

**The SNO+ liquid scintillator response to low-energy electrons
and its effect on the experiment's sensitivity to a future
neutrinoless double beta decay signal**

by

Logan Alfred Boyd Sibley

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Physics
University of Alberta

© Logan Alfred Boyd Sibley, 2016

Abstract

The SNO+ experiment is set to join the international competition of experiments searching for neutrinoless double beta decay. By loading 780 t of liquid scintillator with 0.5% natural tellurium, and with its location 2 km underground at SNOLAB, SNO+ aims to have sensitivity to determining the Majorana nature of the neutrino, a question currently at the forefront of particle physics, approaching the inverted hierarchy of the neutrino masses. To reach this sensitivity, it is crucial that SNO+ understands the response of the liquid scintillator, as systematic uncertainties on the energy scale and resolution, in particular any non-Gaussian shape of the energy resolution, may diminish the experiment's sensitivity in a significant way.

A ^{60}Co calibration source that tags calibration events within the liquid scintillator will enable SNO+ to precisely study the shape of the energy resolution near the endpoint of the ^{130}Te double beta decay. Monte Carlo simulations of the calibration source predict it will measure a 3.24% energy resolution at an energy of 2.51 MeV. Because ^{60}Co emits two gamma-rays upon decaying, whereas the expected signal of neutrinoless double beta decay is the sum of two electrons, it is also crucial for SNO+ to understand how the response of the liquid scintillator depends on particle type and energy. This dissertation provides the first measurement of the SNO+ tellurium-loaded liquid scintillator response to low-energy electrons. Ionization quenching of low-energy electrons in the tellurium-loaded liquid scintillator is small, with Birks' constant measured to be $kB = (4.1 \pm 2.9) \times 10^{-6}$ cm/keV.

Acknowledgements

This dissertation has been under construction for *many* years, requiring both patience and perseverance, not just by me, but by the numerous people who have supported me and my work.

First, thanks is due to my supervisor, Dr. Aksel Hallin, a fantastic mentor and instructor whose process of thought and patient inquisitiveness have always amazed me. I appreciate every opportunity I have had to learn from him. Secondly, thanks to Dr. Carsten Krauss, who trusted enough in my abilities to give me free reign of his laboratory, where I performed nearly all of my work for this dissertation. Dr. James Pinfeld, while not strictly involved in my work here, has always been a champion of mine, providing me opportunities to develop as teacher and scientist that I would not have had otherwise. I truly appreciate his continued support.

This work would not have been possible without support from wonderful technical and professional staff in the Department of Physics at the University of Alberta. Thanks to Paul Zimmerman, Len Wampler, Paul Davis and Chris Ng for helping me design (or for building!) much of the hardware I needed to complete my work. Suzette Chan and Sarah Derr offered assistance and counsel throughout my program and I will always be thankful for the conversation and support they provided. I am also grateful to Dr. Richard Soluk for his technical insight and overall guidance over many years.

I also extend a big thanks to the SNO+ collaboration for harbouring me these last seven years and providing me the experiment on which I based my thesis. Thanks also to SNOLAB for hosting me.

My time as a graduate student would have been incomplete without the support of the other graduate students, post-doctoral fellows, research associates and professors with whom my time overlapped—Rhys, Shalon, Zack, Kevin, Kevin, Matt, Pitam, Sarah, Tom, Jie, Courtney, Pooja, Asif, Aatif and Laura, and Dr.'s Pierre, Aleksandra, Mohammad, Berta, Kalpana, David, Solange, Rob and James, and Dr. Darren Grant—many of whom directly assisted me with my work. Particular mentions go to Tania Wood, the better half of “Team Age and Beauty”, and Pawel MekarSKI, a fervent supporter of “the Logan”. Both have been wonderful friends on whom I could rely for support at any time. I will sorely miss both the conversation and the coffee.

Next, thanks to my good friends Dr. Bob and Cathy deFrece, Benjamin Evans, and Shaun and Karli (and Will!) Andreas for the continued support these many years. Dr. Karyn Klut and Dr. Logan Banadyga have always been two of my strongest advocates and will always be among my greatest of friends. My family has formed my support system since the beginning. Korin, Verne, Delaney and Danica Hinch, and Morgan Sibley, have supplied endless years of love and support (interspersed with times of sibling rivalry), all of which have kept me grounded and secure. A great big thanks to my parents Elmer and Judy Sibley, who instigated my interest in science at a young age, and who supplied never-ending guidance throughout all my ambitions. I love them both deeply. Also, a sincere thanks to my second mom Dr. Val Congdon, whose love and encouragement I appreciate immensely.

Finally, to my wonderful wife Amy. You have supported me in this endeavor for *countless* years and have been my foundation throughout this journey, celebrating my highest achievements and anchoring me in the murkiest waters. I am so excited to continue our journey together! I do love you.

Table of Contents

Introduction	1
1 Neutrinos	5
1.1 Neutrino mass models	7
1.1.1 Dirac mass term	7
1.1.2 Majorana mass term	10
1.1.3 Combined mass term	12
1.2 Neutrino oscillations	14
1.2.1 Vacuum oscillations	16
1.2.2 Matter oscillations	21
1.2.3 Oscillation parameters	25
1.3 Measuring the neutrino mass	26
1.3.1 Neutrinoless double beta decay	29
1.4 Looking ahead	34
2 The SNO+ experiment	37
2.1 SNO+	38
2.1.1 The Sudbury Neutrino Observatory	38
2.1.2 Upgrading SNO \rightarrow SNO+	39

2.2	Physics Program	48
2.2.1	Supernova and solar neutrinos	48
2.2.2	Reactor and geo anti-neutrinos	50
2.2.3	Neutrinoless double beta decay	52
2.3	Experiment phases and timeline	54
3	SNO+ sensitivity systematics	57
3.1	^{150}Nd $0\nu\beta\beta$ sensitivity	60
3.2	Systematics	74
3.2.1	Tail Study	79
4	The ^{60}Co Calibration Source	86
4.1	Source design, construction and testing	89
4.1.1	^{60}Co standard	90
4.1.2	Plastic scintillator	91
4.1.3	Photomultiplier tube	117
4.1.4	Encapsulation	122
4.2	Simulation	132
4.2.1	The simulated ^{60}Co source geometry	133
4.2.2	The PMT simulation	137
4.2.3	Early geometry optimization studies	139
4.2.4	The ^{60}Co source generator	145
4.2.5	Simulation results	156
5	Quenching of low energy electrons	160
5.1	Scintillation and quenching	162
5.2	Apparatus	167

5.2.1	Acrylic chamber	169
5.2.2	Liquid scintillator process system	170
5.2.3	Detector assembly	175
5.2.4	Electronics, trigger logic and data acquisition	179
5.3	Data	183
5.3.1	Calibration	184
5.3.2	Data processing	190
5.4	Analysis	200
5.4.1	Monte Carlo simulation	205
5.5	Results	212
	Conclusion	215
	Bibliography	217
A	^{150}Nd sensitivity method	228
A.1	Robustness of the sensitivity method	228
A.2	Obtaining 90% confidence limits on ^{150}Nd $\tau_{1/2}^{0\nu}$ and $m_{eff}^{0\nu}$	234
B	^{60}Co calibration source	244
B.1	Source drawings	244
B.2	Calibration source procedures	253
B.2.1	Deposition procedure	253
B.2.2	Bonding procedure	256
B.2.3	Button counting and testing procedure	259
B.2.4	Source assembly and leak testing, and Delrin [®] annealing procedures	264
B.3	Simulation	274

B.3.1	Example RAT macro	274
B.3.2	Database fields	275

List of Tables

1.1	Isotopes that undergo $2\nu\beta\beta$	33
2.1	SNO+ calibration sources	45
3.1	Signal and backgrounds for a $0\nu\beta\beta$ sensitivity estimate	62
4.1	Button source dimensions, expected and measured activities	95
4.2	Dry wipe test contamination results	109
4.3	Wet wipe test contamination results	112
4.4	Soak test contamination results	113
4.5	^{60}Co decay branches	150
5.1	PMT SPE parameters	190
5.2	Collected data sets	201
5.3	MC L_0 scaling factors	208
5.4	LS and TeLS kB values	213
A.1	Minuit performance with varying tolerance	230
A.2	Minuit performance with varying fit energy range	230
B.1	^{60}Co source simulation database fields	280

List of Figures

1.1	Solar neutrino Feynman diagrams	23
1.2	Neutrino mass hierarchy and oscillation parameters	26
1.3	Tritium beta decay spectra for varying m_ν	29
1.4	$0\nu\beta\beta$ and $2\nu\beta\beta$ Feynman diagrams	30
1.5	$0\nu\beta\beta$ nuclear matrix elements	32
1.6	Experimental $0\nu\beta\beta$ sensitivities	36
2.1	SNO detector schematic	40
2.2	Camera system photo of SNO+	46
2.3	Expected MC solar neutrino energy spectra	50
2.4	Expected anti-neutrino energy spectra	52
2.5	Expected MC $0\nu\beta\beta$ signal and background energy spectra	55
3.1	Signal and background PDFs	63
3.2	Fit to one year MC data	64
3.3	Distributions of ^{214}Bi and ^{208}Tl scaling parameters and errors	66
3.4	Distributions of ^8B and $2\nu\beta\beta$ scaling parameters and errors	67
3.5	Distributions of the $0\nu\beta\beta$ scaling parameter and error	68
3.6	Correlation between the $0\nu\beta\beta$ scaling parameter and its error	69

3.7	Distribution of the 90% confidence limits on the number of observed $0\nu\beta\beta$ events	69
3.8	Distribution of $\tau_{1/2}^{0\nu}$ of $0\nu\beta\beta$	70
3.9	Distributions of 90% confidence limits on $\tau_{1/2}^{0\nu}$ and $m_{eff}^{0\nu}$ of $0\nu\beta\beta$	71
3.10	90% confidence limit on $0\nu\beta\beta$ $\tau_{1/2}^{0\nu}$ per operational year	72
3.11	90% confidence limit on $0\nu\beta\beta$ $m_{\beta\beta}$ per operational year	73
3.12	Effect of systematic shifts in the energy resolution on $\tau_{1/2}^{0\nu}$	76
3.13	Effect of systematic shifts in the energy resolution on $m_{\beta\beta}$	77
3.14	Effect of systematic shifts in the energy scale on $\tau_{1/2}^{0\nu}$	78
3.15	Effect of systematic shifts in the energy scale on $m_{\beta\beta}$	79
3.16	Function describing a high-energy tail on the energy resolution	81
3.17	Effect of high-energy tail on ^{150}Nd $2\nu\beta\beta$ energy spectrum	82
3.18	Effect of high-energy tail on observed number of $0\nu\beta\beta$ events	83
3.19	Effect of high-energy tail on $\tau_{1/2}^{0\nu}$ per operational year	84
3.20	Effect of high-energy tail on $m_{\beta\beta}$ per operational year	85
4.1	Samples of machined and polished plastic scintillator	93
4.2	Setup of the source deposition fume hood	97
4.3	The bonding jig and oven	99
4.4	Bonding temperature profile	100
4.5	$\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ crystal structure	101
4.6	$\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ residue extent pre- and post- bonding	102
4.7	Post-bonding air gap surrounding residue	103
4.8	Button source 4C microscope image	104
4.9	Button source 1B ^{60}Co rate	105
4.10	Position-dependent HPGGe efficiency	107

4.11	Temperature-dependent leak rate	115
4.12	Button source 1B microscope image	116
4.13	BC-630 soak test data	117
4.14	PMT efficiency	121
4.15	Assembling the copper encapsulation	124
4.16	Leak testing the copper encapsulation	126
4.17	The outer Delrin [®] encapsulation	128
4.18	The SNO+ source connector	130
4.19	Material compatibility test samples	131
4.20	Material compatibility test results	133
4.21	The simulated ⁶⁰ Co source geometry.	136
4.22	Simulated can source geometry	141
4.23	⁶⁰ Co source electron energy deposition and escape spectra	143
4.24	⁶⁰ Co source 1.33 MeV gamma-ray energy deposition and escape spectra	144
4.25	⁶⁰ Co source electron energy deposition versus gamma-ray escape energy for various geometries	145
4.26	Minimum screening correction	153
4.27	⁶⁰ Co electron energy spectra comparisons	155
4.28	⁶⁰ Co decay gamma-ray angular correlation	156
4.29	⁶⁰ Co source shadowing effect	158
4.30	⁶⁰ Co source quenching effect	159
5.1	Molecular excitation levels	164
5.2	Experiment design	168
5.3	Acrylic chamber schematic	170

