal circuits is the open-source SPICE software. While SPICE refers to a particular package, there are many variants of the original framework, designed either for particular analyses or adopted by electronics manufacturers for their own standard tools. The version used in this work is the open-source C++ based Xyce software, developed by a group at Sandia National Laboratory for solving large and complex circuits in parallel [75, 76]. While technically not a SPICE derivative, Xyce is completely compatible with SPICE, and should be indistinguishable from SPICE to most users.

SPICE modelers take as an input a “netlist”. This is a file containing a list of components, their properties, and to what each component is electrically connected, as a complete description of the circuit.

Using the SPICE model we were able to reproduce the basic features of the signals and to design a filter model which can be applied to simulated waveforms.

3.4.3 Electronics Effects on Waveforms

When visually looking at a MAJORANA germanium detector waveform, there are three primary features. The most obvious is the *rising edge*; the detailed shape of this contains information about how the charges moved through the detector before reaching higher field areas. The full rising edge usually occurs within about $1\mu\text{s}$ of the signal onset. The next obvious feature is the *falling edge* or *decaying tail* region of the falling edge. This is from the capacitive coupling of the first and second stages of the preamp, a high-pass filter with a decay time of around $75\mu\text{s}$. Finally, the *baseline* ahead of the rising edge contains no detector signal, and the higher frequency variation from the mean baseline value can be used as a measure of the noise. This baseline is usually at about 0V, but, due to a slow time-constant recovery of the signal, the entire baseline can also shift between subsequent events. The distinguishing point between the baseline and the rising edge is t_0 , the signal start time. Often, the readout is tuned such that the rising edge occurs near the center of the waveform. These three waveform regions are indicated in Figure 3.13.

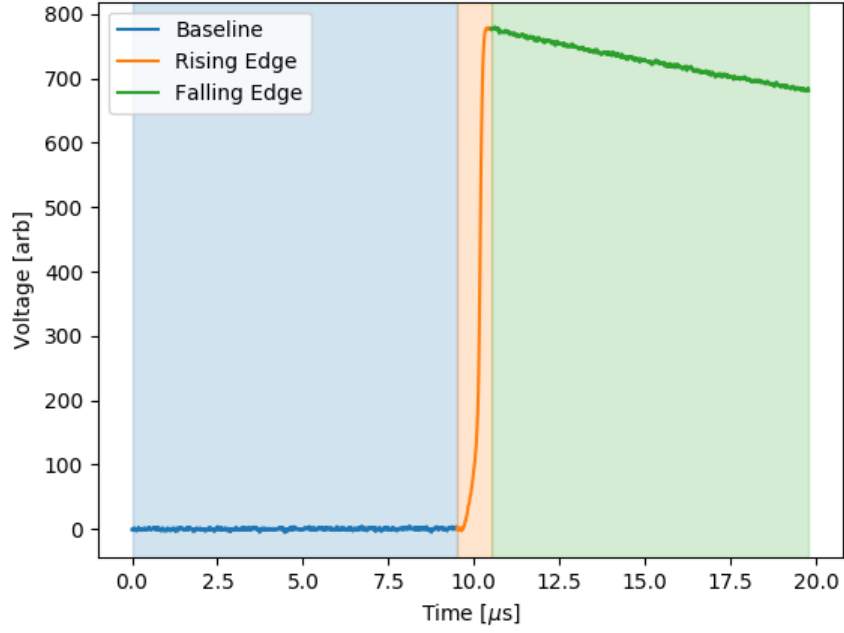


Figure 3.13: Regions of signals waveforms discussed in this chapter. The waveform is a voltage signal in the time domain.

The preamplifier also acts as a long-decay-time high-pass filter with a time constant around 3 ms. Within a single waveform, it is difficult to see the effect of this, however, multiple events in rapid succession will offset the total baseline due to this, so it is necessary to track.

Additionally, the preamplifier, like all real circuits, is incapable of reproducing arbitrarily high frequency signals, and acts as a low-pass filter. By comparing the response to real waveforms and those simulated using SPICE, the response is found to closely match that of a 2nd order filter.

3.5 Digitization, Readout, and Analysis

No analog signal may be recorded with perfect fidelity. Instead, a time varying signal can be recorded as a series of discrete voltage measurements at known time interval. This process is referred to as digitization, and a device which does so is a digitizer.

The digitizers used to read out the detector signals for the MAJORANA DEMONSTRATOR are custom devices that were designed for the GRETINA collaboration [77]. They are in the VME form factor, with 10

differential analog inputs per digitizer. The signals are digitized at 14-bit resolution, with a sampling rate of 100 MSPS (10 ns per sample), and a total maximum trigger rate per card of about 3 kHz.

These VME digitizers are connected over the VME backplane to a Single Board Computer (SBC), which can poll the digitizers for data, perform readout, and buffer data in place. The SBC is also connected to a computer through a network socket over a private ethernet network; this computer controls the run and stores the data file, as will be discussed somewhat more in Section 3.5.1.

As it is technology first developed in the 1980s, the VME form factor is now somewhat dated, but has some implementation advantages here [78], and active development is still done to improve VME systems [79]. The actual readout and implementation is relatively straightforward in a VME system, consisting of a relatively small amount of code to directly access and control the hardware for readout. Additionally, new cards are still produced today by several manufacturers at a modest cost. A more modern system such as μ TCA can handle much higher data rates and strict timing, but such a system is usually more expensive and requires more software overhead to manage a readout.

In the second stage of the preamplifier, the signal is split into two differential signals, a high gain and a low gain. These signals are both digitized (in adjacent channels). This is effectively a way to increase the dynamic range of the digitizers, as the high gain channel is sensitive from the lowest experimental threshold up to about 3 MeV events, whereas the low gain channel does not saturate until approximately 8 MeV with reduced low-energy sensitivity. Together, these two channels can be used to collect data over a wide energy range while still maintaining high sensitivity at lower energies. In general, most physics analyses in the MAJORANA DEMONSTRATOR use the high gain channels.

The digitizer channels are triggered independently using an onboard trapezoidal filter algorithm, similar to that described by Jordanov and Knoll [80]. The trigger threshold is set as an absolute value on the trapezoidal filter output, and the maximum value of the trapezoidal filter can be used as an energy estimate. The trigger then uses an estimated energy as a threshold.

The waveforms that are digitized may be stored in one of several ways. The waveforms are fixed at 2048 samples, for a fully-sampled record of length just over 20 microseconds. While the number of samples in the readout is fixed by the firmware, the digitizer can be configured at runtime to read out waveforms with certain regions “downsampled” to effectively increase the digitized length. This downsampling is accomplished by summing together multiple subsequent samples and storing them as a single sample, then optionally dividing by a multiple of 2. This allows, for example, the rising edge of the waveform to be sampled at the maximum

sampling rate while still sampling the falling edge of the waveform for a long period of time. Within the MAJORANA DEMONSTRATOR, this downsampling was used for some datasets.

The perfect ADC would, given an voltage range (e.g. 0 to 1 Volt) and a number of bits (e.g. 8 bits), split each of the 256 bins evenly among the voltage range; in this example, a single bit increment would then correspond to $1/256$ or 3.9 mV for each step. The 14-bit ADCs that perform the digitization for the DEMONSTRATOR, like all ADCs, exhibit some degree of nonlinearity. This means that out of the 8196 possible voltage step values the signal may occupy, not all “bins” are the same width. Therefore, a given ADC value may in fact correspond to a slightly different true voltage signal. There are various methods to measure and (at least partially) correct for this, measuring known repeated signals. For the work described here, no nonlinearity correction has been applied. This is due to the expectation that the nonlinearity is an effect that is smaller than the magnitude of the noise in most cases.

3.5.1 Readout and Control

The MAJORANA DEMONSTRATOR is a sufficiently complicated experiment that control and readout must be coordinated in an organized way. Rather than having many applications for controlling and monitoring various parts of the experiment, a single multi-featured application was developed, known as the Object-oriented Realtime Control and Acquisition or ORCA [81]. This software is of general use for controlling hardware and experiments, and is publicly available through an SVN repository [82].

Using a graphical interface, ORCA allows users to coordinate settings for all digitizers, control the readout, and store data. Additionally, ORCA instances are used for reading out the muon veto system, as well as monitoring all environmental sensors and controlling the cryogenics and vacuum systems. ORCA stores run data in hybrid file consisting of an XML-based (Apple p-list) header and a body of binary data, which can later be processed. Instead of using these binary run data files, most slow controls data is instead stored in a couchDB database which is replicated offsite and available for near-realtime monitoring.

A more detailed overview of the data acquisition used by the DEMONSTRATOR is given in Appendix C.

3.6 Complete Signal Simulation

Using the models of both detector and electronics described in the previous sections, a realistic simulation of signals has been developed. The total simulation model can be split into several parts, which are described here.

3.6.1 Detector Model

The detector is modeled with several distinct components, including the electric fields within the detector, an impurity gradient, a velocity model for mobilities, and a simple charge trapping estimate.

The simulation of events inside a germanium detector begins by simulating the electric field and weighting potential inside the detector. The impurity profile superimposes an additional field on top of the applied electric field, which helps to fully deplete the detector volume and funnel charges towards the point contact.

In addition to the components considered, there are several other possible additions to the detector model which are not considered here. These could include the modeling the impurity gradient in other ways, such as adding a radial component or a non-linear longitudinal gradient. Additionally, the detector may hold surface charge changing the fields near the surface. The lithiated n^+ surface thickness could also be varied. Consideration of these other components may be valuable for future work.

3.6.2 Field Simulations

The first step in simulating a waveform is calculation of the electric field and associated weighting potential inside the detector volume. The impurity gradient inside the detector is estimable but unknown, and will be determined by fitting waveforms, as will be described in Section 4.3.5.3. Once the impurity value is selected, the field can be calculated directly.

Using the geometry described in the Detector Model, as described by Section 3.6.1, we can numerically solve the Poisson equation to find the electric field and weighting potential inside the detector. The MAJORANA collaboration makes use of David Radford's *siggen* [83] and *fieldgen* [84] packages for this purpose.

The *fieldgen* package performs a numerical relaxation to solve Poisson's equation to find the electric potential and weighting potential in a detector. The electric field can vary depending on many things,

including the exact voltage applied across the detector, the detector temperature, the impurity gradient, and the exact geometry.

Using a python-based version of the *siggen* software package, we solve for the fields. In this adaptation, the actual field calculations are done using the FEniCS framework, a mature toolset for solving partial differential equations (PDEs) using finite element calculations [85]. Using the geometry information determined by Starrett measurement (see section 3.2.3) we can then define the domains of the p^+ and n^+ contacts, the passivated layer, and (for BEGs) the ditch location. Then, we define a simple rectangular mesh over the detector volume. Finding the electric field is the simple case of solving the Poisson equation for a scalar potential Φ ,

$$\vec{E}(r, z, \phi) = -\vec{\nabla}\Phi \quad (3.7)$$

with the assumptions that the n^+ contact is held at an operating voltage V_{op}

$$\begin{aligned} \Phi(p^+) &= 0 \text{ V} \\ \Phi(n^+) &= V_{op} \end{aligned} \quad (3.8)$$

in the presence of a charge impurity distribution held over the top.

This results in an electric field as shown in Figure 3.14, where the field is relatively uniform throughout the bulk of the detector.

The weighting field considers the same detector absent any charge, and therefore the Laplace equation is used. In order to solve for the weighting field, we solve the Laplace equation under the boundary conditions

$$\begin{aligned} \Phi(p^+) &= 1 \text{ V} \\ \Phi(n^+) &= 0 \text{ V}. \end{aligned} \quad (3.9)$$

This results in a weighting potential as shown in Figure 3.6. As it is usually most convenient to perform calculations using the weighting potential rather than the weighting field, the potential is indicated rather than the field.

For detectors with additional readout electrodes a weighting field must be found for each electrode individually leading to fields which are unique for each electrode.

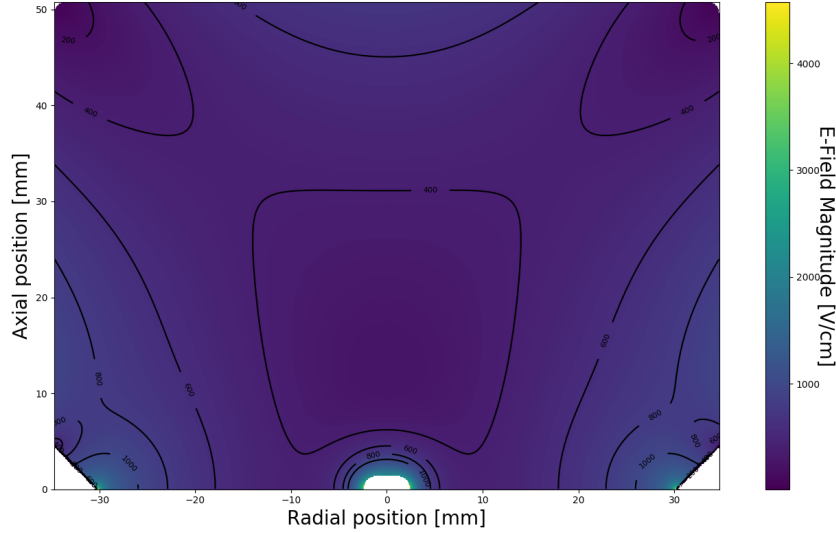


Figure 3.14: The electric field viewed cross-sectionally inside one ORTEC geometry (enriched) PPC detector (P42575A). The highest field location is near the point contact.

3.6.3 Electronics Model

When simulating the electronics response, we must define a set of filters to be applied to the raw waveform. These filters include a series of high- and low-pass filters, to model the known response of the preamplifier.

In the actual circuit, an analog signal is passed through a series of analog filters, and digitized. As we do not have access to the true continuous analog signal (even in simulation), we approximate this in a digital representation. Even before considering the effects of electronics, the simulated waveform is discretized due to the numerical techniques that are employed to calculate the raw waveform signal. The model we use passes this discretized signal through a series of discrete time filters. By numerically sampling at a higher rate then downsampling back to the rate of the digitized signal, we can eliminate any effects from this.

In order to perform the digital filtering required, we use a signal processing technique known as the Z-transform, which is the discrete time equivalent of the laplace transform. For more details about the use of the z-transform in signal processing, see Appendix B

The electronics models considered are indicated in Table 3.6.3. As the *waffle* electronics models are specified in a pole-zero format, there is sufficient flexibility to describe any electronics features, including others not described here. In addition to these electronics models, there are three detector models which

are also used for fitting against, as indicated in Table 3.6.3. Each of these filter models contains parameters which can be fit, as will be described in the following chapters.