5.4	Acrylic chamber photos	171
5.5	Liquid scintillator process system schematic	173
5.6	Quenching detector assembly	176
5.7	Quenching detector frame	178
5.8	Electronics schematic	182
5.9	HPGe calibration	186
5.10	PMT SPE calibration	189
5.11	Sample PMT and HPGe waveforms	192
5.12	PMT Δt distribution	194
5.13	PMT and HPGe time distributions	197
5.14	PMT average PE and HPGe E_{e^-} distributions	198
5.15	PMT average PE versus HPGe E_{e^-} distributions	199
5.16	Effect of assembly procedure on PMT average PE distribution	203
5.17	Complete data sets for LS and TeLS samples	204
5.18	Calculated light output for LS	207
5.19	LS unscaled and scaled MC	209
5.20	LS data versus MC as a function of kB	210
5.21	TeLS data versus MC as a function of kB	211
5.22	LS and TeLS kB dependency on χ^2 computation energy range	214
A.1	Number of fit $0\nu\beta\beta$ events with varying fit energy range . . .	231
A.2	Equivalent χ^2 of good and failed fits	233
A.3	Bias in failed fits	235
A.4	Functional forms of the $\tau_{1/2}^{0\nu}$ and m_{ν}^{eff} confidence limits	243
B.1	Copper sleeve	245
B.2	Copper flange bottom half	246

B.3	Copper flange upper half	247
B.4	Assembled copper encapsulation	248
B.5	Delrin [®] container	249
B.6	Delrin [®] stem	250
B.7	Quick connect	251
B.8	Complete source assembly	252

Introduction

*Ich komme vom Gebirge her,
Es dampft das Tal, es braust das Meer.
Ich wandle still, bin wenig froh,
Und immer fragt der Seufzer, wo?*

- Georg Philipp Schmidt,
set by Franz Schubert, *Der Wanderer*

The mystery of the neutrino has slowly been unravelling over the last half-century. From the first observation of anti-neutrinos in 1956, to the discovery in the late 1990s that neutrinos have mass, this particle continues to challenge our understanding of physics at the most basic level. Although we have learned much about this elusive particle, of the (many) questions remaining, perhaps none is as fundamental as that of the nature of the neutrino. Is this particle Dirac or Majorana—are the neutrino and anti-neutrino distinct or identical? The consequences for physics in either case are significant, and arguments on the side of each camp abound.

The SNO+ experiment is setting out to directly tackle this question. Refurbishing the infrastructure put in place for its predecessor, the Sudbury Neutrino Observatory, SNO+ will load 780 t of liquid scintillator with 4 t of natural tellurium to watch for the decay of its isotope ^{130}Te via the process

of neutrinoless double beta decay, which, if it occurs, is only possible if the neutrino is a Majorana particle. Its sheer size (SNO+ will be one of the world's largest liquid scintillator detectors) and depth underground (it is nearly 2 km underground at SNOLAB, where the cosmic ray background is low) will ensure SNO+ makes a significant contribution among a field of experiments investigating this problem. SNO+ is unique because the concentration of ^{130}Te in the detector is augmentable, which may help validate any observed hints of what would be an *extremely* rare radioactive decay process¹.

My supervisor, Dr. Aksel Hallin, once told me (and I paraphrase) that it is the fundamental questions in science that are those worth the most effort in pursuit, and it is this tenet that really drew me into this particular neutrino mystery. The overarching theme of my work on SNO+ is to understand the response of the SNO+ detector to a potential neutrinoless double beta decay signal. During this investigation, my work took on three distinct projects, each of which forms a chapter of this dissertation. In chapter 1, I present background information on neutrinos, in particular relating to neutrino mass generation schemes and measurement techniques. Chapter 2 contains a description of the SNO+ experiment and detector. The final three chapters highlight my major contributions to SNO+ and indicate where I received significant assistance to complete various tasks. Supervisory guidance should be assumed in all efforts.

First, I performed a Monte Carlo investigation of how the sensitivity of

¹ ^{130}Te is one of a handful of isotopes confirmed to undergo double beta decay, in which two neutrons in the nucleus are converted to two protons via the release of two electrons and two electron anti-neutrinos. The measured half-life for this process is 7.14×10^{24} years. If neutrinoless double beta decay occurs, where only the electrons are emitted in the decay, the half-life would be *even longer*, making this one of the most rare processes to occur on earth, and very difficult to observe experimentally.

SNO+ to observing a neutrinoless double beta decay signal depends on our understanding of systematic alterations in the detector’s energy response. If we incorrectly set the detector’s energy scale or resolution, how do we affect our ability to make a neutrinoless double beta decay measurement? In particular, how do small tails on an otherwise Gaussian energy response mask or mimic a signal? Chapter 3 provides an overview of my work to answer these questions, with supplemental information in appendix A.

One conclusion of this study is that a high-precision calibration of the SNO+ energy response will be necessary to make a strong statement about any supposed neutrinoless double beta decay observation. As my second contribution, I designed a tagged ^{60}Co calibration source for the SNO+ detector. This source will provide a practically background-free set of calibration data, needed to precisely understand the detector energy response. Chapter 4 presents in detail the effort I made in designing, testing and building this calibration source, as well as describes the Monte Carlo simulation of the source I created. For this task, I received support in designing and building the source from technical staff in the electronics and machine shops in the Department of Physics. I also received assistance from colleagues when working with open radioactive sources and in testing encapsulation methods for the ^{60}Co in plastic scintillator. Much of the inspiration for the source encapsulation comes from similar sources built for the SNO detector, although this specific tagged source design is new for SNO+. While most of the design details are complete, the final source will not be built by me. Rather, I have transitioned this work to a SNO+ colleague who is not only taking over the source manufacturing, but who has also updated the Monte Carlo simulation to account for later-stage changes in the source design. I obtained the simulation results in

section 4.2.5 using his updated version of the simulation. Appendix B contains detailed source drawings and manufacturing procedures, as well as additional information relevant to the Monte Carlo source simulation.

While simulating the SNO+ response to the ^{60}Co calibration source, it became clear that it is critical to understand how the light output of the liquid scintillator varies with particle type and energy. Thus, my third major contribution, which chapter 5 outlines, was to measure the change in light output of the liquid scintillator as a function of energy for low-energy electrons. While this had previously been measured for the nominal SNO+ liquid scintillator, it had *not* been measured for the liquid scintillator loaded with natural tellurium, the addition of which may further affect the light output. SNO+ observes neutrinoless double beta decay as the summed light output of two electrons simultaneously interacting in the liquid scintillator. Because some of these electrons are necessarily low in energy, this measurement is essential if we are to be confident in our understanding of the detector response. Again, I received assistance from colleagues and machinists to build the apparatus I used to make this measurement. The collection and analysis of the data, however, I performed as a solo effort.

A brief conclusion brings the major results of my thesis together to close this dissertation.

Chapter 1

Neutrinos

We are happy to inform you that we have definitely detected neutrinos from fission fragments by observing inverse beta decay of protons.

- Frederick Reines and Clyde Cowan,
*In a telegram to Wolfgang Pauli, 1956
and highlighted in Reines 1995 Nobel Lecture*

Thanks for the message. Everything comes to him who knows how to wait.

- Pauli, *In reply to Reines and Cowan*

Physicists studying nuclear beta decay in 1930 faced an acute dilemma. Experimentally, they measured a whole spectrum of energies for the electron emitted in a nuclear transition, in direct opposition to the notion that beta decay was a two-body decay. That the electron did not always have the same, well-defined energy led some to suggest that beta decay violated a fundamental tenet of physics: the law of conservation of energy.

To rescue the laws of physics from this looming peril, W. Pauli took the bold step, against even his own intuition, to postulate that a new, electrically neutral (to conserve charge), nearly massless (because the endpoint of the electron energy spectrum was equal to the expected two-body decay electron energy) particle was emitted in conjunction with the electron, implying beta decay was in fact a three-body process and could be reconciled with data [1]. These neutrinos, as they came to be known, proved to be elusive, however, and more than 20 years would pass before C. Cowan and F. Reines would observe the inverse beta decay signal of an anti-neutrino interaction in a small liquid scintillator detector they had placed near the Savannah River reactor of the U.S. Atomic Energy Commission [2], finally providing unequivocal evidence for the neutrino's existence.

There are three neutrinos in the standard model (SM) of particle physics, one for each of the charged leptons, ν_e , ν_μ and ν_τ (owing to the width of the Z^0 implying three light lepton generations [3]). They have zero charge, no mass, have been observed to only exist in a left-handed helicity state (anti-neutrinos come in a right-handed variety), and interact exclusively via the weak force. However, observations beginning in the late 1990s by Super Kamiokande [4, 5] and the Sudbury Neutrino Observatory (SNO) [6] that a neutrino's flavour oscillates as it traverses space quickly shattered this notion and showed conclusively that neutrinos have non-vanishing mass. Since then, neutrino oscillation physics has rapidly found its place as a mature field. Nevertheless, while we have uncovered many of the neutrino's mysteries, we still have much to learn.

1.1 Neutrino mass models

That neutrinos are able to change flavour in transit implies the state in which they are created is not equivalent to the state in which they propagate. Neutrinos are created, and we observe them experimentally, as flavour eigenstates of the weak force (in fact, we define the neutrino flavour to be that of the charged lepton that accompanies its creation at a weak interaction vertex), but they propagate through space as the mass eigenstates ν_1 , ν_2 and ν_3 . The flavour eigenstates are superpositions of the mass eigenstates (and vice versa) and, because the mass eigenstates propagate at different velocities owing to their different masses, a neutrino created as one flavour will become a superposition of all three flavours as phase differences between the propagating mass eigenstates develop. Knowing this compels us to investigate whether adding a neutrino mass term to the standard model Lagrangian leads to this mixing, and what the consequences of this addition may be.

1.1.1 Dirac mass term

The most general Lorentz invariant mass term we may add to the SM Lagrangian density (while ignoring that the SM admits only left-handed neutrinos or right-handed anti-neutrinos) is, following the formalism in [7] and [8],

$$\mathcal{L}_D^\nu = - \sum_{\alpha, \beta} \nu_{\alpha L}^\dagger m_{\alpha\beta}^D \nu_{\beta R} + h.c. , \quad (1.1)$$

where $m_{\alpha\beta}^D$ is some arbitrary complex matrix, α and β run over the neutrino flavours e , μ and τ , and $h.c.$ refers to the Hermitian conjugate. The $\nu_{\alpha L}$ and $\nu_{\alpha R}$ are left- and right-handed chiral projections of the spinor, call it ν ,

that represents the Dirac fermion field. We have used the standard chirality projection operators

$$P_L = \frac{1 - \gamma_5}{2} \quad \text{and} \quad P_R = \frac{1 + \gamma_5}{2}, \quad (1.2)$$

to write out the left- and right-handed projections of ν as

$$\nu_L = P_L \nu \quad \text{and} \quad \nu_R = P_R \nu, \quad (1.3)$$

with $\nu = \nu_L + \nu_R$ (noting that chirality really only coincides with helicity, from which we have the concept of handedness, for massless particles). Under charge conjugation¹, the components transform as

$$(\nu_L)^c = (\nu^c)_R \quad \text{and} \quad (\nu_R)^c = (\nu^c)_L, \quad (1.4)$$

where the spinors and their charge conjugates give us four degrees of freedom.

We may re-write our arbitrary complex matrix $m_{\alpha\beta}^D$ as a real diagonal matrix m_i^D by applying a bi-unitary transformation

$$m_{\alpha\beta}^D = \sum_i U_{\alpha i}^{L\dagger} m_i^D U_{\beta i}^R, \quad (1.5)$$

where U^L and U^R are unitary (so $U^\dagger U = U U^\dagger = I$). Using these unitary

¹From [8], the charge conjugation takes the form $(\nu_L)^c = C \bar{\nu}_L^T$, where C is the charge conjugate operator and T indicates the transpose.

matrices, we define new neutrino fields

$$\begin{aligned}\nu_{iL} &= \sum_{\alpha} U_{\alpha i}^L \nu_{\alpha L} \\ \nu_{iR} &= \sum_{\alpha} U_{\alpha i}^R \nu_{\alpha R} ,\end{aligned}\tag{1.6}$$

with which the Lagrangian density mass term takes the standard Dirac form

$$\mathcal{L}_D^\nu = - \sum_i m_i^D (\nu_{iL}^\dagger \nu_{iR} + \nu_{iR}^\dagger \nu_{iL}) .\tag{1.7}$$

Here, m_i corresponds to the real mass of neutrino type $i = 1, 2, 3$. Because the Dirac mass term is invariant under the global phase transformation

$$\nu \rightarrow e^{i\phi} \nu \text{ and } \bar{\nu} \rightarrow e^{-i\phi} \bar{\nu} ,\tag{1.8}$$

it implicitly conserves fermion number and generates equal masses for particles and their anti-particles. Also, owing to the unitarity of U^L and U^R , we may invert equation 1.6 and express our neutrino flavour fields as

$$\begin{aligned}\nu_{\alpha L} &= \sum_i U_{\alpha i}^{L*} \nu_{iL} \\ \nu_{\alpha R} &= \sum_i U_{\alpha i}^{R*} \nu_{iR} ,\end{aligned}\tag{1.9}$$

whereby neutrinos with definite flavour (ν_e, ν_μ, ν_τ) become mixtures of neutrinos with definite mass (ν_1, ν_2, ν_3).

Note that a Dirac mass term of this type requires both left- and right-handed neutrinos, where adding right-handed neutrinos would be an extension of the SM. In fact, neutrinos remain massless in the SM precisely because it

excludes right-handed neutrinos. While introducing right-handed neutrinos into the SM seems to be a minimal way to account for neutrino mass, it suffers from a difficulty to account for a *small* neutrino mass, which could require an unnaturally small Yukawa coupling many orders of magnitude below that for charged fermions [9, 10], or the addition of a second Higgs doublet into the SM to provide the neutrino masses [11].

1.1.2 Majorana mass term

In 1937, Ettore Majorana laid the groundwork for the idea that the neutrino and anti-neutrino may actually be identical [12]. Beginning with the left- and right-handed projections of equation 1.3, Majorana's condition is that the right-handed field is equivalent to the charge-conjugate of the left-handed field

$$\nu_{\text{R}} = (\nu_{\text{L}})^c = (\nu^c)_{\text{R}} . \tag{1.10}$$

Allowing for an arbitrary phase factor ϕ ,

$$\begin{aligned} \nu &= \nu_{\text{L}} + e^{i\phi}(\nu_{\text{L}})^c \\ \nu^c &= e^{-i\phi}\nu . \end{aligned} \tag{1.11}$$

Thus, the field ν is equivalent, within a phase factor, to its charge-conjugate, and there are only two degrees of freedom [9].

The Lagrangian density may now have two different mass terms

$$\begin{aligned}\mathcal{L}_M^{\nu_L} &= -\frac{1}{2} \sum_{\alpha,\beta} \nu_{\alpha L}^\dagger m_{\alpha\beta L}^M (\nu_{\beta L})^c + h.c. \\ \mathcal{L}_M^{\nu_R} &= -\frac{1}{2} \sum_{\alpha,\beta} \nu_{\alpha R}^\dagger m_{\alpha\beta R}^M (\nu_{\beta R})^c + h.c. ,\end{aligned}\tag{1.12}$$

where, as a result of the charge conjugation,

$$m_{L,R}^M = (m_{L,R}^M)^T ,\tag{1.13}$$

for each of m_L^M and m_R^M . Because the Majorana mass matrix is symmetric, we may relax the requirement for a bi-unitary transformation to diagonalize it, where

$$m_{\alpha\beta L,R}^M = \sum_i U_{\alpha i}^{L,R\dagger} m_{iL,R}^M U_{\beta i}^{L,R} .\tag{1.14}$$

Once again, we may define new neutrino fields as in equation 1.6, which, with equation 1.11, allow us to write the Lagrangian density as [13]

$$\mathcal{L}_M^{\nu_{L,R}} = -\frac{1}{2} \sum_i m_{iL,R}^M \nu_{iL,R}^\dagger \nu_{iL,R} .\tag{1.15}$$

Again, the flavour states are linear combinations of the mass states, in which case we have, as in equation 1.9,

$$\begin{aligned}\nu_{\alpha L} &= \sum_i U_{\alpha i}^{L*} \nu_{iL} \\ \nu_{\alpha R} &= \sum_i U_{\alpha i}^{R*} \nu_{iR} .\end{aligned}\tag{1.16}$$

If we look more closely at equation 1.12, then we see that both the left- and right-handed Lagrangian density terms appear to be equivalent, so we

could consider dropping the right-handed term altogether and forgo the need to introduce right-handed neutrinos into the model. However, using only the left-handed term in the Lagrangian density demands further non-minimal extensions of the SM in order to generate the Majorana neutrino mass [9]. Not only that, because the Lagrangian density is no longer invariant under a global phase transformation (owing to equation 1.11 and the fact we now only have two degrees of freedom), it no longer conserves lepton number. Rather, the Lagrangian density now violates lepton number by two units.

1.1.3 Combined mass term

Rather than selecting preferentially either the Dirac or Majorana mass terms for the Lagrangian density, the most general mass term we could consider is their sum, whereby our mass term becomes [13]

$$\begin{aligned}
\mathcal{L}_{\text{D+M}}^\nu &= -\frac{1}{2} \sum_{\alpha,\beta} ((\nu_{\alpha\text{L}}^\dagger)^c m_{\alpha\beta\text{L}}^{\text{M}} \nu_{\beta\text{L}} + 2\nu_{\alpha\text{R}}^\dagger m_{\alpha\beta}^{\text{D}} \nu_{\beta\text{L}} + \nu_{\alpha\text{R}}^\dagger m_{\alpha\beta\text{R}}^{\text{M}} (\nu_{\beta\text{R}})^c) + h.c. \\
&= -\frac{1}{2} ((\nu_{\text{L}}^\dagger)^c, \nu_{\text{R}}^\dagger) \begin{pmatrix} m_{\text{L}}^{\text{M}} & (m^{\text{D}})^T \\ m^{\text{D}} & m_{\text{R}}^{\text{M}} \end{pmatrix} \begin{pmatrix} \nu_{\text{L}} \\ (\nu_{\text{R}})^c \end{pmatrix} + h.c.
\end{aligned} \tag{1.17}$$

For clarity, we have removed the explicit matrix multiplication in order to write down the Lagrangian density in a more compact form. We now diagonalize our mass matrix (treating it as a single generation toy-model) to obtain the neutrino mass states

$$\begin{aligned}
\nu_{1\text{L}} &= \cos \theta (\nu_{\text{L}}) - \sin \theta (\nu_{\text{R}})^c & (\nu_{1\text{L}})^c &= \cos \theta (\nu_{\text{L}})^c - \sin \theta (\nu_{\text{R}}) \\
\nu_{2\text{L}} &= \sin \theta (\nu_{\text{L}}) + \cos \theta (\nu_{\text{R}})^c & (\nu_{2\text{L}})^c &= \sin \theta (\nu_{\text{L}})^c + \cos \theta (\nu_{\text{R}}) ,
\end{aligned} \tag{1.18}$$

where the mixing angle θ is given by

$$\tan 2\theta = \frac{2m^{\text{D}}}{m_{\text{R}}^{\text{M}} - m_{\text{L}}^{\text{M}}} . \quad (1.19)$$

Once again, the Majorana terms violate lepton number conservation by two units, while the Dirac term conserves total lepton number (though not lepton flavour number). We end up with two independent Majorana fields with charge-parity (CP) eigenvalues $\epsilon_{1,2}$ and masses [14]

$$m_{1,2} = \epsilon_{1,2} \frac{1}{2} \left((m_{\text{L}}^{\text{M}} + m_{\text{R}}^{\text{M}}) \pm \sqrt{(m_{\text{L}}^{\text{M}} - m_{\text{R}}^{\text{M}})^2 + 4(m^{\text{D}})^2} \right) . \quad (1.20)$$

There are several important scenarios we may extract from the above. If the mixing angle $\theta = 45^\circ$ ($m_{\text{L}}^{\text{M}} = m_{\text{R}}^{\text{M}} = 0$), then $m_{1,2} = m^{\text{D}}$ and we have a pure Dirac field composed of two degenerate Majorana fields [14]. Therefore, we see that the Dirac case is a simplified solution of the Majorana case. Alternatively, if there is no mixing such that $\theta = 0^\circ$ ($m^{\text{D}} = 0$), then $m_{1,2} = m_{\text{L,R}}^{\text{M}}$ and we are left with the case where the neutrinos are purely Majorana.

One interesting scenario is when we assume $m_{\text{L}}^{\text{M}} = 0$ and allow $m_{\text{R}}^{\text{M}} \gg m^{\text{D}}$. In this case, our mixing angle $\theta \ll 1$, $\epsilon_{1,2} = \mp 1$ and we obtain two mass eigenvalues

$$m_1 \approx \frac{(m^{\text{D}})^2}{m_{\text{R}}^{\text{M}}} \quad \text{and} \quad m_2 \approx m_{\text{R}}^{\text{M}} . \quad (1.21)$$

This formalism is the basis for the *see-saw* mechanism of neutrino mass generation [10], where we assume that, in addition to the Higgs mechanism for the generation of the Dirac masses in the SM, there is some beyond-the-SM mechanism that generates the right-handed Majorana mass term. If we assume the

Dirac mass term is of the order of the electroweak scale, so $m^D \sim 100 \text{ GeVc}^{-2}$, and the Majorana mass term is at the scale of new physics, say the grand unification scale, so $m_R^M \sim 10^{15} \text{ GeVc}^{-2}$, then $m_1 \sim 0.01 \text{ eVc}^{-2}$, which is close to the neutrino mass scale indicated by cosmological observations [9]. As such, this scenario leads to a light neutrino mass, and it naturally follows that we associate the physical neutrino mass with m_1 . However, it also requires very heavy right-handed neutrinos. We often refer to these neutrinos as *sterile*, implying they do not interact via the weak force. This may explain why such heavy neutrinos, if they exist, have so far evaded our attention. Lastly, the mixing implies that both the light and heavy neutrinos are Majorana particles, which has consequences that we may explore experimentally (see section 1.3.1).

In a full three-generation scenario, the physical neutrino mass is represented as a symmetric matrix

$$m_\nu \equiv m_1 = m^D (m_R^M)^{-1} (m^D)^T, \quad (1.22)$$

which we may diagonalize as before via a unitary transformation

$$\text{diag}(m_1, m_2, m_3) = U^\dagger m^D (m_R^M)^{-1} (m^D)^T U. \quad (1.23)$$

Again, this unitary transformation would allow us to rewrite our flavour states as combinations of mass states, as previously.

1.2 Neutrino oscillations

Regardless of the mechanism through which neutrinos obtain mass, the formalism we developed in section 1.1 enables us to investigate the physical con-

sequences of there being a non-zero neutrino mass. In any SM weak interaction in which a neutrino is created in conjunction with a charged lepton partner, we define that neutrino to be in a pure flavour state equivalent to that of the charged lepton (so as to conserve lepton flavour). As we saw in section 1.1, this implies the neutrino does not have a definite mass, but is created as some superposition of mass states. The charged lepton, on the other hand, is created in a state of both well-defined mass and flavour, so we say, for example, that a charged lepton with electron-type flavour has a mass of $0.511 \text{ MeV}c^{-2}$. In other words, we identify each charged lepton flavour state with a particular mass state.