Filter Name	Description	Params
High-Pass Filter 1	A 1st order high-pass filter with an approximately $75 \mu\text{s}$ decay time	1
High-Pass Filter 2	A 1st order high-pass filter with an approximately 3 ms decay time	1
Low-Pass Filter	A 2nd order low-pass filter.	2
Overshoot Filter	A 1st order high-pass filter with an approximately $2 \mu\text{s}$ decay time	2
Oscillation Filter	A 2nd order filter to add a small amount of ringing to the waveform	2
Antialiasing Filter	A 4th order low-pass filter designed to look like the digitizer input	2

Table 3.2: Electronics filter models considered for this work.

Model Name	Description	Params
Velocity Model	Charge carrier drift velocity model	4
Impurity Model	A linear gradient impurity profile	2
Trapping	A simple charge trapping model	1

Table 3.3: Detector models considered for this work.

3.6.4 Waveform Simulations

Once the full field and weighting potential are determined, the individual event parameters must be chosen, including the energy and the position in the detector.

Using siggen, we can then place an electron-hole pair in the defined location, and numerically step it through the detector, calculating the induced signal at each point. The total waveform signal is then the collection of induced signals at every time step. After simulating a single raw waveform, the electronics models described in the previous section can be applied, producing a complete simulated signal.

The model of the waveform is described using six parameters, three to describe the location coordinate, one to describe the amplitude, one for the event start time, and one for the size of the charge cloud. These are indicated in Table 3.6.4. Note that simulation of a single waveform depends on the detector parameters described in the previous section, however, these waveform-specific model parameters are expected to be independent between any two waveforms. These parameters will be extracted from waveforms by fitting.

Parameter Name	Description
r	The radial cylindrical coordinate of the event position within the detector volume
z	The cylindrical height coordinate of the event position within the detector volume
phi	The azimuthal angle coordinate of the event position within the detector volume
scale	The amplitude scaling factor
t_0	The event start sample number within a waveform
cloud	The (optional) charge cloud size parameter

Table 3.4: Waveform model parameters considered for this work.

3.7 Conclusions

When radiation is incident upon a germanium radiation detector, it creates an electrical signal which depends on many factors. These include detector effects such as the geometry and impurity profile, and the event-specific effects such as the interaction position and energy. After being amplified and digitized, the signal may somewhat obfuscate its origins. However, using a set of detailed models, we can accurately reproduce signals to match those expected for most of the detector volume.

We will continue in the next chapter to describe a technique for using this set of models to perform machine learning on germanium detector waveforms.

CHAPTER 4

Bayesian Inference, Fitting, and MCMC Methods

In this chapter, we will discuss the generic idea of Bayesian inference, its use in performing a *Bayesian Fit* and the Markov-Chain Monte Carlo (MCMC) methods which are used therein.

4.1 Overview

The model described in the previous chapter shows that there are many details which can contribute to a total waveform signal. In order to determine which specific values of each parameter contributed to producing a specific waveform, we would like to fit the data of each waveform using our model for how this signal is created.

In general, a fit is a way to estimate the parameters of a model by optimizing some figure of merit between the data and a parameterized model, often in an iterative process. Simple fits can be performed using a basic least-squares regression, in which the model parameters are chosen to minimize the error or maximize some goodness-of-fit, however, simple techniques are insufficient for more complex problems, and may become trapped by local error minima and produce results far from the correct or expected values.

The model for a detector waveform signal is highly dimensional (has many parameters) and relatively computationally expensive in that a charge's motion must be simulated through a field in order to produce a waveform, rather than evaluating a simple function at a point. The high dimensionality means that more sophisticated methods must be used to prevent local maxima from trapping the fit into false regions of relatively higher probability. The computational complexity limits the number and type of models that can be practically tested, as well as creating a computational bottleneck in processing large amounts of data. As a result, we will describe alternative methods which can overcome some of these challenges.

4.2 Statistical Methods

4.2.1 Bayesian Statistics

It is worthwhile to begin by introducing some practical points about Bayesian statistics. Many of the concepts that are discussed in this chapter will make use of common Bayesian statistical nomenclature, which are, for clarity, defined here.

Within statistics, Bayesian and Frequentist statistics are two primary philosophies on the analysis and interpretation of results. While their respective analyses may each result in statements of probability, the meaning in each case is distinct.

The result of a frequentist analysis is a statement about the frequency of an event in a high-statistics limit. Frequentist probability is the limiting case of a large number of repeated measurements.

In a Bayesian analysis, the result is a statement of how likely a particular event would be, given the data we have seen and our knowledge or expectation of the system. Bayesian probability is a degree of certainty in a particular result, conditional on this prior knowledge and in the context of a particular model.

There is some discord over whether a Bayesian analysis is too subjective, or a frequentist analysis neglects to include information. In the analysis frequently encountered in the sciences, it is often most natural to perform Bayesian analysis [86].

The fundamental expression used in Bayesian statistics is known as Bayes rule, and it relates the probability of a specific model M with a set of parameters θ given a particular set of data D . It is expressed as

$$\underbrace{P(\theta|D, M)}_{\text{posterior}} = \frac{\overbrace{P(D|\theta, M)}^{\text{likelihood}} \overbrace{P(\theta|M)}^{\text{prior}}}{\underbrace{P(D|M)}_{\text{evidence}}}. \quad (4.1)$$

The posterior probability $P(\theta|D, M)$ is the probability of a specific model M with set of parameters θ being the correct model, given a set of data D . Often the model M is implied and not explicitly written, but for clarity we will include it here. Finding the posterior distribution is generally a primary goal of performing a Bayesian statistical analysis.

The likelihood $P(D|\theta, M)$ is the probability of measuring some data D given a particular model M and a set of parameters θ for that model. This likelihood is the same that would be used in a frequentist maximum likelihood analysis.

The model prior $P(\theta|M)$ is the probability of the model M with a set of parameters before considering any of the data D . The exact choice of this prior is somewhat subjective (indeed, this is the source of the subjective inference in Bayesian analysis), and therefore the cause of many protracted arguments.

The final $P(D|M)$ term is referred to by various names, including the marginal likelihood, the model evidence, and the data probability. It often takes on the variable Z in analogy to the partition function of statistical mechanics. This term is the total probability of measuring the data independent of the specific model under consideration; in this sense, the model has been “marginalized” out. It is generically inconvenient to calculate, as it is the total sum of the joint probability (likelihood times prior) over the entire prior space, and it is rare that any algorithm will truly step over this full space explicitly. The marginal likelihood is a constant which acts to normalize the posterior, and for many analyses its calculation is not strictly required (as the best value of the posterior is independent of the total scale). It is useful when comparing different models to each other. We will discuss methods for calculating its value in Section 4.2.3

In totality, Bayes rule expresses a way to calculate a conditional probability. Given a set of parameters and a model, some value can be found, which can be compared to the data in a likelihood. Based on how likely we thought this value was (the prior) and how likely it was to measure that data at all (the evidence), we can calculate the total probability of a given parameter set, conditional on this information.

The values filling the roles for these terms in our specific case will be outlined later.

4.2.2 Markov Chains

A Markov chain is a basic statistical model consisting of a series of states in a given parameter space. From each state, there is a transition probability to move to another given state, or stay in the same state. The transition probability is not dependent on any previous state history, but instead only depends on where in the prior space you are transitioning to next (and where you are at that moment).

The chain itself consists of the list of all states which have been visited thus far. As is evident from the definition, computation of the next step in a chain depends only on the current value, the last value in the chain at any given time. After the chain has run for a sufficient quantity of steps, it reaches a limiting

distribution, where the spread in parameter values in the state space indicates the uncertainty about a best value. Statistically then, the entire space is sampled.

Using the concept of a Markov chain it is possible to model various processes, or to find the limiting distribution of a sequence. As a tool for Bayesian inference, Markov Chain Monte Carlo (MCMC) allows you to draw samples from an unknown (posterior) distribution by simulating the values in a Markov chain and finding a limiting distribution of the samples.

4.2.2.1 Markov Chain Process

Most MCMC algorithms are based on the Metropolis-Hastings technique [87, 88], with many variations on the particular acceptance criteria and sampling algorithms.

Markov chains begin with a guess point/value θ (usually pulled randomly from the expected prior space); in reality, θ is a vector containing a value for each parameter in the model under investigation. Using the guess value and a model, the likelihood $P(D|\theta, M)$ and prior probability $P(\theta|M)$ are computed. The product of these two terms (together referred to as a joint distribution) is proportional to the posterior value of that guess value θ , and based on the value, the sample can either be saved or thrown away (and then a new sample would be taken).

Using a proposal distribution, the next proposed sample can be found. Usually, this distribution is a gaussian function of some width, centered about the current sample point. The choice of width is a hyperparameter which will affect the convergence for a given model (see Section 4.2.3.3 for more information about hyperparameters).

The choice to keep a proposed sample or not depends on the ratio of the posterior values at the current position and the proposed jump position; this concept is core to the Metropolis-Hastings algorithm and the detailed way in which these proposals are chosen, and accepted or rejected differentiate many MCMC algorithms. In a standard Metropolis-Hastings algorithm, the posterior ratio is compared to a randomly selected number between 0 and 1, accepting jumps to states with sufficiently improved posterior density.

For a large number of samples, this resultant distribution can be used as an estimate of the total posterior distribution.

4.2.3 Nested Sampling

Nested sampling is an MCMC sampling technique for evaluating the marginal likelihood [89]. This technique also samples the posterior distribution, as needed for parameter estimation problems. Aside from the efficient calculation of marginal likelihood, the primary advantages of Nested Sampling over other sampling (such as Gibbs or tempered annealing) are that NS is capable of exploring highly-dimensional and correlated parameter spaces without excessive hyperparameter tuning [90].

While many MCMC methods can be theoretically expected to eventually converge to the true best result, this is cold comfort, as it may take a long time or never actually happen in practice.

The goal of Nested Sampling is to yield Z , the evidence. This quantity is often disregarded due to the challenge in calculation, but is a useful value for model comparison. Without it, the resulting “posterior” is still informative, but lacks normalization; the shape and maximum value are unaffected by normalization. Here, the posterior is a by-product of calculating the evidence; in principle we are primarily interested in the calculation of the posterior (the parameters describing each of our many models), yet Nested Sampling algorithms also provide an efficient way of exploring large parameter spaces.

In recognition of the algorithmic nature of this technique, it is common to express Bayes rule more succinctly as

$$P(\theta) = \frac{\mathcal{L}(\theta)\pi(\theta)}{Z} \quad (4.2)$$

rather than the formal probabilistic notation of equation 4.1. The likelihood $P(D|\theta, M)$ is expressed as $\mathcal{L}(\theta)$, the prior $P(\theta|M)$ as $\pi(\theta)$, and the evidence $P(D|M)$ as simply Z , emphasizing their functional behavior.

Note that this expression can be rearranged

$$Z \times P(\theta) = \mathcal{L}(\theta) \times \pi(\theta). \quad (4.3)$$

In order to maintain normalization, the prior and the posterior must each sum to unity, as

$$\int_{-\infty}^{\infty} \pi(x) dx = 1 \quad (4.4)$$

$$\int_{-\infty}^{\infty} P(x) dx = 1 \quad (4.5)$$

Integrating both sides of Equation 4.3 over the parameter space, but recognizing that Z is a constant, we find

$$Z = \int_{-\infty}^{\infty} \mathcal{L}(\theta)\pi(\theta)d\theta \quad (4.6)$$

Changing variables as $dX = \pi(\theta)d\theta$, we then have

$$Z = \int_0^1 \mathcal{L}(X)dX \quad (4.7)$$

where $\mathcal{L}(X)$ is a likelihood method that accepts a prior mass X as an argument. This dX is a differential element of prior mass, where the prior mass in a region of the prior space is the sum of prior probability values contributing to that region. In order to find Z it is then a “simple” matter to evaluate this integral. The Nested Sampling algorithm is an efficient method to do so statistically.

4.2.3.1 Nested Sampling Algorithm

The evidence Z may be defined by the integral of the “sorted likelihood” over the full prior mass, as in Equation 4.7. Determining the necessary likelihood function mapping to perform this integral is challenging. Were it possible to sort the prior coordinate values by their likelihood values (irrespective of prior coordinate θ itself), you could then draw full contours of constant likelihood in the parameter space. For any given value of likelihood, you would know exactly how much prior mass X was remaining above and excluded below that value; that is, you would have a one-to-one mapping from the parameter space θ to the fraction of prior mass with likelihood greater than that point’s likelihood.

While in a small, low-dimensional problem it may be possible to raster across all values in the parameter space θ and approximate this (or calculate it exactly for a discrete space), in practice it is not possible to sort all points θ in order of their likelihood values, due to the massive parameter space available in most real problems. Instead, it is possible to simulate the sorting statistically, by a constrained sampling method known as *Nested Sampling* [89, 91]¹.

¹Note among the references on Nested Sampling there is no consistent notation for almost any of the parameters of interest. Indeed, some variables are used for similar but distinct quantities between them. Diffusive Nested Sampling, described in the following section, uses also a distinct notation. We have chosen a single set of variables to describe the relevant parameters, which will then be different from other sources.

The technique uses a set of objects known as “live points” or “particles”, which are points randomly drawn from the prior space and evolved through it. The particles are continually constrained to have likelihood values which exceed a given likelihood value. At every step, the particle with the lowest likelihood is identified and replaced with a new sample, while the other particles are sampled about their present locations. The likelihood \mathcal{L}_k of this worst particle defines the k^{th} likelihood contour (level), the hard sampling constraint which all other particles are then subject to. Each level then compresses the sampling into regions of smaller prior mass. Given the required shrinkage between contour levels, the value of prior mass X assigned to this k^{th} level is then

$$X_k = e^{-k/n} \quad (4.8)$$

for n particles. With this X_k , we can assign a weight (area) to this level in a simple Riemann sum as

$$w_k = \mathcal{L}_k(X_{k-1} - X_k) \quad (4.9)$$

with the total evidence Z being given by the total sum of w_k for all m levels, as in

$$\sum_{k=1}^m w_k \rightarrow Z. \quad (4.10)$$

Eventually, it becomes more difficult to increase the likelihood further, causing the value to level off with continued sampling. At the same time, the widths between subsequent X_k decrease, however, the the likelihood will at first increase faster than the widths decrease. Eventually, the likelihood flattens off enough with additional sampling that the level weight again decreases. This leads to a peak being formed in the posterior weight of these levels; these levels contribute the most to the evidence. As will be mentioned in the following chapter, this is *an* indication of convergence, and *may* show a sufficient level of sampling.