Equations 1.9 and 1.16 give us the prescription for writing a neutrino flavour state in terms of mass states. Using Dirac's bra-ket notation and explicitly writing the neutrino states as functions of spacetime coordinates, a neutrino created in flavour state $|\nu_\alpha\rangle$ will evolve in time as

$$|\nu_\alpha(\vec{x}, t)\rangle = \sum_i U_{\alpha i}^* |\nu_i(\vec{x}, t)\rangle . \quad (1.24)$$

If the neutrino is detected after time t , again via some weak interaction so it is detected as a pure flavour state, then the probability we will detect it in the flavour state $|\nu_\beta\rangle$ will be

$$\begin{aligned} P_{\alpha \rightarrow \beta} &= |\langle \nu_\beta | \nu_\alpha(x, t) \rangle|^2 \\ &= \left| \sum_{i,j} U_{\alpha i}^* U_{\beta j} \langle \nu_j | \nu_i(x, t) \rangle \right|^2 , \end{aligned} \quad (1.25)$$

where we have taken the coherent sum over the states to reflect the fact that we cannot distinguish between the different mass states at the time the neutrino is

created. Now, before we can go further, we need to describe the time evolution of the neutrino mass states.

1.2.1 Vacuum oscillations

If we assume the neutrinos are free (propagating in vacuum), then the mass state $|\nu_i\rangle$ is a stationary state and will evolve as a plane wave, where

$$|\nu_i(\vec{x}, t)\rangle = e^{-i(E_i t - \vec{p}_i \cdot \vec{x})} |\nu_i(\vec{x}, 0)\rangle . \quad (1.26)$$

Now, if we assume the neutrino is highly relativistic, then in time t it will travel a distance $L = t$ (now working in units where $\hbar = c = 1$). We may also approximate the neutrino's momentum as

$$p_i = \sqrt{E_i^2 - m_i^2} = E_i \sqrt{1 - \frac{m_i^2}{E_i^2}} \approx E_i - \frac{m_i^2}{2E_i} , \quad (1.27)$$

in which case

$$E_i t - p_i x \approx E_i L - \left(E_i - \frac{m_i^2}{2E_i} \right) L = \frac{m_i^2 L}{2E_i} \approx \frac{m_i^2 L}{2E} , \quad (1.28)$$

where we make the assumption that each state i is created with the same energy $E \equiv E_i$ (there is no good reason to assume this, particularly because we have evolved the neutrino as a plane wave [15]). Taking advantage of the

orthonormality of the states, $\langle \nu_\alpha | \nu_\beta \rangle = \delta_{\alpha\beta}$, we have

$$\begin{aligned}
P_{\alpha \rightarrow \beta} &= \sum_{i,j} U_{\alpha i} U_{\alpha j}^* U_{\beta i}^* U_{\beta j} e^{-im_i^2 L/2E} e^{im_j^2 L/2E} |\langle \nu_j | \nu_i \rangle|^2 \\
&= \delta_{\alpha\beta} - 4 \sum_{i>j} \text{Re}(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2 \left(\Delta m_{ij}^2 \frac{L}{4E} \right) \\
&\quad + 2 \sum_{i>j} \text{Im}(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin \left(\Delta m_{ij}^2 \frac{L}{2E} \right) , \quad (1.29)
\end{aligned}$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$. For an anti-neutrino, we may simply replace U with U^* . We see that the probability of transition between states oscillates as a function of the baseline distance L . It is then convenient to define an oscillation wavelength to give a scale of the spatial extent of the oscillation

$$L_{ij}^0 \equiv \frac{4\pi\hbar c E}{\Delta m_{ij}^2 c^4} \simeq 2.48 \text{ km} \frac{E(\text{GeV})}{\Delta m_{ij}^2(\text{eV}^2)} , \quad (1.30)$$

where we have put back in the values for c and \hbar to obtain the constant.

These oscillation probabilities do not hold, though, if the neutrino is created in a particular mass state. Because the flavour state is a mixture of the mass states, the probability for selecting a specific mass state $|\nu_i\rangle$ from a flavour state $|\nu_\alpha\rangle$ is just

$$P_i = |\langle \nu_i | \nu_\alpha \rangle|^2 = |U_{i\alpha}|^2 . \quad (1.31)$$

The probability the neutrino is then detected some time later in state $|\nu_\beta\rangle$ is

$$\begin{aligned}
P_{\alpha \rightarrow \beta} &= \sum_i |\langle \nu_\beta | \nu_i \rangle e^{-im_i^2 L/2E} \langle \nu_i | \nu_\alpha \rangle|^2 \\
&= \sum_i |U_{\alpha i}|^2 |U_{\beta i}|^2 . \quad (1.32)
\end{aligned}$$

So, the probability to observe the neutrino as a different flavour state from

that in which it was created is simply a function of the mixing parameters and is not dependent on the coordinates. In this case, we take the incoherent sum of the states, a result of knowing the exact initial state of the neutrino.

Up to now, we have not said anything about the mixing matrix U . If we assume three neutrino generations, then we let $\alpha \in \{e, \mu, \tau\}$ and $i \in \{1, 2, 3\}$, and the mixing matrix U has nine components. An arbitrary complex 3×3 matrix has $2 \times 3^2 = 18$ real parameters; however, because our mixing matrix is unitary, nine of these parameters are constrained by the conditions

$$\begin{aligned} \sum_{\alpha} U_{\alpha i} U_{\alpha i}^* &= 1 \quad i \in \{1, 2, 3\} \\ \sum_{\beta \neq \alpha} U_{\alpha i} U_{\beta i}^* &= 0 \quad i \in \{1, 2, 3\}, \alpha \in \{e, \mu, \tau\}. \end{aligned} \quad (1.33)$$

Of the remaining nine parameters, three correspond to rotations and six are phases, three of which will be absorbed into the phases of the charged leptons in the interaction Lagrangian [16]. In the case of Dirac neutrinos, two of the remaining phases may be absorbed into the global phase transformation of equation 1.8, where the remaining phase cannot be absorbed and is physical. Using these four parameters, we may write U , which we typically refer to as the Pontecorvo-Maki-Nakagaya-Sakata (PMNS) mixing matrix, as (in its common parametrization)

$$\begin{aligned}
U &= \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau3} & U_{\tau3} \end{pmatrix} \\
&= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \\
&= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}, \quad (1.34)
\end{aligned}$$

where $c_{ij} \equiv \cos \theta_{ij}$, $s_{ij} \equiv \sin \theta_{ij}$ with θ_{ij} the mixing angle, and δ is the phase.

For Majorana neutrinos, because the fields are not invariant under a global phase transformation, we are left with three physical phases. One of these phases is absorbed into δ in the Dirac case above, where we refer to the other two, call them ϕ_1 and ϕ_2 , as Majorana phases. We may add the Majorana phases into the PMNS matrix as

$$U \rightarrow U \cdot \begin{pmatrix} e^{i\phi_1/2} & 0 & 0 \\ 0 & e^{i\phi_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.35)$$

There are several interesting conclusions we may draw from equation 1.29 that have important physical consequences. The matrix containing the Majorana phases simply cancels out in the oscillation probability, so we cannot use neutrino oscillations to reveal whether the neutrino is Majorana in nature.

Another consequence is that, given a transition probability $P_{\alpha\rightarrow\beta}$ and its CP conjugate $P_{\bar{\alpha}\rightarrow\bar{\beta}}$, the difference between these rates is non-zero only if CP is not conserved, such that $\delta \neq 0, \pi$. This is a direct result of the third term of equation 1.29 flipping sign when taking the CP conjugate of a transition probability. Thus, this third term encodes all of the information regarding CP violation in neutrino oscillations [17].

Lastly, neutrinos may only exhibit oscillations if they are massive where at least two mass states are not degenerate so that $\Delta m_{ij}^2 \neq 0$. The mass states are different so they will propagate at different velocities, leading to interference effects and oscillations. For $\Delta m_{ij}^2 \ll 1$ eV, the states remain coherent over very long distances and L_{ij}^0 becomes large (for relatively low energies). Interestingly, then, we should also note that given a baseline $L \gg L_{ij}^0$, the mass states will have enough time to move sufficiently far apart to become decoherent, leading to a probability for observing a particular flavour according to equation 1.32, where oscillations cease to occur. This is particularly relevant for extra-galactic neutrino sources where the baseline to an earth-based experiment is extremely long. Also, if $\Delta m_{ij}^2 \sim 1$ eV, then L_{ij}^0 can be very short, on the order of metres (which is one reason why we do not realistically consider charged lepton oscillations, which have $\Delta m_{ij}^2 \gg 1$ eV [18]). So, in the scenario where very heavy right-handed sterile neutrinos contribute to the mixing, we may expect to see rapid neutrino flavour oscillations on a scale the size of a typical particle detector.

We also see that oscillations only allow us to measure the differences between the squares of the neutrino masses, not the masses themselves. In other words, neutrino oscillations give us no information as to the absolute neutrino mass scale. With three neutrino generations, there are two independent mass

splittings, which we may choose to be Δm_{12}^2 and Δm_{23}^2 . These values then constrain the value of the third splitting according to the relationship [19]

$$\Delta m_{12}^2 + \Delta m_{23}^2 + \Delta m_{31}^2 = 0 . \quad (1.36)$$

Often, information about the mass splittings is presented using the alternate parameters Δm^2 and δm^2 , where

$$\Delta m^2 = m_3^2 - (m_1^2 + m_2^2)/2 \quad \text{and} \quad \delta m^2 \approx \Delta m_{12}^2 . \quad (1.37)$$

The information about the neutrino masses is actually even more limited. Looking at equation 1.29, if we make the assumption that the CP-violating phase δ is close to 0 or π , then the oscillation probability ceases to depend on the sign of the mass splitting (neutrino oscillations in matter (section 1.2.2) are sensitive to the sign of the mass splittings, so we may glean information from there). Currently, only the sign of the mass difference Δm_{21}^2 is known to be positive [20]. As a result, it is unclear whether $(m_3 > m_2 > m_1)$ or $(m_2 > m_1 > m_3)$, scenarios we refer to as the normal and inverted hierarchy, respectively. If nature does choose the inverted hierarchy, then neutrinos would be the only particles whose masses would uncharacteristically not increase with generation number [21].

1.2.2 Matter oscillations

We have seen how, in vacuum, neutrino oscillations occur as a direct result of phase differences arising between neutrino mass states as they propagate with different velocities. As with any fundamental particle, however, we may

also (or, in some cases, *moreso*) be interested in what happens when that particle travels through and interacts with matter. If we assume a neutrino is propagating through matter with an interaction characterized by a potential V , then the total energy of the neutrino mass state is $E+V$, where E is the energy it would have if propagating in vacuum. If different neutrino flavours feel different potentials, such that $V_\alpha \neq V_\beta$, then phase differences may also develop and neutrino oscillations may be incited through these matter interactions [19].

In matter, neutrinos may coherently interact via charged current or neutral current interactions. From the weak interaction Hamiltonian [22, 9, 13], the potential a neutrino experiences while travelling through electrically neutral matter comprised of electrons, protons and neutrons is

$$V_\alpha = \pm\sqrt{2}G_F \left(n_e\delta_{\alpha e} - \frac{1}{2}n_n \right) , \quad (1.38)$$

where $G_F = 1.16637 \text{ GeV}^{-2}$ is the Fermi coupling constant, n_e is the electron density, n_n is the nucleon density, the positive sign refers to neutrinos and the negative sign to anti-neutrinos. The left term is the charged current contribution and the right term is the neutral current contribution. For solar neutrinos propagating in the sun, only electron neutrinos may interact via charged currents, since neither muon nor tau neutrinos have enough energy to create a physical muon or tauon. All types, though, may interact through neutral currents, as the Feynman diagrams for these processes in figure 1.1 show. Because the neutral current contribution is common to all three neutrino flavours, it simply adds a phase to the oscillation probability, so only the charged current component of the potential is important. Also, if the matter density changes along a neutrinos's propagation path, then n_e and n_n explicitly

become functions of the neutrino path length.