There are several software packages that can presently perform nested sampling, including Nestle, CPNest, dynesty, UltraNest, PyMultiNest, DNest4, and PyPolyChord, each having some variation on the basic algorithm. The Diffusive Nested Sampling algorithm used by the DNest4 package has been employed for this work, and is described in the following section.




Figure 4.1: The distributions sampled by both Classic Nested Sampling and Diffusive Nested Sampling as the chain proceeds forward. The Diffusive Nested Sampling algorithm allows for continuous mixing at shells of lower likelihood. Adapted from [90]

4.2.3.2 Diffusive Nested Sampling

Diffusive Nested Sampling (DNest) is one further revision of the nested sampling technique developed primarily by Brewer[90, 92]. It follows the basic technique of the standard Nested Sampling algorithm, with some important variations. The “diffusive” nature of the method comes by periodically accepting points outside of the inner likelihood “nest” with a length scale denoted Λ . This backtracking is designed to allow for continuous exploration of the lower likelihood levels while still in search of new, higher likelihood levels. In Figure 4.1, the standard and diffusive nested sampling algorithms are compared; in standard nested sampling, once a higher likelihood nest has been defined, it is no longer possible to sample again back in the areas believed to have lower likelihoods. The color in this figure indicates the relative amount of sampling done in a given region.

A peculiarity of Diffusive Nested Sampling is that it internally keeps track of how many samples have been acquired in each level, and compresses level spacing accordingly, but also can reject samples in order to sufficiently explore other regions. In this way, it is not a true Markov Chain, as it relies on past knowledge of how the chain has progressed; however, despite this, as the chain progresses it asymptotically approaches a Markov chain [90].

The Diffusive Nested Sampling algorithm is ideal for use in this study. The algorithm requires more effort to specify a model than most other MCMC tools, however, it is robust and relatively efficient, particularly for use in high-dimensional fits. Other compelling sampling techniques such as Gibbs or Hamiltonian MCMC are

more likely to give poor results when faced with multimodal posteriors; annealing and tempering algorithms can in principle handle this better, but often require excessive tuning of parameters [92].

4.2.3.3 Hyperparameters

Hyperparameters are parameters that define how an algorithm will be performed, but are not explicitly related to the model being fit itself². They are often ways of directing the fit or learning technique to proceed, such as the temperature profile in a simulated annealing or the layer structure in an artificial neural network. The values used for these do not generally depend on the data explicitly and generally must be manually tuned for a particular problem.

Within Diffusive Nested Sampling, the backtracking scale length is described by a hyperparameter Λ . Smaller numbers can be chosen for simple fits, with higher values being slower but more fail-safe; the exact choice is rarely important in deciding whether a fit will succeed or fail, but may impact the relative efficiency of convergence. In addition, there is an *equal weight enforcement* hyperparameter β , which allows the fit to enforce the desired level weight mixing of $w_j \propto 1$.

The number of particles used varies based on the computational resources available. By parallelizing the DNest4 algorithm, processor cores can each be assigned to calculating likelihoods for a single data waveform and single particle. In this way, increasing the number of particles is expected to increase the speed at which sampling can be performed.

The *new level interval* hyperparameter is the number of samples with likelihood greater than the current constraint which must be accrued before creating a new level. Other hyperparameters include the sampling interval between levels.

Within Diffusive Nested Sampling, these default parameters are sufficient for almost all tasks. Unlike in other techniques such as simulated annealing, the levels are all compressed by the same factor, so there is no need to have a “temperature profile” to describe the full problem ahead of time.

4.3 Fitting waveforms

From Chapter 3 it is clear that the realistic simulation of a waveform includes many detailed models and parameters. It would be ideal to be able to recover the values of all the included parameters after taking

²The usage of the term *hyperparameter* is a bit different in Bayesian statistics and Machine Learning, but we are using the Machine Learning definition here.

Symbol	Description	Value used
Λ	Backtracking scale length	10
β	Equal weight enforcement	100
—	Number of particles	~ 100
—	Thread Steps	100
—	New Level Interval	10,000

Table 4.1: Hyperparameters used in DNest4, with the values used in this study indicated.

data, to ensure that the presumed model is indeed reasonable. Unfortunately, the larger number of models and parameters gives a high dimensional parameter space with possibly large and unknown correlations between parameters. Furthermore, determining the signal created by a given set of parameters is relatively computationally expensive. For these reasons, an advanced fitting technique that is capable of navigating the parameter space efficiently – such as is described in the previous sections – is necessary in order to perform fits on germanium detector waveforms.

4.3.1 Data Processing

The DEMONSTRATOR experiment produces raw ORCA data run files, approximately hourly in background running, and every few minutes during calibrations. The process of converting this data into a format for performing fits is described below.

First, the raw ORCA data is processed using the python-based *pygama* package to produce “Tier 1” data (T1). This converts the binary waveform data from the Gretina digitizers into Hierarchical Data Format (HDF5) files using custom built decoders. HDF is a modern self-descriptive (binary) data format designed for general scientific use with large amounts of data³. The Tier 1 data is effectively just a data formatting step, and prepares the data for manipulation and analysis using the *pandas* python framework. Tier 1 data includes the full information available in the raw data packets, including the raw waveforms, the event timestamp, and the onboard energy value.

The Tier 2 data processing performs a series of transforms and calculations on each waveform to produce a Tier 2 HDF file. Among other things, it performs an energy calibration, calculates a series of timepoints, calculates the maximum current, and calculates the baseline slope and intercept.

³While probably no data format is perfect, HDF5 is certainly adequate for this use, and is expected to be supported for the long-term future. Perhaps other formats, such as Advanced Scientific Data Format (ASDF) may be appropriate for consideration in the future, but for now they are probably not mature enough.

The details of the selection and processing done for training is outlined in Section 4.3.5.2.

4.3.2 Machine Learning Techniques

Many problems have sufficient complexity that it is necessary to algorithmically perform some steps of the analysis, rather than giving explicit descriptions of the procedure.

A detailed waveform simulation requires knowledge of parameters which can be classified in two types: those that are common to all waveforms from that detector, and those that may vary for any given waveform. The common parameters would include things like the detector geometry, the impurity gradient, and the applied voltage. The parameters which may vary between subsequent waveforms could be the event position in the detector or the energy deposited.

It would be both computationally challenging and wasteful to try to recover all of these parameters for each waveform independently, particularly when it is known that many of the parameters will be shared. Due to their large number and relative complexity, waveforms are a good candidate for machine learning.

The simulation of a given waveform needs both the shared and individual parameters, but the fit should be able to find common values for all waveforms for the shared parameters; the values should not change from waveform to waveform, so forcing them to find common values should provide stability to the fit. Ideally, we would be able to fit a large number of waveforms for a given detector, possibly all or a large subset of all waveforms in a dataset. It would be both computationally intractable and unnecessary to fit all of these waveforms simultaneously, as that would be a many-thousand dimensional fit. Instead of fitting all waveforms simultaneously, we use a form of machine learning in order to deduce the shared parameters, then fit the individual waveforms with knowledge of the shared parameters.

Machine learning (ML) is a very broad term, which encompasses a number of very different techniques. They may be supervised or unsupervised, and the outputs can be a variety of data types (as in clustering data into groups, or identifying a single trend).

Some ML techniques (including many artificial neural networks) are black-box in nature, akin to fitting data with a 10^{th} -order polynomial with disregard for the meaning of each coefficient. Other techniques are model-based, and allow an algorithmic technique to establish patterns in data within a given model.

The machine learning architecture used here consists of three primary parts: model selection, training and waveform fitting.

Model selection is a process of identifying which detector and waveform models are most ideal for accurately reproducing the observed waveforms, regardless of their exact model parameters. While this step is an input into the fits, it is not explicitly a machine learning step. This will be further discussed in Section 4.4.

Training is done by selecting a small set of waveforms from calibration data for a given detector. The selected waveforms are fit as a single simultaneous fit, with the detector parameters held in common. The result of this step is a set of detector parameters which should be applicable for any waveform from this detector. This step is considered “training” because it can find the parameters which will best suit the detector, and is the precursor step necessary to understand a given detector. The outputs of this step are valuable on their own, as they can identify detector characteristics.

Single waveform fits are the fastest step, and consist of lower-dimensional fits which use the detector model produced by the training step as fixed inputs rather than fit parameters. The output of this step is the set of parameters which vary among subsequent waveforms.

4.3.3 Bayesian Inference of Waveforms

Bayesian inference is the process of using Bayes’ Theorem to estimate parameter values based on data but conditional on prior knowledge. As described in Section 4.2.2.1, the machinery of MCMC can be used to make this inference feasible on waveforms.

Here, the θ vector becomes the collection of all parameters that feed into a waveform model, which are then used to simulate a waveform. These will include the shared detector parameters, and the set of waveform parameters for that particular event, as indicated in Section 3.6.1 and 3.6.4. An example of such a set of parameters used for fitting is shown in Figure 4.2.

The likelihood of a waveform is found by simulating a waveform and calculating a gaussian likelihood between each of the n sample points in the waveform, as in

$$\mathcal{L}_i = \text{Normal}(d_i - m_i, \sigma_d^2) \quad (4.11)$$

where d_i and m_i are the i^{th} elements of the data and model waveforms, and σ_d is the model error, taken as the baseline RMS noise of the waveform. For the total waveform the likelihood is the product of each sample’s likelihood, giving

$$\ln(\mathcal{L}_{wf}) = \sum_{i=0}^n \ln(\mathcal{L}_i). \quad (4.12)$$

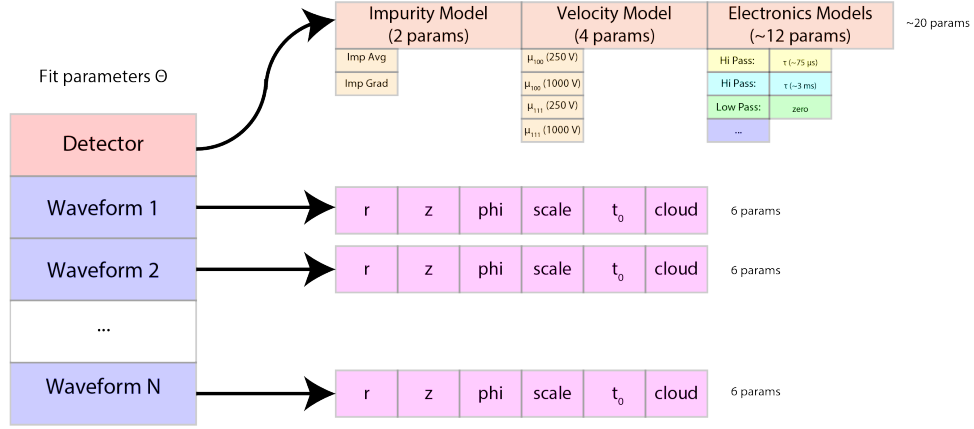


Figure 4.2: Data structure for training fits. Other parameters (such as detector geometry) are not allowed to float. When fitting a single waveform, the fit uses the shared detector parameters without floating them and the single waveform parameters are fit to match that waveform. Electronics models are indicated in Table 3.6.3.

A total training model containing multiple waveforms then has a total log likelihood given by the sum of each contributing waveform's log likelihood.

4.3.4 Waffle

In order to organize and simplify the fitting of waveforms, a python package known as *waffle* has been developed. The implementation of a model for use with DNest4 is considered to not be easy, according to the software's authors [92]. The *waffle* software is designed as a modular tool for describing the full set of models to be fit, as well as managing fits (both training and waveform fits) and postprocessing their results. This software is used for performing all fits described in this thesis.

4.3.5 Training Fits

4.3.5.1 Parameters

Training fits identify detector-specific parameters. As described in Table 3.6.3 and Table 3.6.3, the parameters that are fit within the context of the detector training model are:

- The **electronics model** parameters, described in Table 3.6.3.
- The **impurity profile**, described by 2 parameters, the impurity concentration at each end of the detector.

- The hole **velocity model**, described by 4 parameters (corresponding to the low and high field mobility in the $\langle 100 \rangle$ and $\langle 111 \rangle$ directions).
- A simple **charge trapping** parameter, described by 1 parameter, which reduces the remaining charge by a constant fraction every step of the signal simulation.

The priors used for these models are defined by various different expected values. The electronics model parameters are defined by pole and zero positions, with a uniform prior in a reasonable range to select them. The high-pass filter parameters are estimated during the data processing step, and the mean value is used to set this range. The velocity model parameters are defined with gaussian priors about the values expected for germanium, as described in Section 3.3.0.1. The impurity profile parameters are the impurity values at each end of the detector, and the prior that describes their values is a uniform prior along a range of reasonable values, as determined from the detector vendor measurements.

4.3.5.2 Data Selection

Training on a particular detector begins with selecting the waveforms to be used. When training is performed, it is necessary that a subset of data must be used to represent the full population. Therefore, this small number of waveforms that is chosen must be chosen carefully to avoid introducing biases into the training process. In most cases, there are later validation steps which could identify poorly selected samples, but it is nevertheless ideal to avoid the problem before spending computation time fitting them.

The data should be free of any events that are not produced by the expected standard signal process, including electrical noise and digitizing artifacts. Ideally, the data would have a high signal-to-noise ratio (SNR), and be free of pileup. For these reasons, we have made no attempt to perform fits using low energy waveforms.

Waveforms are selected using an A/E cut to remove multisite events, and a time cut to eliminate events which occur too recently after another event. The A/E value is defined by the maximum of the current pulse divided by the energy of that pulse. A pulse with a high A/E collected its charge within a short period of time, and a low A/E indicates that the charge collection is spread in time. The most common use case for A/E is for identifying events where the charge is split between two current pulses in near-coincidence, with the total energy of the waveform approximately given by the sum of that from each contributing event.

Additionally, the drift time is used to eliminate far outliers. The drift time is calculated as the time difference between the estimated signal start time (t_0) and the time at which the signal reaches 99% of its maximum value, known as t_{99} .

4.3.5.3 Impurity

When navigating the available parameter space, the field inside the detector will in general depend on the parameters chosen. We do not allow the physical dimensions of the detector to be floated in the fit, but we do fit the impurity gradient. As a practical matter, we would like to be able to simulate any fields for a continuous set of impurity gradients while running fits, but doing the full field simulation is too computationally expensive to perform in the likelihood evaluation. Instead, we can precompute a range of reasonable possible fields with different impurity gradients, and interpolate between them as needed. The impurity gradient is defined by the impurity concentration at each end of the crystal, and by choosing a range of around 5 values at each end, we can calculate around 25 fields which should contain the true value inside.

4.3.5.4 Number of waveforms

Ideally, as many waveforms as possible would be used during the training, in order to pool data from as many distinct parts of the detector as possible. This is in tension with the need to make estimating the likelihood for a set of parameters as efficient as possible.

Training fits have been run with as few as 2 waveforms, and as many as 32. When investigating convergence, it was found that training with only a small number of waveforms, say 4 or 8, the results were able to match those of larger waveform sets. Therefore, to aid with efficient testing of many model sets and many training datasets, most training has been done using 4 or 8 waveforms at once. This will be discussed somewhat more in Section 5.2.

4.3.6 Waveform Fits

Single waveform fits proceed using the detector model parameters found in the training step as true values, without floating them. Instead, only the six parameters indicated in Table 3.6.4 are allowed to float. The selection of waveforms for these fits is less stringent than on training fits, as fitting a single waveform is less computationally expensive, and we can therefore risk trying to recover parameters from more waveforms.

A typical example of a single waveform which has been fit is shown in Figure 4.3. The data and simulated




Figure 4.3: Waveform fit comparison for an average waveform. The horizontal scale indicates a time parameter. The vertical scale on the top figure is a voltage, and the bottom figure indicates the residual difference between a data waveform and the proposed simulated value. The data and fit overlap considerably, with a residual feature at the rising edge.

waveforms overlap nearly perfectly, with a maximum residual at less than 1% of the waveform itself. The rising edge contains the most detailed information about the relative motion of electrons and holes in the detector, and therefore it is not surprising that the residual is greatest at this point. When investigating the residuals from many waveforms together, it is clear that they usually add constructively around this rising edge feature, which indicates that the model may lack some detector physics. The overall fit is quite good for most waveforms, however.

4.4 Model Testing

The models described in the previous sections and chapter describe the physics of signal creation, and the modification of the signal electrically before it is digitized.

The model selection process compares different sets of models to determine which is the best to describe the system under investigation. In general, it is not possible to identify the true correct model, but rather a model which is sufficiently good to provide predictive power in the parameter region of interest. In addition, some model selection schemes aim to identify models which will perform well on other data which has not been observed.

There are really two sets of models to consider validating, the training model and the waveform model. The waveform model contains only basic parameterization, and is well established from other studies. The training model contains several waveform models as a subset, and the variation between training models is between different parameterizations of the electronics and detector. For this reason, we only consider model comparison between training models.

There are various metrics to use for model selection, including Bayes factors, cross-validation, and various information criteria such as AIC, WAIC, DIC, and BIC. Each has various implementation and interpretation advantages, and we will briefly discuss only two options here. Future studies may benefit from considering other techniques.

4.4.1 Bayes Factor

The Bayes factor K is a simple metric, corresponding to

$$K_{12} = K(M1, M2) = \frac{p(D|M1)}{p(D|M2)}, \quad (4.13)$$

the ratio of the Bayesian evidence values of each of the respective models. Gelman states [93] that Bayes factor model comparison is helpful in situations in which each of the alternative models makes scientific sense, and there are no obvious intermediate models in between. The BF is considered somewhat sensitive to the choice of prior, and is generally considered challenging to explicitly compute, as the evidence is often not computed at all in Bayesian inference problems. When using a sampling algorithm such as Nested Sampling, which gives the evidence value as an output, it becomes easy to execute such an analysis.

A model with a higher BF is more likely to be the more correct model. The degree to which the BF ratio is larger than unity determines the degree of significance to which that model is preferred, and there are various estimates to estimate the significance. A BF of 2 is said to be of minor significance and a BF of 10 is considered relatively definitive [94], however an exact numerical cutoff is not recommended.

As an example of a model comparison test, the same detector was trained on two sets of four waveforms, using two different sets of electronics models. The only difference between the two models is that model A contained an antialiasing filter, while model B did not. The results of this model comparison can be seen in Table 4.2. Clearly, the model set described by model B was vastly preferred, despite the fact that model A

	WF Set 1	WF Set 2
K_{AB}	0.068	0.054
K_{BA}	14.7	18.5

Table 4.2: Example of Bayes Factor comparison between models, for two sets of models, with two sets of waveforms indicating a strong preference for model B.

contained more free parameters. Proceeding in this fashion, each proposed model can be tested to verify its applicability to the prescribed data.

4.4.2 Cross-Validation

Cross validation segments the data into a training set and a validation set, performing the training analysis on the training set and evaluating the performance on the validation set. This helps to avoid overfitting, by rewarding the ability to perform well on previously unseen data. Generally, this is a computationally expensive process, as there are many different ways to partition and segment the data. This technique is powerful, especially when limited training data is available, but has not yet been used for this study.

4.5 Conclusions

As the fitting techniques we are using give the evidence Z as an output, we find the use of the Bayes factor to be most convenient in comparing model performance. We will apply the statistical techniques described in this chapter to waveform data. The results of this are described in the following chapter.

CHAPTER 5

Performance and Results

In this chapter, we will discuss the output of training and waveform fits, and the performance and stability of these fits. In addition, we discuss the validation techniques used to verify fit performance, and other future methods which could further validate them.

5.1 Convergence Metrics

There are several standard metrics used to evaluate the convergence of individual fits. For example, one popular technique advocated by Gelman [93] is to check MCMC chain convergence using \hat{R} , an autocorrelation metric.

The direct output of a nested sampling fit is a chain of samples with weights to apply for each level and a marginal likelihood – a posterior is accessible after processing the chain. Within the Diffusive Nested Sampling paradigm, the most basic check is to see that each level has been explored, and that adequate mixing of levels has occurred following this, as indicated in Figure 5.1. Here, we have obtained the desired number of levels, and the particles have begun exploring the full mixture of levels. Usually, the fit should be allowed to continue fully mixing for as long as feasible.

In addition to this simple exploration metric, the level weight may be used to determine convergence. A peak of weight as a function of remaining prior mass X indicates that the majority of the posterior mass may have been found. An example of this is shown in Figure 5.2. Here, the top figure shows the log-likelihood as a function of the enclosed prior mass; this indicates that past a certain point, the log-likelihood no longer appears to increase. The bottom figure of Figure 5.2 indicates a peak in posterior weight as a function of the remaining enclosed prior mass. The peak is expected in every successful run, as it indicates that there is not a continuously increasing amount of posterior weight during continued exploration. The fit should be run for long enough that no new posterior weight peaks are discovered at decreasing $\log(X)$. If an insufficient number of contour levels have been specified, the fit may fail to form a full peak, leaving the posterior weight




Figure 5.1: Beginning of level exploration for a fit using 105 particles per each of 4 waveforms over 750 levels. Each level is successively made, compressing the remaining prior mass in each subsequent level, and after all levels have been created, the particles diffuse into exploring through all existing levels.




Figure 5.2: Top figure shows the log-likelihood as a function of the enclosed prior mass. Bottom figure shows the weights of the saved particles, shown by their enclosed prior mass values; particles to the left of the peak are not expected to contribute to the evidence significantly.




Figure 5.3: Fit result of impurity at each end of the detector for detector P42575A in dataset 1. The right plots are chains of the posterior values with the mean indicated, plotted as a function of sample number. The left plots are KDEs of the indicated chain, with the median and mean estimates indicated.

all at unspecified higher likelihood levels, with the enclosed prior mass all lying in just the innermost contours. This can be avoided by using a larger number of levels, however, using too many levels is computationally wasteful and time consuming.

5.2 Training Performance

As explained in the previous chapter, the output of a training fit is a set of model parameters unique to a particular detector as defined by the set of models use for training. These usually include the impurity gradient and charge trapping parameters, as well as the various electronics models to describe the preamplifier and front-end electronics. An example of the training result for one model in one detector is should in Figure 5.3. Here, the two parameters are the impurity concentration at each end of the detector in units of 10^{10} impurities per cubic centimeter.

In order to reduce the amount of bias in the estimate, a Kernel Density Estimator (KDE) with a gaussian kernel is used as an alternative to histogramming. This helps to eliminate the binning effects that can arise in reading histograms.

In order to establish the reproducibility of these model parameters for a given detector, we have fit the same detector using independent sets of data waveforms. If the fit is indeed finding a true good value, it is expected to produce the same set of parameters back, regardless of which particular set of waveforms are used for training. An example of this for the impurity gradient parameters is shown in Figure 5.4. The spread in parameter values when fit with multiple waveform sets can be expected to represent an uncertainty in the fit result values.

In the case of this detector, P42575A, the impurity gradient reported by the manufacturer was -0.971×10^{10} at the seed end of the crystal and -0.7×10^{10} at the tail. The fit values estimated by comparing the three training sets shown are impurities of -0.72×10^{10} and -0.33×10^{10} , for the same values. While the values themselves differ by an offset of approximately 0.3×10^{10} , the impurity gradient is found to be consistent with the detector vendor values.

In addition, the detector depletion voltage is an experimentally determined quantity which is well-determined and calculable given a particular impurity profile. Therefore, we can use the depletion voltage as a way of evaluating the impurity profile determined by performing training fits. Using the impurity profile indicated in Figure 5.4, the depletion voltages is found to be 1802 V. This compares extremely closely to the experimentally determined 1800 V value, and is therefore an indication that the field we determine from the training is an accurate representation of the true field in the detector.

Both the natural BEGe detectors and the enriched “ORTEC” detectors have been used for training. The fits from both detector types successfully converge, even using the same basic models (aside from geometry).

5.3 Waveform Fit Performance

Using the techniques described in the previous chapter, we are able to fit waveforms from trained detectors as an automated Tier 3 data processing step. While the fits are much less computationally intensive than the training fits, they still require at least several minutes of single-core processing time per waveform, and up to several hours. In addition, it is much more challenging to verify the convergence of a large number of waveforms manually; instead, as the results are processed into a dataframe, the convergence metrics are

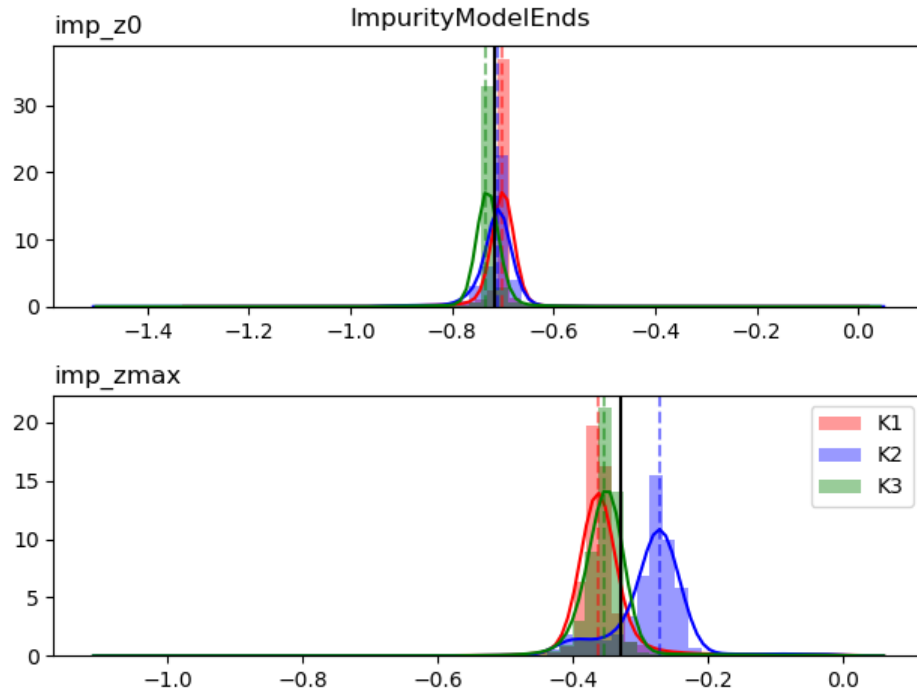


Figure 5.4: Fit result of impurity at each end of the detector for detector P42575A in dataset 1. The three different colors each represent a training fit using 4 different waveforms, and the black line is the average of the posterior maxima.




Figure 5.5: Comparison of energy parameters between fit data and trapezoidal filter estimation. The two metrics are largely in agreement, although some waveforms differ by as much as 10 keV.

added to allow for cuts to be applied. So, while the mechanism for performing this type of single waveform fit is now in place, it has not yet been done on a large-scale set of waveforms, as from a full calibration or background dataset.

Doing a simple single-point calibration at the 2615 keV photopeak, we can compare the Tier 3 fit energy to that which is found by a more conventional trapezoidal filter and spectral calibration. In Figure 5.5, the energy value that is found by the conventional analysis is plotted against the simple waveform fit energy value. Most of the spectrum agrees between these two methods, however, a small number of events differ by as much as 10 keV. Although this is a low-statistics comparison, this large difference merits additional study, as this may easily push events into or out of the ROI. Also note that the distribution of energy residuals in this figure is non-symmetric, with the traditional fit value being a slightly higher value, but with the T3 fit energies all outlying low.

In the joint plot of Figure 5.6 the scale is zoomed to just allow the 2615 photopeak. Here it is possible to compare the energies found by each method, and the energy resolutions are shown to be comparable. The resolution is indeed somewhat degraded for both methods when compared to that found in the traditional




Figure 5.6: Comparison of energy parameters between fit data and trapezoidal filter estimation, in a joint KDE plot centered over the 2615 keV photopeak.

MAJORANA analysis, most likely because the ADC nonlinearity correction is not included here. Charge-trapping effects should correctly be contained within the detector model. This is further discussed in Section 6.3.2.

Although we do not know the true values for each of the event locations within the detector, we can approximate the approximate distribution of event positions within the detector bulk. The calibration source is positioned around the outside of the array. Simple mass attenuation for a 2.5 MeV photon will attenuate the flux by approximately one half in 3.5 cm of germanium, the radius of detector P425757A [95]. Therefore, we expect that the center of the detector will have somewhat lower volumetric count rate than the rate at the radial extreme.




Figure 5.7: Fitted positions for 112 calibration events in a single detector. Note that radius value is squared to correct for the cylindrical volume element. Events show some higher density near the radial extreme of the detector and some clustering in a small region near the surface. The events which have slower drift times are found to correctly be located farther from the point contact ($r = 0, z = 0$).

When investigating waveforms fit with these training parameters, the hits were not uniformly distributed, and clustered somewhat towards a point near the outside top edge of the detector, as indicated in Figure 5.7. This seems to indicate that the found parameters are locating a local maximum, or that the convergence is not complete for some events.

The fit waveforms can be compared to the simulated distribution indicated in Figure 5.8. Here, there is slightly higher density at the radial maximum, and a relatively uniform distribution of events throughout the bulk. The rate decreases with increasing axial position (increasing z), presumably due to the source location.