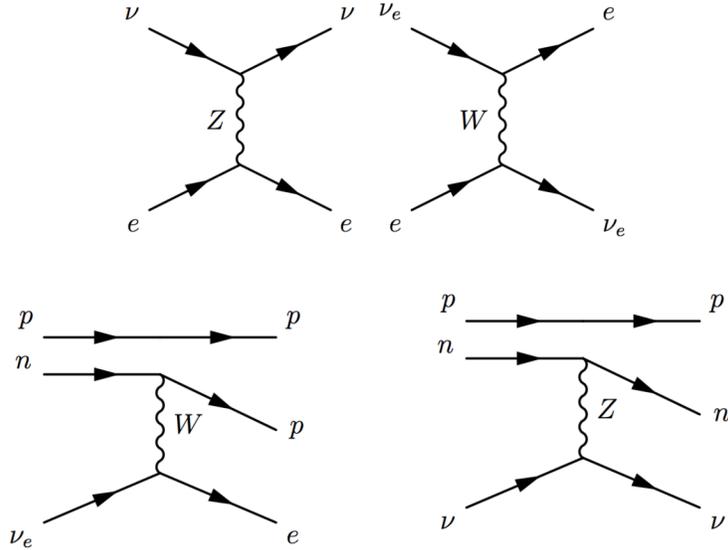


Figure 1.1: Feynman diagrams for elastic scattering (top), charged current (bottom, left) and neutral current (bottom, right) interactions of solar neutrinos in the SNO D₂O detector. Note that, in the sun, only electron neutrinos interact via charged current interactions, since muon and tau neutrinos at solar neutrino energies are unable to produce their charged lepton partners [19]. The neutral current interactions and a fraction of the total elastic scatters are sensitive to all neutrino flavours. This property is what allowed SNO to measure the total neutrino flux relative to the electron neutrino flux emanating from the sun.

Taking this potential into account breaks the diagonality of the mass matrix, implying the mass eigenstates in vacuum are not eigenstates in matter [14, 19]. To obtain the new mass eigenstates, we must diagonalize the new mass matrix via some unitary transformation, as in equation 1.5, characterized by a single rotation angle θ_m (working with only two flavours). Defining $A = \pm 2\sqrt{2}G_F E n_e$, the mass eigenstates become

$$m_{1m,2m}^2 = \frac{1}{2} \left((\Sigma m^2 + A) \mp \sqrt{(A - \Delta m^2 \cos 2\theta)^2 + (\Delta m^2)^2 \sin^2 2\theta} \right), \quad (1.39)$$

with a mixing angle

$$\sin 2\theta_m = \frac{\sin 2\theta}{\sqrt{(A/\Delta m^2 - \cos 2\theta)^2 + \sin^2 2\theta}}. \quad (1.40)$$

To recover the case of vacuum oscillations, we let $A \rightarrow 0$, whereby we see that both $m_{1,2m}^2 \rightarrow m_{1,2}^2$ and $\theta_m \rightarrow \theta$, as we should expect [14].

Like the case of vacuum oscillations, there are several interesting conclusions we may form simply by examining equations 1.39 and 1.40. If the vacuum mixing angle $\theta = 0$, then the matter mixing angle $\theta_m = 0$ always. We conclude that matter oscillations require the existence of vacuum oscillations, which we know from section 1.2.1 are a result of neutrinos having mass.

Next, the value of θ_m is sensitive to a change in sign of Δm^2 , a result of the first term under the square root in the denominator of equation 1.40. Therefore, matter oscillations give us some handle on determining the mass hierarchy, whereas, as we have already seen in section 1.2.1, we cannot get this information by measuring vacuum oscillations only [19].

Lastly, we have already seen that allowing $A \rightarrow 0$ reduces matter oscillations to the vacuum scenario. If a neutrino propagates through very dense matter, where $A \rightarrow \infty$, then this forces $\theta_m \rightarrow 0$, implying that oscillations will not occur. However, we also see from the form of equation 1.40 that there is a resonance condition, $A/\Delta m^2 = \cos 2\theta$, where θ_m becomes maximal at 45° . So, there is some value for the electron density n_e that causes a 100% probability for oscillation, even if the probability for oscillation in vacuum is small. Thus, we may get *matter enhanced* oscillations, which we refer to as the Mikheyev-Smirnov-Wolfenstein (MSW) effect [22]. This scenario has a dramatic effect on neutrinos emanating from the solar core, many of which (the MSW effect

depends on the neutrino energy, as well) oscillate to muon neutrinos within the sun, remaining as such along their entire journey to earth. It is this effect that enabled SNO (with data from KamLAND) to determine the sign of the Δm_{12}^2 mass splitting [23].

1.2.3 Oscillation parameters

We describe neutrino oscillations with three mixing angles, θ_{12} , θ_{23} and θ_{13} , two mass differences, Δm_{12}^2 and Δm_{23}^2 (or Δm^2 and δm^2), and one CP-violating complex phase, δ , for a total of six parameters (or, if neutrinos are Majorana particles, so we include two additional phases, ϕ_1 and ϕ_2 , then eight parameters). Equation 1.34 portrays the neutrino mixing in a convenient manner, where the mixing matrix is the product of three separate matrices, each of which depends on a single mixing angle. The left-most matrix, depending on θ_{23} , describes cases when L/E is small, namely atmospheric and accelerator neutrino oscillations. The right-most matrix depends on θ_{12} , describing solar neutrino and reactor anti-neutrino mixing, which have large L/E . The third mixing angle, θ_{13} , is measured by investigating the disappearance of electron anti-neutrinos in reactor anti-neutrino experiments and the appearance of electron neutrinos in muon neutrino beams.

A recent global fit [24] sets $\sin^2 \theta_{23} = 0.437_{-0.020}^{+0.033}$ ($0.569_{-0.051}^{+0.028}$) and $|\Delta m^2| = 2.50_{-0.04}^{+0.04} \times 10^{-3} \text{ eV}^2$ ($2.46_{-0.04}^{+0.05} \times 10^{-3} \text{ eV}^2$) for the normal (inverse) hierarchy. The global fit returns $\sin^2 \theta_{12} = 0.297_{-0.017}^{+0.017}$ and $\delta m^2 = 7.37_{-0.16}^{+0.17} \times 10^{-5} \text{ eV}^2$. Lastly, $\sin^2 \theta_{13} = 0.0214_{-0.0009}^{+0.0011}$ ($0.0218_{-0.0012}^{+0.0009}$) for the normal (inverted) hierarchy. The global fit sets $\delta/\pi = 1.35_{-0.22}^{+0.29}$ ($1.32_{-0.25}^{+0.35}$) for the normal (inverted) hierarchy, showing that a value for $\delta \neq \pi$ is favoured at the 1σ confidence

level (though at 3σ , the fit still allows δ to be anywhere between 0 and 2π).

Figure 1.2 condenses the oscillation parameters into graphical form [25].

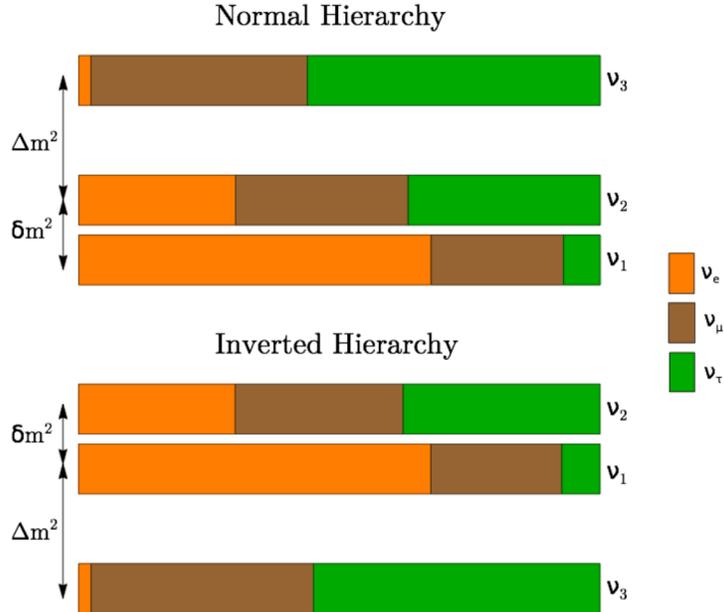


Figure 1.2: This diagram succinctly summarizes the neutrino oscillation parameters (with the exception of the phase δ). The colours denote the relative flavour constitution of each mass state, corresponding to the values of the mixing angles θ_{ij} , with yellow for ν_e , brown for ν_μ and green for ν_τ . We clearly see the smallness of θ_{13} relative to the other mixing angles by the minute contribution of ν_e to the ν_3 mass state. Because the sign of Δm_{23}^2 is unknown, the plot shows both the normal and inverted hierarchy scenarios, with the smallness of δm^2 relative to Δm^2 shown not to scale [26].

1.3 Measuring the neutrino mass

While the phenomenon of neutrino oscillations clearly indicates that neutrinos have non-zero mass, neutrino oscillations cannot, unfortunately, tell us what the neutrino masses *are*. As we saw in section 1.2, neutrino oscillations only give us the difference between the squares of the masses, with the added com-

plication that determining the absolute sign of those differences is non-trivial experimentally. Ultimately, however, we would like to know the values (or, at the very least, the scale) of the neutrino masses.

Fortunately, there are several well-posed experimental methods we may use to attempt to measure the neutrino mass. At the largest of scales, cosmology is able to offer insight on how large the neutrino mass may be. Because the neutrino is such a prevalent particle in the universe, second in number only to the ubiquitous photon, too large a neutrino mass may have a significant effect on the large scale structure of the universe. As a result, we may set an upper bound on the sum of the neutrino masses via model-dependent cosmological arguments, with a current upper limit on the sum of the masses of $0.23 \text{ eV}/c^2$ [27].

Examining the kinematics of beta decay may also give us insight into the neutrino mass scale [28, 29]. In a beta decay, the energy of the emitted electron takes on one of any spectrum of values given by

$$\frac{dN}{dE} = \sqrt{E^2 - m_e^2} E (E_0 - E) \sqrt{(E_0 - E)^2 - m_\nu^2} F(Z, E) , \quad (1.41)$$

where m_e and E are the electron mass and energy, E_0 is the total energy released in the decay, and $F(Z, E)$ is the Fermi function that corrects for distortions in the spectrum resulting from interactions between the daughter nucleus, with atomic number Z , and the outgoing electron². Because the neutrino mass m_ν is so small, we may generally disregard it when considering a typical radioactive decay, since it will have a marginal effect on the electron endpoint energy. However, if E_0 is small enough, then the fractional change in

²Several other corrections are possible. See section 4.2.4.

the spectrum as a result of a finite neutrino mass becomes large and potentially measurable. Taking the small neutrino mass into account leads to a downward shift in the electron endpoint energy and steepens the approach of the energy spectrum to the endpoint, as figure 1.3 shows [30].

Because we know (by definition) that an electron anti-neutrino is released in a beta decay, we cannot be certain as to what mass state the anti-neutrino is created in. If an experiment had perfect energy resolution, then it would actually be able to measure three different endpoints lying atop one another corresponding to each of the anti-neutrino mass states. However, because the differences between the masses are much smaller than even the energy scale of a low-endpoint beta decay, an experiment generally only measures an incoherent sum of the neutrino masses, where

$$m_{\nu_e}^2 = \sum_i |U_{ei}|^2 m_i^2 , \quad (1.42)$$

and U is the PMNS mixing matrix in equation 1.34. Thus, kinematic measurements of m_ν are only able to measure the neutrino mass scale.

Taking advantage of the small 18.59 keV endpoint energy of tritium, the Mainz-Troitsk experiments have set the best upper bound on the mass of the electron anti-neutrino with $m_{\bar{\nu}_e} < 2.05 \text{ eV}/c^2$ at a 95% confidence limit [3, 28]. The Karlsruhe TRItium Neutrino (KATRIN) experiment will also use tritium to look for distortions in the shape of the electron spectrum, with an estimated sensitivity of $m_{\nu_e} < 0.2 \text{ eV}/c^2$ [31].

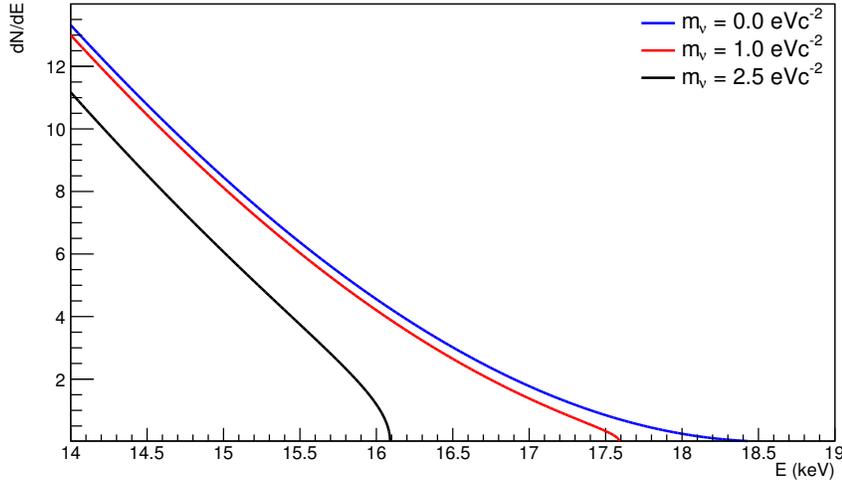


Figure 1.3: When $m_\nu \neq 0$ there is a downward shift in the endpoint energy of electrons emitted in beta decay. This graph shows the electron energy spectrum of tritium decay (neglecting the Fermi function), with $E_0 = 18.59$ keV.

1.3.1 Neutrinoless double beta decay

It was 1935 when Goeppert-Mayer noted that in rare cases, when a nucleus is kinematically forbidden to beta decay, it may instead undergo the process of double beta decay ($2\nu\beta\beta$), where the nucleus releases two electrons and two electron anti-neutrinos, increasing its atomic number Z by two in the process [32]. Taking this together with Majorana's idea that the neutrino and anti-neutrino may be identical (see section 1.1.2), Furry pointed out just a few years later in 1939 that $2\nu\beta\beta$ could also proceed, though more rarely, as neutrinoless double beta decay ($0\nu\beta\beta$). In this scenario, the two anti-neutrinos (requiring they have non-zero masses), each other's anti-particles under the Majorana condition of equation 1.10, are virtually exchanged, resulting in an end state that contains only two electrons with a fixed sum-energy. Figure 1.4 shows a Feynman diagram of the $0\nu\beta\beta$ process [33].

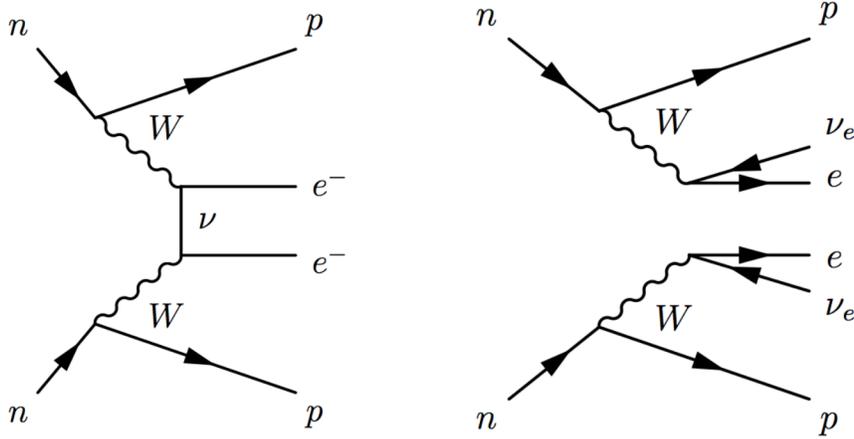


Figure 1.4: The Feynman diagram of $0\nu\beta\beta$ (left) shows how the decay proceeds via the exchange of a virtual neutrino, violating lepton number by two units in the process and emitting just two electrons [33]. The Feynman diagram of $2\nu\beta\beta$ (right) shows how two anti-neutrinos are also emitted in the final state.

We may write the decay rate for $2\nu\beta\beta$ as [34]

$$\lambda_{2\nu} = \frac{1}{\tau_{2\nu}} = G_{2\nu}(Q, Z)g_A^4 |m_e c^2 M^{(2\nu)}|^2 . \quad (1.43)$$

Using a Primakoff-Rosen approximation, this rate becomes an explicit function of the decay Q -value (the amount of available energy for the decay products), whereby

$$\lambda_{2\nu} \approx Q^7 \left(1 + \frac{Q}{2} + \frac{Q^2}{9} + \frac{Q^3}{90} + \frac{Q^4}{1980} \right) . \quad (1.44)$$

Here, $G_{2\nu}$ is the phase-space factor, a function of the Q -value and the atomic number Z , $M^{(2\nu)}$ is the matrix element, $g_A = 1.269$ is the axial coupling constant, and $\tau_{2\nu}$ is the $2\nu\beta\beta$ lifetime. The decay rate is proportional to Q^{11} .

For $0\nu\beta\beta$, the situation is somewhat different. We now have a neutrino propagator in our matrix element, unlike the $2\nu\beta\beta$ case, and, because only left-handed fields exist in the weak Hamiltonian, when we integrate over all internal

neutrino momenta, a term proportional to the neutrino mass remains [14, 35]. The resulting decay rate for this process is

$$\lambda_{0\nu} = \frac{1}{\tau_{0\nu}} = G_{0\nu}(Q, Z)g_A^4 |M^{(0\nu)}|^2 \left(\frac{m_{\beta\beta}}{m_e}\right)^2, \quad (1.45)$$

which again depends on a phase-space factor $G_{0\nu}$ and matrix element $M^{(0\nu)}$. $\tau_{0\nu}$ is the $0\nu\beta\beta$ lifetime. There is also an explicit dependence on the square of the effective neutrino mass $m_{\beta\beta}$ (which, as in the case for beta decay, is what we are measuring), where [33]

$$\begin{aligned} m_{\beta\beta} &= \left| \sum_i U_{ei}^2 m_i \right| \\ &= \left| \cos^2 \theta_{12} \cos^2 \theta_{13} m_1 + e^{i\phi_1/2} \sin^2 \theta_{12} \cos^2 \theta_{13} m_2 + e^{i\phi_2/2} \sin^2 \theta_{13} m_3 \right| \end{aligned} \quad (1.46)$$

and ϕ_1 and ϕ_2 are the Majorana phases we saw in equation 1.35. Interestingly, depending on their values, the Majorana phases may lead to interference in the sum and cause the effective neutrino mass to become zero [14]. This interference is non-existent in the case of beta decay with a Dirac neutrino.

Also of note in equation 1.45 is that the effective neutrino mass is inversely proportional to the $0\nu\beta\beta$ lifetime. This is of profound significance because, in the event an experiment observes $0\nu\beta\beta$, the number of $0\nu\beta\beta$ events it observes during its period of data collection gives us the effective neutrino mass to within a set of constants made up of a calculable phase-space factor and matrix element. One challenge, though, lies in the calculation of the matrix element, which involves complex nuclear physics and which may be approached using a variety of methods that typically yield different results, as figure 1.5 shows. While these calculations continue to be refined, our knowledge of the

matrix element provides a large uncertainty in any determination of an effective neutrino mass. Also note that this effective mass is only valid for Majorana neutrinos [14].

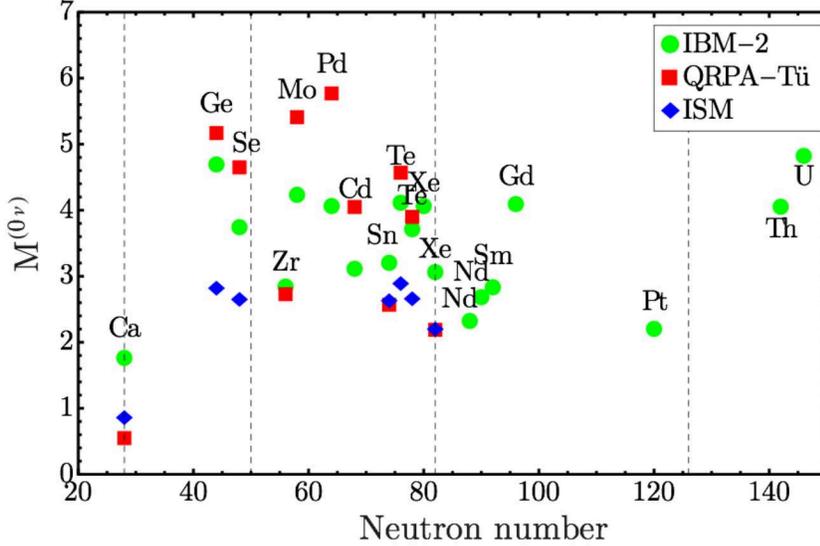


Figure 1.5: Calculating $M^{(0\nu)}$ for $0\nu\beta\beta$ is challenging, and different methods of calculation lead to a spread in the values for all $2\nu\beta\beta$ isotopes. This spread presents a large uncertainty in relating a $0\nu\beta\beta$ decay rate to an effective neutrino mass via equation 1.45.

Again, assuming a Primakoff-Rosen approximation, as a function of Q the $0\nu\beta\beta$ rate is [14]

$$\lambda_{0\nu} \approx \frac{Q^5}{30} - \frac{2Q^2}{3} + Q - \frac{2}{5}, \quad (1.47)$$

where the rate only scales as Q^5 , a factor of 10^{-6} lower than that of the $2\nu\beta\beta$ scenario. This implies that the $0\nu\beta\beta$ process, if it exists, will be significantly *more rare* than the already rare $2\nu\beta\beta$ process. Table 1.1 lists some of the isotopes that are confirmed to undergo $2\nu\beta\beta$ [36, 34, 37], along with some of their properties.

A host of experiments have performed searches for $0\nu\beta\beta$ using a variety

Isotope	Abundance (%)	Q -value (keV)	$G^{0\nu}$ (10^{-15}yr^{-1})	$M^{(0\nu)}$
^{48}Ca	0.187	4273.7	24.81	1.98
^{76}Ge	7.8	2039.1	2.363	5.42
^{82}Se	9.2	2995.5	10.16	4.37
^{96}Zr	2.8	3347.7	20.58	2.53
^{100}Mo	9.6	3035.0	15.92	3.73
^{110}Pd	11.8	2004.0	4.815	3.62
^{116}Cd	7.6	2809.1	16.70	2.78
^{124}Sn	5.6	2287.7	9.040	3.50
^{130}Te	34.5	2530.3	14.22	4.03
^{136}Xe	8.9	2461.9	14.58	3.33
^{150}Nd	5.6	3367.3	63.03	2.32

Table 1.1: Isotopes known to undergo $2\nu\beta\beta$ vary substantially in natural abundance and Q -values for ground-state-to-ground-state transitions. Also shown are phase-space factors [34] and nuclear matrix elements (also see figure 1.5) [37] calculated for the $0\nu\beta\beta$ of these isotopes.

of the isotopes table 1.1 lists, and many more are either being constructed or designed [26]. Of all the experiments, however, just one subgroup of the Heidelberg-Moscow collaboration, led by Klapdor-Kleingrothaus, claims to have observed $0\nu\beta\beta$ using ^{76}Ge . Their analysis results in a best fit for the $0\nu\beta\beta$ half-life of $(2.23_{-0.31}^{+0.44}) \times 10^{25}$ yr and effective neutrino mass of 320 meV [38, 39]. This claim is now completely ruled out by other experiments, the best limits of which figure 1.6 shows [40]. The most stringent limits come from KamLAND-Zen, a kilotonne-scale liquid scintillator detector loaded with ^{136}Xe , where they measure a $0\nu\beta\beta$ half-life greater than 1.1×10^{26} yr at 90% confidence, placing an upper bound on the effective neutrino mass in the range 60–161 meV [41]. Other ^{76}Ge experiments (GERDA, IGEX and Heidelberg-Moscow) rule out the claimed $0\nu\beta\beta$ observation with GERDA supplying the best lower bound on the $0\nu\beta\beta$ half-life of 5.2×10^{25} yr at 90% confidence [42, 43]. The bolometric ^{130}Te experiment CUORE (as Cuoricino and CUORE-0) [44] also rules out the claimed $0\nu\beta\beta$ observation, with a combined bound on the $0\nu\beta\beta$ half-life

of 4.0×10^{24} yr and effective neutrino mass in the range 270–760 meV, at 90% confidence.