5.3.1 Other Detectors and Datasets

Within this thesis, the focus is largely on data taken in the first dataset of the DEMONSTRATOR, and for a single detector. In the development of this work, training fits were performed on approximately a dozen detectors, within 4 datasets. While these fits were instructive in establishing a framework and workflow for

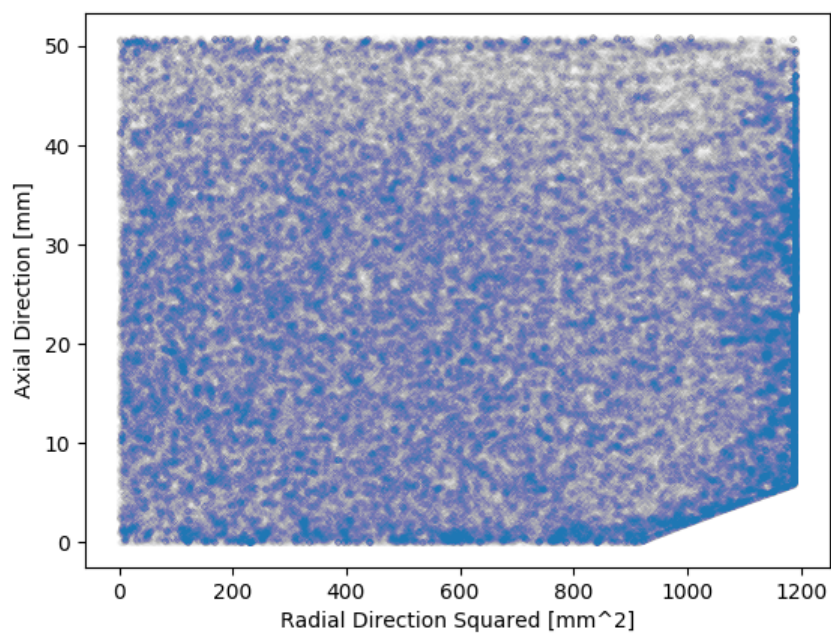


Figure 5.8: Simulated event interaction locations for 630,525 calibration source events in detector P42575A, with each blue point indicating a different event. Note that radius value is squared to correct for the cylindrical volume element. A higher color density indicates a relatively higher number of events in that location. Events show some higher density near the radial extreme of the detector and the density decreases with increasing axial position. The event distribution is otherwise nearly uniform.

performing fits, many of these fits did not reach convergence. This was later found to be resolvable using additional computing resources. Therefore, the results of these fits are not included here.

Among detectors, fits were performed among both the enriched-Ge ORTEC detectors and the natural-Ge Canberra BEGe detectors. All BEGe detectors have the same physical dimensions, and use germanium of natural isotopic abundance. Successful training fits have been performed on both styles of detectors, and the existing models appear to encompass both styles adequately.

5.3.2 Future Validation

Rather than fitting simulated data, the ideal validation technique would make use of signals of known position to deposit energy at specific points within the detector, and then verify that these positions are recovered. One way to perform this technique would be to use a scanning table to scan a highly collimated x-ray or gamma ray source along the detector. This configuration presents several challenges: it requires a different experimental configuration, which must be constructed, and it may not translate directly to the true configuration. However, it would also be very valuable as a test to validate the results of this method, as well as possibly an input to training methods with known position coordinates.

For somewhat different reasons, a campaign was initially performed within MAJORANA to characterize detectors in an azimuthal source-scanning table. This was done in an effort to identify the crystal axes *in situ* for a dark matter analysis, but ultimately required too much experimental time to be performed on each detector. A similar experimental technique may be useful for these studies.

Additionally, it may be possible to use a spectrum analyzer with a tracking generator to characterize the preamplifier response directly on a lab bench, and then translate this into the experimental configuration. This has not been explored, but would be interesting.

CHAPTER 6

Applications

In this chapter, we will discuss the applications for advanced event reconstruction with germanium detectors, including some potential future applications.

6.1 Overview

The expected outcome of an advanced fitting technique for analysis of detector waveforms is the improved characterization and understanding of individual detectors and waveforms. That is, by training algorithms to understanding each detector, we can then model the electronics and detector characteristics for that detector.

6.2 Inverse Filtering

There are applications in which it is desirable to do analysis on the waveforms with the electronics effects removed. In principle, these recovered waveforms should be a close approximation to the actual raw signal produced by the detector.

In a discrete time system, a Z-transform can be used to model filters, as described in Appendix B. In order to recover the raw waveform from a filtered one, it is necessary to apply an inverse filter¹. Such a filter inverts the transfer function of the original filter, either literally, as

$$H_{inv}[z] = \frac{1}{H[z]} \quad (6.1)$$

or effectively, by other means such as spectral deconvolution. This would then convert a high-pass filter into a low-pass, and vice versa.

¹Note that an inverse filter is distinct from an inverse transform, such as an inverse Fourier or Z-transform, which reverses the signal domain from its transformed domain back into its original domain. In the inverse filters described here, all signals are in the time domain.

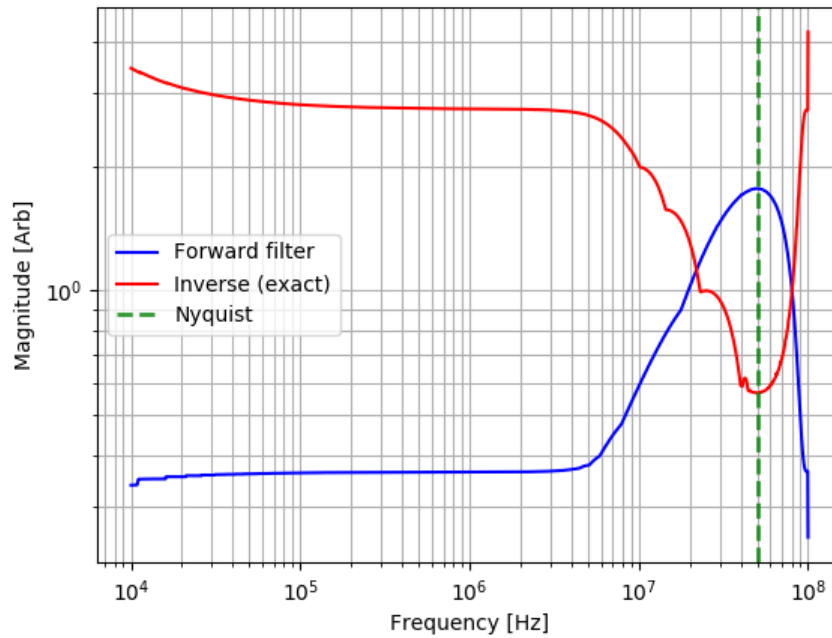


Figure 6.1: Frequency response of forward and inverse transfer functions. The forward and inverse filters are nicely matched, indicating that much of the frequency-space power can be recovered.

While in principle it is possible to apply an inverse to any filter, it is not always useful in practice. Consider a filter which entirely removes some part of the signal: applying an inverse filter to such a waveform could never hope to recover it. In most cases, a filter will just decrease the amplitude in the frequency domain of particular parts of a signal, usually by several orders of magnitude. Inverting one of these filters may add considerable noise by amplifying a signal which no longer exists.

Another way to consider this problem is by recognizing that zeros, which are roots of the numerator, become poles when inverted according to Equation 6.1. Poles outside the unit circle in the z -plane will always be unstable, giving rise to undamped oscillations. This means that a stable forward filter may have an unstable inverse filter, though if all poles and zeros lie within the unit circle, the filter is said to be *minimum phase* and invertible.

Therefore, developing inverse filters is often a compromise which cannot always be expected to reproduce a perfect raw waveform. Regardless, it is still valuable to remove as much of the filtering as possible. In the case of MAJORANA waveforms, much of the shaping can be removed by inverse filtering. An example of the forward and reverse filter frequency responses are indicated in Figure 6.1. These applications perform well,




Figure 6.2: Waveform with electronics effects removed. Note that the falling edge of the waveform is corrected to a flat top, as expected.

with the theoretical forward and inverse filters well matched to each other, as expected. The effect of the reverse filter on the waveform is indicated in Figure 6.2, where the falling edge is corrected, eliminating the obvious high-pass filter effect.

6.3 Future applications

There are several other applications which have been explored, but not yet fully demonstrated.

6.3.1 Electronics Stability Performance

The output of a training fit is a set of parameters describing the electrical characteristics of a given detector. In general, variations of the performance between detectors is believed to result from the slight differences in the components between front ends, as well as the detector's electrical characteristics. We believe the gain or other electrical properties may drift over time; this could be the result of component values shifting over time, or from the laboratory environment itself changing.

The "warm electronics" outside the cryostat (the preamplifier and digitizer) may experience some drifting performance from temperature shifts. The "cold" front end may be sensitive to changes in the power supply voltages.

Any type of instability within a run or dataset is undesirable, as it may deteriorate energy resolution and decrease our ability to extract information from individual waveforms. It is therefore of interest to try to identify when this is occurring, and possibly also to correct for it.

The anticipated technique for identifying this instability using waveform fits would be to perform training fits on the same detector at different points in time, and compare the trained values to look for trends or variation.

6.3.2 Energy estimation

The energy resolution of germanium detectors is excellent, as discussed in Chapter 3. The value of the energy resolution is essentially determined by estimating the FWHM of a peak in the ROI.

In a traditional Majorana analysis, the energy of a given single detector hit is found by a process where each waveform is corrected for ADC nonlinearity, pole-zero, and charge trapping [96]. Then the total spectrum is calibrated using a ^{228}Th calibration source. By correcting for charge trapping effects within the detector, the energy resolution can be improved by nearly a factor of 2 over the naive calibration in many detectors. In the MAJORANA DEMONSTRATOR, this is done in a combined analysis with the pole zero correction, as both effects cause exponential decaying effects.

Using the *pygama* frameworks, we can perform a spectral calibration on each detector, fitting to multiple peaks from calibration data. Here, we find the energy resolution to be around 5 keV at the 2615 keV thallium line. This corresponds to a 0.2% energy resolution, and is approximately a factor of two worse than the expected best value as found by our traditional analysis.

By performing a tier 3 fit of waveforms around the 2615 photopeak, the scaling factor (amplitude) of each waveform may be determined as an energy predictor. This fit amplitude parameter corresponds to the number of electron-hole pairs formed in the initial event, and is scaled according to the various models applied to the waveform.

In the analysis performed within this study, an ADC nonlinearity correction was not applied. While the amplitude difference of a given sample in the waveform due to nonlinearity is smaller than the noise amplitude, the integral nonlinearity over a region of the waveform can be significant. The framework has

been developed to allow this nonlinearity correction, and it is interesting to consider the possible outcomes on these fits.

6.3.3 Likelihood Cut

In a spectral analysis problem, the goal is to identify and quantify spectral features. As mentioned previously, background events act to occlude the signals of interest. In order to remove background events we apply analysis cuts, which are often calculated heuristics with a threshold applied to accept or reject particular events.

With a full detector and waveform model in place and with the best estimate for the waveform parameters found, we can then give a single total likelihood value to evaluate how well a given waveform matches the model. This can be used as a measure of how well the presented model fits to the specific waveform. A waveform which fits poorly to the model is expected to be of a distinct class of events which is not being modeled, as in a multisite event.

6.3.4 Compton Imaging

When a photon is incident on a detector (or any material) it has some probability of compton scattering, leaving some energy in the detector, and then being emitted back out with the remaining energy. If two events are registered nearly simultaneously on nearby detectors and their energies sum to a known photopeak value, then the measured energy corresponds to a particular angular direction (cone) which the photon must have entered from. Compton imaging uses multiple detectors (or segments of detectors) to do this reconstruction to backproject the possible location of incident photons, indicating the likely source location.

Using the techniques described here, it is in principle possible to identify for a single event its coordinate angle ϕ within the detector. Due to the 8-fold symmetry of the axes, it is only possible to place events within a single (unknown) octant of the detector volume. Within the MAJORANA DEMONSTRATOR this is of limited utility because the crystal axes of individual detectors are unknown; therefore, any identification of an event's axial position within the detector has no corresponding direction within which a lab-frame direction can be specified. It may be possible to “calibrate” the axis orientation by investigating the coincident compton-scattered event waveforms, but this has not been attempted. For this reason, we consider this a low-value application for this particular detector system, although the application is of general interest.

6.3.5 Other future applications

The ultimate capability of this waveform fitting technique would be to identify particular event signatures for individual events. These events might include the obvious double-beta decay, but also identifying such events as Lightly Ionizing Particles (LIPs) or muons, which are largely vertical-going and should deposit a line of ionized charge in the detector.

In addition to this, we expect that by using a finite mixing model, it would be straightforward to fit multisite and pileup events, recovering the parameters from each contributing event. Similarly, it may also be possible to guess from the unsaturated parts of a saturated waveform how large the final amplitude was going to get; this would be of value if it could extend the dynamic range of our data acquisition system.

Currently, the digitizers used by the DEMONSTRATOR use a simple trapezoidal trigger to trigger the readout. An interesting idea to consider is using the training parameters as part of a triggering algorithm wherein we could potentially even estimate some of the waveform parameters (or a proxy, such as drift time) to make trigger decisions.

CHAPTER 7

Conclusions

Within this document, a technique has been presented for extracting information from detector signals. Using a set of detailed models contributing to p-type point contact germanium detector signal formation, we are able to evaluate the validity of proposed model components using a flexible Bayesian framework. This framework also allows us to determine model parameters to describe individual detectors and waveforms. Outside the scope of this work, the author believes that no groups are using such a technique now for routine analysis of germanium detector waveforms.

The described framework is organized primarily into two software packages, *waffle* and *pygama*. Contributions to *pygama* allowed digitized waveforms from the MAJORANA DEMONSTRATOR detectors to be analyzed in a python environment convenient for this study but also of general value. The *waffle* framework was extended for this work to allow for a third tier of data processing to efficiently perform waveform fits with only a few user commands, storing the fit results in a format parallel to the existing data parameters.

7.0.1 Results

- Using Dnest4 it is possible to explore a complicated parameter space, although with computational expense which is somewhat limiting
- Using a detailed detector model, it is possible to recover some detector and electronics parameter values which are reasonable.
- We have introduced a more flexible model for the preamplifier electronics, which should be able to match any observed features.
- We have contributed to the fitting framework to improve the single waveform fit pipeline, as well as improving convergence metrics and fit postprocessing.

- When fitting back the event energy scale parameter, some events deviate significantly from the expected value. This effect merits additional study.
- Additional fitting and validation is required to fully understand where deficiencies lie.

7.1 Extensibility

The advances described in this work are specifically oriented towards the improvement of models for P-type Point Contact germanium detectors, as used by a neutrinoless double-beta decay experiment. Despite the intent of this specific use case, the ideas presented are generic enough to be used with other detector systems or waveform analysis.