$0\nu\beta\beta$ experiments also have some ability to constrain the mass hierarchy of Majorana neutrinos, as figure 1.6 shows. Looking at equation 1.47, we see that the effective neutrino mass depends on the values of each neutrino mass state m_1 , m_2 and m_3 . In the inverted hierarchy where m_3 is the lightest neutrino mass state, the effective neutrino mass is one order of magnitude larger than in the normal hierarchy when the lightest neutrino mass state has a mass less than approximately 1 meV. As experiments continue to push down the upper limit on the effective neutrino mass, one of two scenarios will arise. Either experiments will rule out the inverted hierarchy region completely (if they observe no evidence for $0\nu\beta\beta$ at very long $0\nu\beta\beta$ half-life limits), or they will observe $0\nu\beta\beta$ and confirm the inverted hierarchy is that found in nature. If the first scenario unfolds, then we will know the neutrino masses follow a normal hierarchy if neutrinos are Majorana particles. In this case, though, the effective Majorana mass can be effectively 0 meV, rendering $0\nu\beta\beta$ impossible to observe [26].

1.4 Looking ahead

While experiments in the last quarter century have determined many of the properties of neutrinos, notably the structure and real elements of the PMNS mixing matrix and the relative difference between the neutrino mass states, there are still many questions concerning the fundamental nature of neutrinos that remain unanswered.

We have yet to measure the absolute neutrino mass scale, and cannot yet

explain, but only hypothesize, why the neutrino masses are so much smaller than those of the charged leptons. In addition, until we are able to measure the sign of the oscillation parameter Δm_{23}^2 , we will not know whether the neutrinos follow a normal mass hierarchy, with the highest generation number mass state also being the heaviest, or an inverted hierarchy, where the highest generation number mass state is uncharacteristically the lightest. Whether the neutrino is a Dirac or a Majorana particle also remains a mystery. The best solution for testing the latter case is poised to come from $0\nu\beta\beta$ experiments, but the measurement is difficult, and the interpretation of data is a non-trivial task. However, that this process demands lepton number violation makes it a worthwhile measurement. Lastly, the value of the CP-violating phase δ remains unknown, though constraints on its value are continually narrowing.

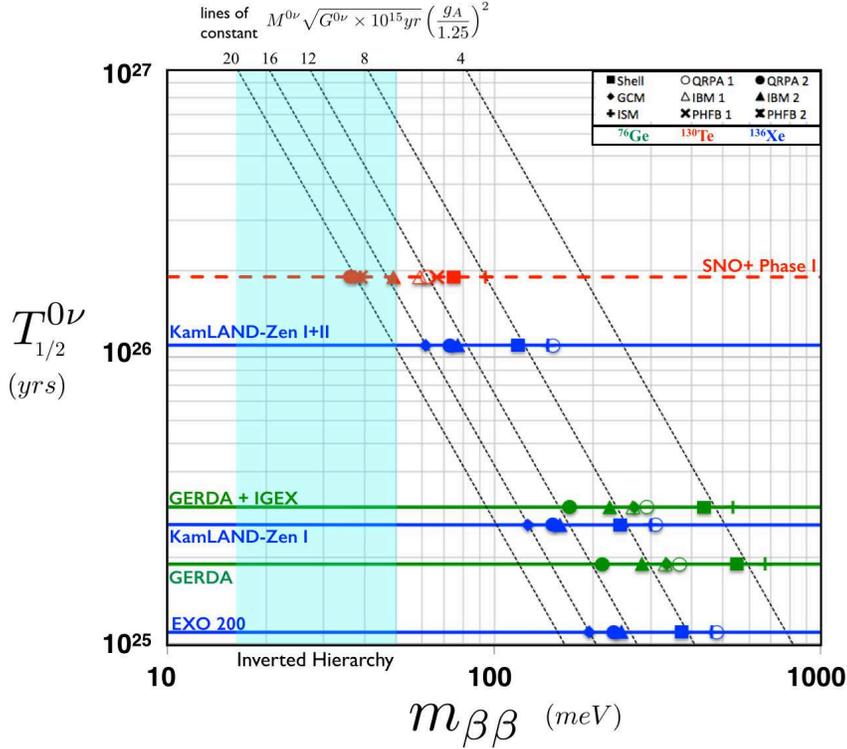


Figure 1.6: Setting longer limits on the $0\nu\beta\beta$ half-life enables experiments to push down the upper limit on the effective neutrino mass $m_{\beta\beta}$. In this plot [40], the most stringent half-life limits from experiments using ^{76}Ge (green) and ^{136}Xe (blue) are shown against the projected limits of the SNO+ experiment, which will use ^{130}Te (red). The spread in $m_{\beta\beta}$ resulting from the variety of calculations of $M^{(0\nu)}$ is evident. The diagonal, dashed black lines are lines of constant $G_{0\nu}g_A^4 |M^{(0\nu)}|^2$ from equation 1.45, which show how changes to $G_{0\nu}$, $M^{(0\nu)}$ or g_A affect an experiment's sensitivity. The cyan region indicates values of $m_{\beta\beta}$ that correspond to the inverted hierarchy solution. If $m_{\beta\beta}$ lies below this region, then the inverted hierarchy solution will be ruled out experimentally (if the neutrino is a Majorana particle).

Chapter 2

The SNO+ experiment

To illustrate my point, one would not write a scientific paper describing an experiment in which an experimenter stood on a mountain and reached for the moon, and concluded that the moon was more than eight feet from the top of the mountain.

- Referee comment to Ray Davis, *Phys. Rev.* **97**, 766–769

Joining the current generation of neutrinoless double beta decay experiments, SNO+ is a kilotonne-scale liquid scintillator (LS) particle detector housed at SNOLAB 2 km underground in the Vale Canada Ltd. Creighton Mine near Sudbury, Ontario, Canada. The detector largely reuses the components of its predecessor, the Sudbury Neutrino Observatory (SNO) [45], but incorporates many design changes that will enable it to obtain a lower energy threshold and, consequently, increase its physics reach [46, 47, 48, 49].

SNO+ is an international collaboration¹ consisting of more than 100 scientists (scientists and engineers at various stages in their respective careers) from Canada, Germany, Mexico, Portugal, the United States and the United

¹SNO+ typically describes three distinct entities: a detector, an experiment and a collaboration.

Kingdom. As with any large scientific collaboration, each member of SNO+ contributes in a unique way to the experiment, and each member performs service duties that directly affect the detector’s operability. SNO+, as an experiment, is the achievement of the collective work of all of these individuals. Credit for the experiment’s success is due to the collaboration as a whole.

In the sections below, I briefly describe the SNO+ experiment. Section 2.1.1 gives a brief overview of SNO, with section 2.1.2 outlining the upgrades that transform SNO into SNO+. Specifically, this section describes the SNO+ LS, including its composition, optical properties and purification, the SNO+ calibration hardware and the Monte Carlo (MC) simulation of the detector. The physics SNO+ will investigate is described in section 2.2, highlighting the investigation of $0\nu\beta\beta$ using ^{130}Te . A description of the experiment phases (and physics measurements to be made in each phase) follows in section 2.3.

2.1 SNO+

2.1.1 The Sudbury Neutrino Observatory

SNO consisted of a 12 m diameter acrylic vessel surrounded by 9438 inward-facing photomultiplier tubes (PMTs)² held on a 17.8 m diameter PMT support structure (PSUP). Each PMT was mounted inside a 27 cm diameter light concentrator, which, together with the PMTs, provided SNO a 54% effective photocathode coverage. 1000 t of extremely pure heavy water (D_2O) filled the acrylic vessel (AV), while a bath of ultra pure water (UPW) filled the remaining extent of the SNO cavern, including the space between the PSUP

²The PMTs are 8” Hamamatsu R1408.

and AV. Because of the density difference of these two media, a set of 10 ropes attached at the equator held the AV in place during the experiment's lifetime. SNO detected neutrinos by observing Čerenkov light resulting from neutrino interactions in the D₂O. The use of D₂O enabled SNO to observe interactions from all three neutrino flavours (leading to the confirmation of neutrinos having non-zero mass [6] and a share of the 2015 Nobel Prize in Physics [50]). Figure 2.1 shows a cartoon schematic of the SNO detector [45].

The UPW surrounding the AV acted as a shield to reduce backgrounds from sources external to the AV, such as the PMTs and surrounding rock walls of the SNO cavern. The rock walls were also coated in a concrete material covered in a radon-impermeable Urylon liner, which limited the extent to which the UPW in the cavity could leach radioactivity from the rock and limited the ingress of radon gas into the cavity, further reducing background levels. These layers also acted as a water seal to prevent gross leakage of UPW from the cavity into the surrounding rock. The 6000 metres water equivalent (MWE) rock overburden above the SNO detector acted as an effective cosmic ray shield, limiting the number of muons passing through the detector to just 70 per day [49]. An additional 91 outward-facing PMTs were used to veto cosmic muons passing through the cavity.

2.1.2 Upgrading SNO → SNO+

The change defining the move from SNO to SNO+ is the replacement of the D₂O with LS (section 2.1.2 describes the LS in detail). The LS density is lower than that of water (much lower than D₂O) at 0.86 g/cm³, so the water shielding will exert a large buoyant force on the filled AV. Like the ropes

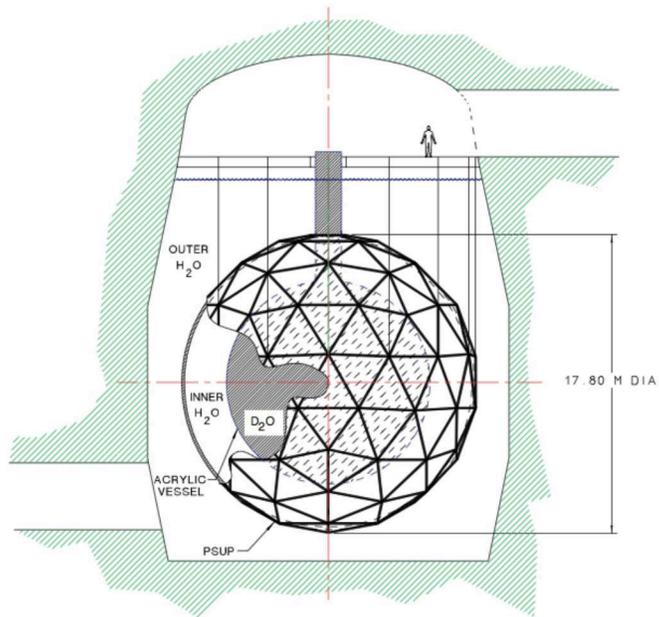


Figure 2.1: This cartoon schematic of the SNO detector shows the AV and PSUP sitting relative to the SNO cavern. A human is drawn on the upper deck for size reference.

that held up the SNO AV, a new hold-down rope net has been installed over the AV to counteract this 1.25 MN of buoyant force³ [51]. The rope net is anchored to the cavity floor at 20 locations. A physical barrier (in the form of an ‘umbrella’ suspended from the PSUP base) separated the construction work during the anchor installation on the cavity floor from the clean AV. Underwater load cells near each anchor point enable the monitoring of the tension in each arm of the rope net. The rope material is TensylonTM, chosen because it contains sufficiently low levels of ⁴⁰K contamination, a significant background for SNO+ solar neutrino physics goals. New TensylonTM ropes have also replaced the original hold-up ropes. The rope net is visible in the

³I assisted with laboratory measurements of the strain a sample rope arm would place on a test piece of acrylic cut from a spare acrylic panel used in the SNO AV construction.

picture of SNO+ in figure 2.2.