In order to develop for another application, a generic physics model is needed for creating a waveform. Further, there must be a way to process and organize the raw waveform data, making it available to the fitting framework. Using these, then the fitting software must describe the prior estimates and distributions, and a means of evaluating the likelihood for a given parameter proposal.

While describing a signal model in this way is non-trivial, doing so may be of particular value in those cases where the model is high-dimensional and understanding the underlying parameters is of value.

7.2 Future Work

The work outlined in this thesis leaves several additional avenues for future exploration.

7.2.1 Near Term Work

The first steps which should be considered are:

Add multisampled waveforms This is nearly complete in *pygama*. Adding this will allow for data from all datasets to be easily processed.

Process data with ADC nonlinearity correction This will remove distortion effects from the digitization which we know contribute to waveforms, improving the fit.

Train multiple detectors on each dataset This will allow us to identify trends in model parameters and build confidence in our model set.

Choose a model set By trying multiple sets of detector and electronics models and picking the best fitting one, we can identify the single set of models which can best describe the true system. This allows us to perform future fits with more confidence that the parameters should return meaningful values.

7.2.2 Further Validation

This analysis is largely a passive experimental effort based on the data which has been already been taken by the DEMONSTRATOR. Several experiments could be designed which would aid in this type of analysis.

As described in Section 5.3.2, the first experiments that would be interesting to try would be experimentally validating fit performance using a highly collimated gamma ray source to scan different positions through the detector, and fit back waveforms with at a position coordinate already known (either r or z would be known, and possibly the azimuthal angle as well). It may be necessary to do other measurements to absolutely determine the azimuthal angle, such as by x-ray diffraction or polarized raman spectroscopy.

Also, direct measurement of the preamplifier transfer function could be useful to validate or estimate the front end properties. Operating a preamplifier on a bench top, we could more carefully evaluate its response without the complicating effects of cables, flanges, or multiple temperatures.

Additionally, it would be interesting to operate detectors using different digitizers and preamplifiers. The *pygama* framework is flexible enough that by writing a decoder for the new digitizer, the waveform objects should be directly comparable to those made for the Gretina digitizer currently in use. Additionally, by testing different front ends and preamplifiers, we may be able to more readily isolate detector effects.

Within the DEMONSTRATOR it may also be possible to perform analysis using both the high- and low-gain channels together, for events of the appropriate energy windows. This may help to avoid channel dependent nonlinearity and noise effects. Doing so would probably require a new software class of waveform object which can hold two different “views” of the same event.

7.2.3 Model Improvement

7.2.3.1 Improve the impurity gradient model

As described in Section 3.2.0.2, the impurity gradient is a crucial part of understanding the field, and therefore signal. The model used in this work is a simple axial linear gradient model, however, it is known that this is not a complete description [97]. It may be valuable to introduce additional parameterization in

order to use an exponential gradient, as well as introducing a radial component to the impurity [59]. The effect is expected to be small, but within reach without great effort. It may be also possible to work with the detector vendor to obtain a better understanding here.

7.2.3.2 Improve Velocity Model

The hole mobility parameters are allowed to float in the fits described here, introducing additional computation and uncertainty. A measurement of the charge carrier mobility in isotopically enriched ^{76}Ge at liquid nitrogen temperatures would allow these parameters to be frozen at or held near their true values. An experimental technique that allowed measurement of mobility for individual detectors as a characterization step would be of even more value.

7.2.4 Extend fits

After deciding on a final set of models to use for fitting, the final goal will be to perform training fits on all detectors in every dataset. This is a near-term analysis goal for the DEMONSTRATOR collaboration.

7.2.5 Improve Computational Performance

The fits performed here are computationally expensive. A single training fit using only 4 waveforms has a typical computation of approximately 3 days using approximately 500 Xeon Broadwell-EP CPU cores. The training fits were performed by giving each DNest particle it's own core, with the same number of particles designated for each waveform. Fits performed for longer would be interesting to ensure convergence, but are probably not necessary. The single waveform fits were generally allowed to run on a single or small number of cores for several hours.

In order to improve the computational performance, there are two main directions to work from. First, we can try to reduce the needed computation by simplifying the model. For example, if the velocity model parameters and impurity profile were known more precisely from other means, it may not be necessary to float them in a fit, and their values could be treated as exact. Additionally, it may be possible by further detailed study of the preamplifier circuit to reduce the electronics model to a simplified yet complete version. Reduction of parameters in this way reduces the dimensionality of the fit.

The other way to improve the computational performance is by speeding up the computations that are currently performed. This may mean porting the existing code to run on a GPU, improving the calculation algorithm, or using other sampling algorithms.

Additional parallelization may be a possible method to improve fit speed, however, we are already at the practical limit of current computing resources.

7.3 Conclusions

Physics is a results-driven field, with a continuous need for modern analysis solutions to increasingly challenging experimental problems. Neutrinoless double-beta decay experiments must work to improve their understanding of individual events as a mechanism for further reducing experimental backgrounds; low-background experimental design alone can not conceivably reduce backgrounds by several additional orders of magnitude at this time. The techniques described in this thesis are a step towards improving the analysis of germanium detector experiments, with particular applications for neutrinoless double-beta decay using germanium detectors. These techniques will soon be refined enough to contribute to limit-setting analysis, identify spurious backgrounds, and inform future detector design.

APPENDIX A

IMPURITY CONCENTRATION

This appendix serves as a simple manipulation of units used when discussing the density of impurities in crystals.

The density of germanium is about 5.323 grams per cubic centimeter. The atomic mass of natural germanium is 72.630 grams per mole.

This gives a molar density of

$$\frac{5.323 \text{ g/cc}}{72.6 \text{ g/mol}} = 0.0733 \text{ mol/cc.} \quad (\text{A.1})$$

for germanium.

A single mole is 6.02×10^{23} atoms. This then gives $(0.0733 \text{ mol/cc})(6.02 \times 10^{23} \text{ atoms/mol})$, or 4.4×10^{22} atoms per cubic centimeter.

The impurity concentration of High Purity Germanium (HPGe) detectors is frequently stated as a concentration impurities per cubic centimeter. Detector grade HPGe can often have an impurity concentration lower than 10^{10} impurities per cubic centimeter. Using the value stated above, we find that this corresponds to a net impurity concentration

$$[\text{Impurity}] = \frac{10^{10} \text{ impurities/cc}}{4.4 \times 10^{22} \text{ atoms/cc}} = 2.3 \times 10^{-12}. \quad (\text{A.2})$$

That is, the impurity is at a level of about 2 parts per trillion, or 0.002 parts per billion! Indeed, HPGe is one of the most chemically pure materials ever fabricated.

APPENDIX B

Z-TRANSFORMS

This appendix serves as an introductory overview of the z-transform as it is used in signal processing applications.

There is no shortage of resources to learn about digital signal processing. The methods of [98] are clear; Oppenheim is the author of a classic text [99], which is considered a standard reference. The author is partial to the free Micromodeler DSP online software [100], which provides a convenient interface to explore the results of any arrangement of poles and zeros on digital system.

B.1 The Z-Transform

In general, the role of a filter is to enhance or diminish a signal at particular frequencies or ranges of frequencies. Filters may be as simple as an RC high- or low-pass filter, or may be more complicated filters such as the Chebyshev or Elliptic filters. In order to model a digital filter, we make use of the Z-transform. The Z-transform is a discrete-time equivalent to the Laplace transform, and is used extensively in digital signal processing for modeling and applying filters.

The z-transform of a sequence $x[n]$ is defined as:

$$X(z) = \sum_{n=-\infty}^{\infty} x[n]z^{-n} \quad (\text{B.1})$$

A Z-transform describes a filter's transfer function as a quotient of polynomials, usually in negative powers of z, as

$$H[z] = \frac{B_0 + B_1z^{-1} + B_2z^{-2} + \dots + B_Nz^{-N}}{A_0 + A_1z^{-1} + A_2z^{-2} + \dots + A_Mz^{-M}} \quad (\text{B.2})$$

Adding confusion to the field, some disciplines (notably geophysics) choose to use positive powers of z, with the appropriate transformations to yield the same utility. This transfer function can then be applied directly to a signal $X[z]$ to give a transformed signal $Y[z]$, as

$$Y[z] = H[z]X[z] \quad (\text{B.3})$$

Roots of the numerator describe zeros, and roots of the denominator describe poles. A system is stable when its poles are all contained within the unit circle, and is invertible if its zeros are within the unit circle.

The z-transform describes a complex phasor spinning about the origin with the system's sampling frequency. A given transform describes filters only relative to the sampling frequency, and has no knowledge of particular rolloff frequencies except as a fraction of the sampling frequency. In this way, the $\phi = 0$ direction corresponds to a DC signal, and $\phi = \frac{\pi}{2}$ corresponds to one fourth of the sampling frequency. As $\phi = \pi$ corresponds to the Nyquist frequency, all complex-valued poles and zeros will necessarily be mirrored about the real axis.

An equivalent way to consider the z-transform is as a linear difference equation [98, pg. 110]. In this way, a transfer function describes a difference equation which can feedback and feedforward scaled sample values to accomplish a particular filter. The powers in the numerator add to the value according to their respective coefficients, and the powers in the denominator will subtract from it, as:

$$y[n] = \sum_{k=0}^M b_k x[n-k] - \sum_{k=1}^N a_k y[n-k] \quad (\text{B.4})$$

Using the first numerator coefficient a_0 as a normalizing coefficient, this could equally be described as:

$$a_0 * y_n = b_0 * x_n + b_1 * x_{n-1} + \dots + b_M * x_{n-M} \quad (\text{B.5})$$

$$-a_1 * y_{n-1} - \dots - a_N * y_{n-N} \quad (\text{B.6})$$

It is this convolutional filter that is used for applying a particular z-transform to a signal, as by the `scipy.signal.lfilter` method.

APPENDIX C

DATA ACQUISITION FOR THE MAJORANA DEMONSTRATOR

This appendix serves as an overview of the data acquisition system used by the Demonstrator.

C.1 ORCA

ORCA provides a framework for controlling hardware objects and acquiring data from them. The design is modular and object-oriented, written in Objective-C for the Mac OS using the model-view-controller design pattern. The application was targeted to Mac OS using Objective-C because, at the time of its development, there were no simple options to support across platforms, and that option provided the best experience for the primary target user. Further, because it relies on the Apple Cocoa framework, it continually has the appearance of modern applications in the current operating system with minimal development effort.

ORCA is designed as a graphical interface with many drag-and-drop features, allowing users to change any experimental parameters and topologies without recompiling the application.

C.1.1 Objects in ORCA

Each hardware object in ORCA consists of three classes, a model, view, and controller, each separated into its own files (this is a standard MVC design pattern). The *model* contains all the necessary code to interact with and control the device. The *view* is an Apple NIB or XIB file, which is a graphical description of the object as it will be interacted with by a user. The *controller* is a class which binds user actions in the *view* to methods in the *model*. The bulk of the work that goes into adding an object into the ORCA hardware catalog is in the model code, as the model must contain code to control every option that a device is capable of, and some complex objects, such as digitizers, have hundreds of registers and channels to keep track of and strict sequences of events which must occur to allow for data readout. In general, ORCA models are written to encompass all features of each hardware device, allowing the user to later access the relevant features.

C.2 The MAJORANA DEMONSTRATOR

The DAQ system used by the MAJORANA DEMONSTRATOR is graphically demonstrated in Figure C.1. Digitizers in VME crates are time-synchronized using custom “trigger cards”. Using a serial interface on




Figure C.1: The Data Acquisition system for the Majorana Demonstrator.

the front panel of the digitizers, the baseline voltages may be read out from the controller cards of the preamplifiers. Data is read from the digitizer cards over a VME backplane, where it is buffered to a Single Board Computer (SBC) acting as the controller in each VME crate. The High Voltage system is controlled by MPOD-based high voltage cards. Both the VME SBC cards and the MPOD controllers are accessed over a private network connected to a DAQ readout computer locally underground.

A separate veto system monitors acrylic veto panels by PMT signals on QDCs for the appropriate logic condition to indicate a veto is necessary. Using a VME latching scaler, the clock signal from the trigger card is counted and the timestamp is recorded at times which indicate a veto is necessary. Vetoed digitizer events are still recorded, and are removed in analysis by comparing the event time to the recorded veto time.

A separate Slow Controls computer also runs ORCA and manages the secondary DAQ systems, such as environmental monitoring. These processes include clean room particle count, radon, humidity, and temperatures, the glovebox nitrogen purges, as well as tracking and controlling the liquid nitrogen levels in each module. Each module's cryovacuum system is controlled by an independent computer running ORCA. The status of slow controls systems are recorded in a CouchDB database which is synchronized to a server at UNC, where a web server displays the current status of all systems.

BIBLIOGRAPHY

- [1] Wikimedia Commons. Nuclear decay mode as a function of neutron and proton number., 2017. URL <https://commons.wikimedia.org/wiki/File:DecayModeNuDat2.png>. File: DecayModeNuDat2.png. xi, 2
- [2] Brookhaven National Laboratory National Nuclear Data Center. Nudat (nuclear structure and decay data), March 18, 2008 2008. xi, 4
- [3] Laurie M. Brown. The idea of the neutrino. *Physics Today*, 31(9):23–28, Sep 1978. doi: 10.1063/1.2995181. URL <http://dx.doi.org/10.1063/1.2995181>. 3
- [4] David Griffiths. *Introduction To Elementary Particles*. Wiley-VCH Verlag GmbH, Weinheim, 1 edition, 2004. ISBN 978-0-471-60386-3. 3
- [5] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: a confirmation. *Science*, 124(3212):103–104, 1956. ISSN 0036-8075. doi: 10.1126/science.124.3212.103. URL <http://science.sciencemag.org/content/124/3212/103>. 5
- [6] John N. Bahcall. Solar neutrinos. i. theoretical. *Physical Review Letters*, 12(11):300–302, Mar 1964. doi: 10.1103/physrevlett.12.300. URL <http://dx.doi.org/10.1103/PhysRevLett.12.300>. 6
- [7] Raymond Davis. Solar neutrinos. ii. experimental. *Physical Review Letters*, 12(11):303–305, Mar 1964. doi: 10.1103/physrevlett.12.303. URL <http://dx.doi.org/10.1103/PhysRevLett.12.303>. 6, 15
- [8] Y. Fukuda, T. Hayakawa, E. Ichihara, K. Inoue, K. Ishihara, H. Ishino, Y. Itow, T. Kajita, J. Kameda, S. Kasuga, and et al. Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 81(8):1562–1567, Aug 1998. doi: 10.1103/physrevlett.81.1562. URL <http://dx.doi.org/10.1103/PhysRevLett.81.1562>. 6
- [9] LA Science. Los alamos science: Celebrating the neutrino, 1997. URL <https://la-science.lanl.gov/lascience25.shtml>. 6
- [10] B. Pontecorvo. Neutrino experiments and the problem of conservation of leptonic charge. *Sov. Phys. JETP*, 26:984–988, 1968. URL http://www.jetp.ac.ru/files/pontecorvo1968_en.pdf. Zh. Eksp. Teor. Fiz.53,1717(1967). 6
- [11] M. Tanabashi, K. Hagiwara, K. Hikasa, K. Nakamura, Y. Sumino, F. Takahashi, J. Tanaka, K. Agashe, G. Aielli, C. Amsler, and et al. Review of particle physics. *Physical Review D*, 98(3), Aug 2018. doi: 10.1103/physrevd.98.030001. URL <http://dx.doi.org/10.1103/PhysRevD.98.030001>. 6
- [12] Planck Collaboration. Planck 2015 results - xiii. cosmological parameters. *A&A*, 594:A13, 2016. doi: 10.1051/0004-6361/201525830. URL <https://doi.org/10.1051/0004-6361/201525830>. 7
- [13] J. F. Wilkerson and R. G. H. Robertson. *Direct Measurements of Neutrino Mass*, pages 39–64. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001. ISBN 978-3-662-04597-8. doi: 10.1007/978-3-662-04597-8_3. URL http://dx.doi.org/10.1007/978-3-662-04597-8_3. 7

- [14] A. Osipowicz et al. KATRIN: A Next generation tritium beta decay experiment with sub-eV sensitivity for the electron neutrino mass. Letter of intent. *arXiv*, 2001. 7
- [15] Benjamin Monreal and Joseph A. Formaggio. Relativistic cyclotron radiation detection of tritium decay electrons as a new technique for measuring the neutrino mass. *Physical Review D*, 80(5), Sep 2009. doi: 10.1103/physrevd.80.051301. URL <http://dx.doi.org/10.1103/PhysRevD.80.051301>. 7
- [16] Gary Steigman. Observational tests of antimatter cosmologies. *Annual Review of Astronomy and Astrophysics*, 14(1):339–372, Sep 1976. doi: 10.1146/annurev.aa.14.090176.002011. URL <http://dx.doi.org/10.1146/annurev.aa.14.090176.002011>. 7
- [17] Vladimir I Tretyak and Yuri G Zdesenko. Tables of double beta decay data:an update. *Atomic Data and Nuclear Data Tables*, 80(1):83–116, Jan 2002. doi: 10.1006/adnd.2001.0873. URL <http://dx.doi.org/10.1006/adnd.2001.0873>. 8
- [18] B. Pritychenko. Systematics of evaluated half-lives of double-beta decay. *Nuclear Data Sheets*, 120: 102–105, Jun 2014. doi: 10.1016/j.nds.2014.07.018. URL <http://dx.doi.org/10.1016/j.nds.2014.07.018>. 8
- [19] Boris Kayser. Are neutrinos their own antiparticles? *Journal of Physics: Conference Series*, 173(1): 012013, 2009. URL <http://stacks.iop.org/1742-6596/173/i=1/a=012013>. 9
- [20] A. Gando, Y. Gando, T. Hachiya, A. Hayashi, S. Hayashida, H. Ikeda, K. Inoue, K. Ishidoshiro, Y. Karino, M. Koga, and et al. Search for majorana neutrinos near the inverted mass hierarchy region with kamland-zen. *Physical Review Letters*, 117(8), Aug 2016. doi: 10.1103/physrevlett.117.082503. URL <http://dx.doi.org/10.1103/PhysRevLett.117.082503>. 9
- [21] J. Kotila and F. Iachello. Phase-space factors for double- β decay. *Phys. Rev. C*, 85:034316, Mar 2012. doi: 10.1103/PhysRevC.85.034316. URL <https://link.aps.org/doi/10.1103/PhysRevC.85.034316>. 10
- [22] F. Cappuzzello, C. Agodi, M. Cavallaro, D. Carbone, S. Tudisco, D. Lo Presti, J. R. B. Oliveira, P. Finocchiaro, M. Colonna, D. Rifuggiato, and et al. The NUMEN project: NUClear Matrix Elements for Neutrinoless double beta decay. *The European Physical Journal A*, 54(5), May 2018. doi: 10.1140/epja/i2018-12509-3. URL <http://dx.doi.org/10.1140/epja/i2018-12509-3>. 11
- [23] Jonathan Engel and Javier Menéndez. Status and future of nuclear matrix elements for neutrinoless double-beta decay: a review. *Reports on Progress in Physics*, 80(4):46301, Mar 2017. doi: 10.1088/1361-6633/aa5bc5. URL <http://dx.doi.org/10.1088/1361-6633/aa5bc5>. 11
- [24] Michelle J. Dolinski, Alan W. P. Poon, and Werner Rodejohann. Neutrinoless Double-Beta Decay: Status and Prospects. *arXiv*, 2019. 11, 13
- [25] Stefano DellOro. Neutrinoless double beta decay: expectations and uncertainties. *Nuclear and Particle Physics Proceedings*, 265-266:31–33, Aug 2015. doi: 10.1016/j.nuclphysbps.2015.06.008. URL <http://dx.doi.org/10.1016/j.nuclphysbps.2015.06.008>. 11
- [26] Fedor Šimkovic. Neutrinoless double beta decay: Theoretical challenges. *Nuclear and Particle Physics Proceedings*, 265-266:19–24, Aug 2015. doi: 10.1016/j.nuclphysbps.2015.06.005. URL <http://dx.doi.org/10.1016/j.nuclphysbps.2015.06.005>. 11

- [27] Felix Boehm and Petr Vogel. *Physics of Massive Neutrinos*. Cambridge University Press, 2 edition, 1992. ISBN 0-521-41824-0. 11, 13
- [28] Kai Zuber. *Neutrino Physics*. IOP Publishing, 2004. ISBN 0-7503-0750-1. 11, 13
- [29] Giovanni Benato. Effective majorana mass and neutrinoless double beta decay. *The European Physical Journal C*, 75(11), Nov 2015. doi: 10.1140/epjc/s10052-015-3802-1. URL <http://dx.doi.org/10.1140/epjc/s10052-015-3802-1>. 12
- [30] Frank Deppisch, Chris Jackson, Irina Nasteva, and Stefan S oldner-Rembold. Probing the mechanism of neutrinoless double beta decay with supernemo. *Progress in Particle and Nuclear Physics*, 64(2): 278–280, Apr 2010. doi: 10.1016/j.pnpnp.2009.12.028. URL <http://dx.doi.org/10.1016/j.pnpnp.2009.12.028>. 12
- [31] Frank F. Deppisch, Lukas Graf, Julia Harz, and Wei-Chih Huang. Neutrinoless double beta decay and the baryon asymmetry of the universe. *Physical Review D*, 98(5), Sep 2018. doi: 10.1103/physrevd.98.055029. URL <http://dx.doi.org/10.1103/PhysRevD.98.055029>. 12
- [32] J. Schechter and J. W. F. Valle. Neutrinoless double- β decay in $su(2) \times u(1)$ theories. *Physical Review D*, 25(11):2951–2954, Jun 1982. doi: 10.1103/physrevd.25.2951. URL <http://dx.doi.org/10.1103/PhysRevD.25.2951>. 12
- [33] A. Baha Balantekin and Boris Kayser. On the properties of neutrinos. *Annual Review of Nuclear and Particle Science*, 68(1):313–338, Oct 2018. doi: 10.1146/annurev-nucl-101916-123044. URL <http://dx.doi.org/10.1146/annurev-nucl-101916-123044>. 12
- [34] Steven R. Elliott and Petr Vogel. Double beta decay. *Annual Review of Nuclear and Particle Science*, 52(1):115–151, Dec 2002. doi: 10.1146/annurev.nucl.52.050102.090641. URL <http://dx.doi.org/10.1146/annurev.nucl.52.050102.090641>. 13
- [35] J.J Gómez-Cadenas, J Martín-Albo, M Sorel, P Ferrario, F Monrabal, J Muñoz, P Novella, and A Poves. Sense and sensitivity of double beta decay experiments. *Journal of Cosmology and Astroparticle Physics*, 2011(6):007–007, Jun 2011. doi: 10.1088/1475-7516/2011/06/007. URL <http://dx.doi.org/10.1088/1475-7516/2011/06/007>. 13
- [36] Matteo Agostini, Giovanni Benato, and Jason A. Detwiler. Discovery probability of next-generation neutrinoless double- β decay experiments. *Physical Review D*, 96(5), Sep 2017. doi: 10.1103/physrevd.96.053001. URL <http://dx.doi.org/10.1103/PhysRevD.96.053001>. 13
- [37] F. Alessandria, R. Ardito, D. R. Artusa, F. T. Avignone III, O. Azzolini, M. Balata, T. I. Banks, and et. al. Sensitivity and discovery potential of cuore to neutrinoless double-beta decay. *arXiv*, 2013. URL <https://arxiv.org/abs/1109.0494>. 13
- [38] R. G. H. Robertson. Empirical survey of neutrinoless double beta decay matrix elements. *Modern Physics Letters A*, 28(8):1350021, Mar 2013. doi: 10.1142/s0217732313500211. URL <http://dx.doi.org/10.1142/S0217732313500211>. 13
- [39] Reyco Henning. Current status of neutrinoless double-beta decay searches. *Reviews in Physics*, 1: 29–35, Nov 2016. doi: 10.1016/j.revip.2016.03.001. URL <http://dx.doi.org/10.1016/j.revip.2016.03.001>. 13

- [40] N. Abgrall, I.J. Arnquist, III Avignone, F.T., H.O. Back, A.S. Barabash, F.E. Bertrand, M. Boswell, A.W. Bradley, V. Brudanin, M. Busch, and et al. The MAJORANA DEMONSTRATOR radioassay program. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 828:22–36, Aug 2016. doi: 10.1016/j.nima.2016.04.070. URL <http://dx.doi.org/10.1016/j.nima.2016.04.070>. 14
- [41] C. Cuesta, N. Abgrall, E. Aguayo, III Avignone, F.T., A.S. Barabash, F.E. Bertrand, M. Boswell, V. Brudanin, M. Busch, D. Byram, and et al. Background model for the majorana demonstrator. *Physics Procedia*, 61:821–827, 2015. doi: 10.1016/j.phpro.2015.06.001. URL <http://dx.doi.org/10.1016/j.phpro.2015.06.001>. 14
- [42] N. Abgrall, E. Aguayo, III Avignone, F.T., A.S. Barabash, F.E. Bertrand, A.W. Bradley, V. Brudanin, M. Busch, M. Buuck, D. Byram, and et al. Muon flux measurements at the davis campus of the sanford underground research facility with the majorana demonstrator veto system. *Astroparticle Physics*, 93:70–75, Jul 2017. doi: 10.1016/j.astropartphys.2017.01.013. URL <http://dx.doi.org/10.1016/j.astropartphys.2017.01.013>. 15
- [43] C.E. Aalseth, N. Abgrall, E. Aguayo, S.I. Alvis, M. Amman, I.J. Arnquist, F.T. Avignone, H.O. Back, A.S. Barabash, P.S. Barbeau, and et al. Search for neutrinoless double- β decay in ^{76}Ge with the majorana demonstrator. *Physical Review Letters*, 120(13), Mar 2018. doi: 10.1103/physrevlett.120.132502. URL <http://dx.doi.org/10.1103/PhysRevLett.120.132502>. 16
- [44] S.I. Alvis, I.J. Arnquist, F.T. Avignone III, A.S. Barabash, C.J. Barton, V. Basu, F.E. Bertrand, B. Bos, M. Busch, and et. al. A Search for Neutrinoless Double-Beta Decay in ^{76}Ge with 26 kg-yr of Exposure from the MAJORANA DEMONSTRATOR . *arXiv*, 2019. 16
- [45] N. Abgrall, A. Abramov, N. Abrosimov, I. Abt, M. Agostini, M. Agartioglu, A. Ajjaq, S. I. Alvis, III Avignone, F. T., X. Bai, and et al. The large enriched germanium experiment for neutrinoless double beta decay (legend). *AIP Conference Proceedings*, 1894(1):020027, 2017. doi: 10.1063/1.5007652. URL <https://aip.scitation.org/doi/abs/10.1063/1.5007652>. 16
- [46] Paul M. Koenraad and Michael E. Flatté. Single dopants in semiconductors. *Nature Materials*, 10(2): 91–100, Jan 2011. doi: 10.1038/nmat2940. URL <http://dx.doi.org/10.1038/nmat2940>. 18
- [47] Wikimedia Commons. Solid state electronic band structure, 2017. URL https://commons.wikimedia.org/wiki/File:Solid_state_electronic_band_structure.svg. File:Solid state electronic band structure.svg. xi, 19
- [48] Ioffe. Band structure and carrier concentration, 2019. URL <http://www.ioffe.ru/SVA/NSM/Semicond/Ge/bandstr.html>. ”Website <http://www.ioffe.ru/SVA/NSM/Semicond/Ge/bandstr.html>”. xi, 20
- [49] Frank Herman. The electronic energy band structure of silicon and germanium. *Proceedings of the IRE*, 43(12):1703–1732, 1955. doi: 10.1109/jrproc.1955.278039. URL <http://dx.doi.org/10.1109/JRPROC.1955.278039>. 20
- [50] Glenn F. Knoll. *Radiation Detection and Measurement*. John Wiley & Sons, Inc, Hoboken, NJ, 4th edition, 2010. 21, 24, 28, 31, 34
- [51] Helmuth Spieler. *Semiconductor Detector Systems*. Oxford University Press, Aug 2005. doi: 10.1093/acprof:oso/9780198527848.001.0001. URL <http://dx.doi.org/10.1093/acprof:oso/9780198527848.001.0001>. 21, 26

- [52] Cor Claeys and Eddy Simoen, editors. *Germanium-Based Technologies: From Materials to Devices*. Elsevier, London, 1st edition, 2007. ISBN 0-08-044953-0. doi: 10.1016/B978-0-08-044953-1.X5000-5. URL <https://doi.org/10.1016/B978-0-08-044953-1.X5000-5>. 21, 23
- [53] N. Abgrall, I.J. Arnquist, F.T. Avignone III, A.S. Barabash, F.E. Bertrand, A.W. Bradley, V. Brudanin, M. Busch, M. Buuck, J. Caja, and et al. The processing of enriched germanium for the majorana demonstrator and r&d for a next generation double-beta decay experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 877:314–322, Jan 2018. doi: 10.1016/j.nima.2017.09.036. URL <http://dx.doi.org/10.1016/j.nima.2017.09.036>. 22
- [54] Ichiro Yonenaga. *Germanium crystals*, pages 89–127. Elsevier, 2019. doi: 10.1016/b978-0-08-102096-8.00004-5. URL <http://dx.doi.org/10.1016/B978-0-08-102096-8.00004-5>. 22
- [55] *Improving Germanium Detector Resolution and Reliability*, 09 2008. URL <https://apps.dtic.mil/docs/citations/ADA516274>. 2008 Monitoring Research Review: Ground-Based Nuclear Explosion Monitoring Technologies. 24
- [56] Jingtian Fang, Mahmud Reaz, Stephanie L. Weeden-Wright, Ronald D. Schrimpf, Robert A. Reed, Robert A. Weller, Massimo V. Fischetti, and Sokrates T. Pantelides. Understanding the average electron/hole pair-creation energy in silicon and germanium based on full-band monte carlo simulations. *IEEE Transactions on Nuclear Science*, 66(1):444–451, Jan 2019. doi: 10.1109/tns.2018.2879593. URL <http://dx.doi.org/10.1109/TNS.2018.2879593>. 24
- [57] Zhong He. Review of the shockley-ramo theorem and its application in semiconductor gamma-ray detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 463(1-2):250 – 267, 2001. ISSN 0168-9002. doi: 10.1016/S0168-9002(01)00223-6. URL [https://doi.org/10.1016/S0168-9002\(01\)00223-6](https://doi.org/10.1016/S0168-9002(01)00223-6). 25
- [58] P. N. Luke, F. S. Goulding, N. W. Madden, and R. H. Pehl. Low capacitance large volume shaped-field germanium detector. *IEEE Transactions on Nuclear Science*, 36(1):926–930, Feb 1989. ISSN 0018-9499. doi: 10.1109/23.34577. 25, 29, 30
- [59] S. Mertens, A. Hegai, D.C. Radford, N. Abgrall, Y.-D. Chan, R.D. Martin, A.W.P. Poon, and C. Schmitt. Characterization of high purity germanium point contact detectors with low net impurity concentration. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 921:81–88, Mar 2019. doi: 10.1016/j.nima.2018.09.012. URL <http://dx.doi.org/10.1016/j.nima.2018.09.012>. 28, 30, 85
- [60] R.D. Martin, M. Amman, Y.D. Chan, J.A. Detwiler, J.C. Loach, Q. Looker, P.N. Luke, A.W.P. Poon, J. Qian, K. Vetter, and H. Yaver. Determining the drift time of charge carriers in p-type point-contact {HPGe} detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 678:98 – 104, 2012. ISSN 0168-9002. doi: <https://doi.org/10.1016/j.nima.2012.02.047>. URL <http://www.sciencedirect.com/science/article/pii/S0168900212002811>. 29
- [61] W. Xu, N. Abgrall, E. Aguayo, F.T. Avignone, A.S. Barabash, F.E. Bertrand, M. Boswell, and et. al. Testing the ge detectors for the majorana demonstrator. *Physics Procedia*, 61:807 – 815, 2015. ISSN 1875-3892. doi: <http://dx.doi.org/10.1016/j.phpro.2014.12.104>. URL <http://www.sciencedirect.com/science/article/pii/S1875389214007172>. 30

- [62] S.I. Alvis, I.J. Arnquist, F.T. Avignone III, A.S. Barabash, C.J. Barton, F.E. Bertrand, B. Bos, and et. al. Multi-site event discrimination for the majorana demonstrator. *arXiv*, 2019. URL <https://arxiv.org/abs/1901.05388>. 30
- [63] C. Canali, G. Majni, R. Minder, and G. Ottaviani. Electron and hole drift velocity measurements in silicon and their empirical relation to electric field and temperature. *IEEE Transactions on Electron Devices*, 22(11):1045–1047, Nov 1975. doi: 10.1109/t-ed.1975.18267. URL <http://dx.doi.org/10.1109/T-ED.1975.18267>. 31
- [64] Bart Bruyneel, Peter Reiter, and Gheorghe Pascovici. Characterization of large volume {HPGe} detectors. part i: Electron and hole mobility parameterization. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 569(3):764 – 773, 2006. ISSN 0168-9002. doi: <https://doi.org/10.1016/j.nima.2006.08.130>. URL <http://www.sciencedirect.com/science/article/pii/S0168900206015166>. 31, 34
- [65] G. Ottaviani, C. Canali, and A. Alberigi Quaranta. Charge carrier transport properties of semiconductor materials suitable for nuclear radiation detectors. *IEEE Transactions on Nuclear Science*, 22(1):192–204, 1975. doi: 10.1109/tns.1975.4327640. URL <http://dx.doi.org/10.1109/TNS.1975.4327640>. 31
- [66] D. M. Chang and J. G. Ruch. Measurement of the velocity field characteristic of electrons in germanium. *Applied Physics Letters*, 12(3):111–112, Feb 1968. doi: 10.1063/1.1651896. URL <http://dx.doi.org/10.1063/1.1651896>. 31
- [67] L. Reggiani, C. Canali, F. Nava, and G. Ottaviani. Hole drift velocity in germanium. *Physical Review B*, 16(6):2781–2791, Sep 1977. doi: 10.1103/physrevb.16.2781. URL <http://dx.doi.org/10.1103/PhysRevB.16.2781>. 34
- [68] A. Alberigi Quaranta, C. Jacoboni, and G. Ottaviani. Negative differential mobility in iii-v and ii-vi semiconducting compounds. *La Rivista Del Nuovo Cimento*, 1(4):445–495, Oct 1971. doi: 10.1007/bf02747246. URL <http://dx.doi.org/10.1007/BF02747246>. 34
- [69] L. Mihailescu, W. Gast, R.M. Lieder, H. Brands, and H. Jager. The influence of anisotropic electron drift velocity on the signal shapes of closed-end hpge detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 447(3):350?360, Jun 2000. doi: 10.1016/s0168-9002(99)01286-3. URL [http://dx.doi.org/10.1016/S0168-9002\(99\)01286-3](http://dx.doi.org/10.1016/S0168-9002(99)01286-3). 34
- [70] V. I. Ozhogin, A. V. Inyushkin, A. N. Taldenkov, A. V. Tikhomirov, G. E. Popov, E. Haller, and K. Itoh. Isotope effect in the thermal conductivity of germanium single crystals. *Journal of Experimental and Theoretical Physics Letters*, 63(6):490–494, Mar 1996. doi: 10.1134/1.567053. URL <http://dx.doi.org/10.1134/1.567053>. 34
- [71] Kohei Itoh, W.L. Hansen, E.E. Haller, J.W. Farmer, V.I. Ozhogin, A. Rudnev, and A. Tikhomirov. High purity isotopically enriched ⁷⁰ge and ⁷⁴ge single crystals: Isotope separation, growth, and properties. *Journal of Materials Research*, 8(6):1341–1347, Jun 1993. doi: 10.1557/jmr.1993.1341. URL <http://dx.doi.org/10.1557/JMR.1993.1341>. 34
- [72] Eugene E. Haller. Isotopically controlled semiconductors. *Journal of Nuclear Science and Technology*, 39(4):382–385, Apr 2002. doi: 10.1080/18811248.2002.9715208. URL <http://dx.doi.org/10.1080/18811248.2002.9715208>. 34

- [73] Eugene E. Haller. Isotopically controlled semiconductors. *MRS Bulletin*, 31(7):547–553, Jul 2006. doi: 10.1557/mrs2006.141. URL <http://dx.doi.org/10.1557/mrs2006.141>. 34
- [74] Wai-Kai Chen. *Circuits and filters handbook*. CRC Press, Boca Raton, Fla., 1st edition, 1995. ISBN 0-8493-8341-2. 38
- [75] S. Hutchinson, E. Keiter, R. Hoekstra, H. Watts, A. Waters, T. Russo, R. Schells, S. Wix, and C. Bogdan. The xyceTM parallel electronic simulator – an overview. In *Parallel Computing*. Published by Imperial College Press And Distributed By World Scientific Publishing Co., Jul 2002. doi: 10.1142/9781860949630_0021. URL http://dx.doi.org/10.1142/9781860949630_0021. 39
- [76] Eric R. Keiter, Scott A. Hutchinson, Robert J. Hoekstra, Thomas V. Russo, and Lon J. Waters. XyceTM parallel electronic simulator design: Mathematical formulation, version 2.0, 2004. URL https://xyce.sandia.gov/downloads/_assets/documents/Xyce_Math_Formulation.pdf. 39
- [77] John Anderson, Renato Brito, Dionisio Doering, Todd Hayden, Bryan Holmes, John Joseph, Harold Yaver, and Sergio Zimmermann. Data acquisition and trigger system of the gamma ray energy tracking in-beam nuclear array (gretina). *IEEE Transactions on Nuclear Science*, 56(1):258–265, Feb 2009. doi: 10.1109/tns.2008.2009444. URL <http://dx.doi.org/10.1109/TNS.2008.2009444>. 40
- [78] Wade D. Peterson. *The VMEbus Handbook, third Edition*. VFEA International Trade Association, 1993. 41
- [79] M. Munch, J. H. Jensen, B. Loher, H. Tornqvist, and H. T. Johansson. Vme readout at and below the conversion time limit. *IEEE Transactions on Nuclear Science*, 66(2):575–584, Feb 2019. doi: 10.1109/tns.2018.2884979. URL <http://dx.doi.org/10.1109/TNS.2018.2884979>. 41
- [80] Valentin T. Jordanov, Glenn F. Knoll, Alan C. Huber, and John A. Pantazis. Digital techniques for real-time pulse shaping in radiation measurements. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 353(1-3):261–264, Dec 1994. doi: 10.1016/0168-9002(94)91652-7. URL [http://dx.doi.org/10.1016/0168-9002\(94\)91652-7](http://dx.doi.org/10.1016/0168-9002(94)91652-7). 41
- [81] M.A. Howe, G.A. Cox, P.J. Harvey, F. McGirt, K. Rielage, J.F. Wilkerson, and J.M. Wouters. Sudbury neutrino observatory neutral current detector acquisition software overview. *IEEE Transactions on Nuclear Science*, 51(3):878–883, Jun 2004. doi: 10.1109/tns.2004.829527. URL <http://dx.doi.org/10.1109/TNS.2004.829527>. 42
- [82] ORCA. Object-oriented realtime control and acquisition, 2019. URL <http://orca.physics.unc.edu/>. Website. 42
- [83] David C. Radford, 2017. URL http://radware.phy.ornl.gov/MJ/mjd_siggen/. 43
- [84] David C. Radford, 2010. URL <http://radware.phy.ornl.gov/gretina/fieldgen/>. 43
- [85] Martin Alnæs, Jan Blechta, Johan Hake, August Johansson, Benjamin Kehlet, Anders Logg, Chris Richardson, Johannes Ring, Marie E Rognes, and Garth N Wells. The FEniCS project version 1.5. *Archive of Numerical Software*, Vol 3, 2015. doi: 10.11588/ans.2015.100.20553. URL <http://journals.ub.uni-heidelberg.de/index.php/ans/article/view/20553>. 44

- [86] Robert D. Cousins. Why isn't every physicist a bayesian? *American Journal of Physics*, 63(5): 398–410, May 1995. doi: 10.1119/1.17901. URL <http://dx.doi.org/10.1119/1.17901>. 49
- [87] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, Jun 1953. doi: 10.1063/1.1699114. URL <http://dx.doi.org/10.1063/1.1699114>. 51
- [88] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, Apr 1970. doi: 10.1093/biomet/57.1.97. URL <http://dx.doi.org/10.1093/biomet/57.1.97>. 51
- [89] John Skilling. Nested sampling for general bayesian computation. *Bayesian Analysis*, 1(4):833–859, Dec 2006. doi: 10.1214/06-ba127. URL <http://dx.doi.org/10.1214/06-ba127>. 52, 53
- [90] Brendon J. Brewer, Livia B. Pártay, and Gábor Csányi. Diffusive nested sampling. *Statistics and Computing*, 21(4):649–656, Aug 2010. doi: 10.1007/s11222-010-9198-8. URL <http://dx.doi.org/10.1007/s11222-010-9198-8>. xii, 52, 55
- [91] D.S. Sivia and J. Skilling. *Data Analysis: A Bayesian Tutorial*. Oxford University Press, 2006. ISBN 978-0-19-856832-2. 53
- [92] Brendon J. Brewer and Daniel Foreman-Mackey. Dnest4: Diffusive nested sampling in c++ and python. *Journal of Statistical Software*, 86(7), 2018. doi: 10.18637/jss.v086.i07. URL <http://dx.doi.org/10.18637/jss.v086.i07>. 55, 56, 60
- [93] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis, Third Edition*. Chapman and Hall/CRC, Nov 2013. doi: 10.1201/b16018. URL <http://dx.doi.org/10.1201/b16018>. 64, 66
- [94] Robert E. Kass and Adrian E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795, Jun 1995. doi: 10.1080/01621459.1995.10476572. URL <http://dx.doi.org/10.1080/01621459.1995.10476572>. 64
- [95] J. H. Hubbell and S. M. Seltzer. Nist: X-ray mass attenuation coefficients - germanium, July 2004. URL <https://dx.doi.org/10.18434/T4D01F>. NIST Standard Reference Database 126. 72
- [96] Pinghan Chu. Unidoc of Energy Performance of MAJORANA DEMONSTRATOR. Unidoc M-TECHDOCUNIDOC-2017-012, MAJORANA Collaboration, 2018. URL <https://mjdoc.npl.washington.edu/record/1856?ln=en>. Internal Document. 79
- [97] Gang Yang, Guojian Wang, Wenchang Xiang, Yutong Guan, Yongchen Sun, Dongming Mei, Bruce Gray, and Yuen-Dat Chan. Radial and axial impurity distribution in high-purity germanium crystals. *Journal of Crystal Growth*, 352(1):43–46, Aug 2012. doi: 10.1016/j.jcrysgro.2011.12.042. URL <http://dx.doi.org/10.1016/j.jcrysgro.2011.12.042>. 84
- [98] Dimitris G. Manolakis and Vinay K. Ingle. *Applied Digital Signal Processing*. Cambridge University Press, 2009. doi: 10.1017/cbo9780511835261. URL <http://dx.doi.org/10.1017/CBO9780511835261>. 88, 89
- [99] Alan V. Oppenheim and Ronald W. Schaffer. *Discrete-Time Signal Processing*. Pearson, 3 edition, 2010. 88

- [100] MicroModeler. MicroModeler DSP, 2019. URL <http://www.micromodeler.com/dsp/>. [Online; accessed 25-April-2019]. 88