Because the LS will produce on the order of 100 times more light than the amount of Čerenkov light produced in SNO, SNO+ will be able to collect data at a very low energy threshold (hundreds of keV). This means SNO+ will experience a much higher event rate from low energy events and will have more PMT signals to process per event at nominal energies (several MeV). As such, the SNO electronics and data acquisition have been upgraded to handle the increased throughput requirements [52]. A series of dry runs in 2012 and 2014 allowed for testing of the new electronics, when high voltage was supplied to the PMTs and triggered test-data were collected. The AV was filled either with air (meaning not filled) or partially filled with water during these runs.

SNO+ will reuse the PMTs from SNO. Damaged PMTs have been repaired on the bottom half of the PSUP, with plans to replace damaged top-half PMTs from a boat as the cavity is filled with water and as the experiment's schedule permits. Approximately 9400 PMTs will be operational for the start of the SNO+ data collection.

The internal walls and top external surface of the AV have been completely cleaned to remove dust and debris that settled there in the period after draining the SNO cavity when the AV was exposed to the laboratory environment. This exposure also allowed radon daughters to accumulate on the internal surface of the AV in a significant way⁴. To prevent the migration of radon gas into the AV going forward, a newly-installed cover gas system aims to reduce by a factor of 10^5 the ingress of radon from laboratory air into the AV during

⁴Some effort was expended to devise an automated system to remove these embedded radon daughters; however, the decision by SNO+ to prioritize a $0\nu\beta\beta$ measurement led to the cessation of this effort. This surface contamination will not affect $0\nu\beta\beta$ and low-energy ^8B solar neutrino measurements.

operation. It consists of a series of buffer bags filled with high-purity nitrogen gas that are connected to the volume between the LS and a universal interface (UI) that seals the neck of the AV from the laboratory environment. The bags will accommodate changes in the mine air pressure to provide a robust physical barrier between the LS and the laboratory air.

Liquid scintillator

The SNO+ LS consists of a linear alkylbenzene (LAB) solvent containing 2 g/L of the fluor 2,5-diphenyloxazole (PPO). The choice of this LS is due mainly to its excellent optical properties and chemical compatibility with the AV, both critical for the experiment. In addition, the LS has a low flash point and is environmentally safe, making it suitable for use in the SNOLAB underground laboratory environment.

The de-oxygenated LS has a measured light yield of 11900 photons/MeV [53]. This response is linear down to 400 keV, below which the loss of Čerenkov light and ionization quenching (see chapter 5) lead to a reduction in the overall light yield. As with other organic liquid scintillators, the LS exhibits two distinct timing profiles depending on the type of particle interacting in the scintillator. Electrons produce mostly prompt light, with a measured decay constant near 4.5 ns, while alphas produce a long-lived scintillation component, with a measured decay constant close to 180 ns [54]. This allows for particle identification via pulse shape discrimination (PSD), where SNO+ expects to reject > 99.9% of alpha interactions, produced exclusively by backgrounds, while keeping > 99.9% of the electron interactions, such as those from ^{130}Te $2\nu\beta\beta$.

To investigate the $0\nu\beta\beta$ of ^{130}Te , SNO+ will load natural tellurium into the LS, with loading levels up to several percent demonstrated. SNO+ has investi-

gated two distinct methods to achieve this loading. In the preferred method, tellurium in the form of telluric acid $\text{Te}(\text{OH})_6$ is reacted with 1,2 butanediol to form a diol complex (Te-diol) that is then dissolved into the LS (TeLS). The other method also begins with $\text{Te}(\text{OH})_6$, but the acid is first dissolved in water before being added to the LS using the surfactant isopropylamine dodecylbenzene sulfonate [55] (commercially available as PRS)⁵. In both methods, adding 30 mg/L of the wavelength shifter 1,4-Bis(2-methylstyryl)benzene (bis-MSB) to the TeLS shifts the wavelength of the light emitted by the PPO to between 390–430 nm, closer to the PMT maximum quantum efficiency. This also reduces self-absorption of the scintillation light, which helps to improve the energy resolution of the SNO+ detector. This is particularly important as the Te loading decreases the absolute light yield of the TeLS relative to the LS. The Te-diol LS has a measured light yield of 6650 photons/MeV with 0.5% loading of Te-diol, by mass [56], whereas the surfactant loading method yields 9245 photons/MeV [53]. Alpha interactions also produce a lower proportion of late light in the TeLS, reducing the ability to perform PSD.

To minimize the contaminant level in the LS, a new LS processing and purification plant has been constructed underground near the SNO+ cavity [57], and the UPW production plant has been upgraded. Vacuum distillation, solvent-solvent water extraction, gas stripping using steam and nitrogen, and metal scavengers all contribute to removing heavy metal contaminants and ultra-fine particulates from the LAB and PPO. The gas stripping process also helps remove oxygen, which would otherwise act to quench the LS light output (see chapter 5). The water extraction process will also enable the removal

⁵This method is less favourable because of the significant challenge involved in purifying the PRS to the required radiopurity levels in the underground SNOLAB environment.

of Te (or Nd) from the LS for recovery of the experiment's $2\nu\beta\beta$ material⁶. Purification of the Te to remove heavy metal contaminants will occur before feeding it into the processing plant to add to the LS [58].

Calibration

The change from D₂O to LS also demands a new calibration program to characterize the SNO+ detector's behaviour. A suite of new or refurbished deployed radioactive sources, which table 2.1 lists, will determine the SNO+ energy scale and resolution, and linearity of the detector response in the energy range 0.1–6.1 MeV. They will also provide a check for determining systematic uncertainties in the software-based reconstruction of event energies and positions.

In addition to these radioactive sources, a deployable light-diffusing sphere (laserball) and Čerenkov source (formed by enveloping the SNO ⁸Li calibration source in an acrylic ball) will calibrate the PMT timing and gain, and will enable the decoupling of the Čerenkov and scintillation light output of the LS in the reconstruction of events. A series of LED and laser-based optical sources are also distributed uniformly on the PSUP [59]. As they are permanently in place, this allows for their frequent use. It also mitigates the risk of contaminating the detector inherent in the deployment of external media inside the detector's active volume (and which limits the frequency with which SNO+ will be able to perform deployed calibrations). These optical sources enable *in situ* measurements of the scattering and absorption lengths of the LS as a function of time.

Substantial changes to the SNO calibration systems are necessary to meet

⁶This is important in the event SNO+ decides to change the $2\nu\beta\beta$ isotope to confirm any signal observed with the Te, or to remove the $2\nu\beta\beta$ isotope to proceed with pure LS running and pursue further solar neutrino physics goals.

Source	Output	Energy (MeV)	Phase
^{16}N	γ	6.1	H_2O
^{24}Na	γ	4.1 (sum)	LS & TeLS
^{48}Sc	γ	3.3 (sum)	LS & TeLS
^{57}Co	γ	0.122	LS & TeLS
^{60}Co	γ	2.5 (sum)	LS & TeLS
$^{220,222}\text{Rn}$	α, β, γ	various	LS & TeLS
AmBe	n, γ	2.2, 4.4 (γ)	H_2O & LS & TeLS

Table 2.1: SNO+ will rely on a series of radioactive calibration sources to determine the energy scale and resolution, and to test the linearity of the detector response. The planned sources cover the entire energy range of interest of the experiment. The ^{60}Co calibration source is the subject of chapter 4.

the strict SNO+ radiopurity requirements and to ensure the chemical compatibility of all calibration hardware materials with the LS. The source deployment system, including a universal interface between the AV and the laboratory, has been redesigned according to vacuum standards to limit the ingress of radon from the lab air into the detector volume. A sealed source container will store calibration sources when not in use, ensuring the surface cleanliness of calibration sources is maintained. It will connect to the main detector volume via a sealed port, again limiting radon ingress into the detector. Calibration sources will connect to an umbilical that allows their deployment inside the AV, and which allows for communication of the source with the laboratory (such as by passing high voltage, gas, or light via fibre optic cables to the sources, as required, and by reading responses from sources that host a PMT). Each source connects to the umbilical via a custom-designed source connector that simplifies the process of attaching the two parts inside a glove box, with its finite space to maneuver and limited visibility.

Six cameras in watertight enclosures will also be distributed uniformly across the PSUP [60]. The cameras will enable the accurate and precise tri-

angulation of the position of calibration sources inside the detector. They will also complement the load cells on the rope-net arms in monitoring the position of the hold-down rope net by visually verifying any shifts of the rope net during the experiment's lifetime. Figure 2.2 shows a picture taken of SNO+ by a camera mounted on the bottom half of the PSUP.

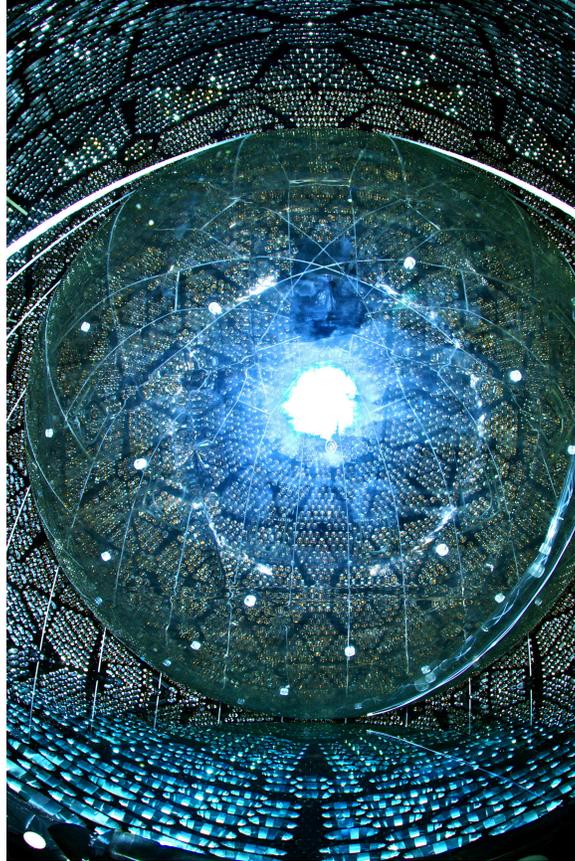


Figure 2.2: SNO+ is visible in this photo taken by one of the calibration system cameras mounted to the bottom of the PSUP. The water level is just below the bottom of the AV, recognizable by the reflections of PMTs on the water surface. The rope net is also visible as the white lines surrounding the AV. From the point of view of the camera, the AV neck is pointing backward, just above the large reflection of the camera light on the AV surface. The structure of the rope net surrounding the AV neck is visible.

Simulation

A complete Geant4-based [61, 62] MC simulation, RAT (RAT is an Analysis Tool), enables the simulation and reconstruction of physics events in the SNO+ detector. The simulation includes the full detector geometry and simulates particle interactions with the detector. It also simulates the detector optics, including the production of scintillation and Čerenkov light, the attenuation of those optical photons in the LS, acrylic and water, and the production of photoelectrons in the PMTs. Radioactive and optical calibration sources are also modelled. In addition, RAT simulates the entire SNO+ electronics chain, including the PMT response, trigger and data acquisition (DAQ). It handles data and analysis using ROOT [63].

RAT simulates all types of physics events relevant to the SNO+ physics program (see section 2.2). In particular, event generators exist to simulate the radioactive decays of all expected backgrounds, including the $0\nu\beta\beta$ and $2\nu\beta\beta$ decays of a number of isotopes, specifically ^{130}Te . The nuclear decays have been verified against existing data [64, 65] to ensure accuracy in the production of associated gamma-rays, electrons and alphas (see section 4.2.4).

The results of laboratory measurements constantly feed back to the simulation to ensure the simulated detector responds according to the best knowledge of its components. For example, the optical properties of the LS, like light yield and scattering length, are updated as laboratory measurements are refined and the LS components are tweaked. This information, as well as information regarding the detector's operational state (for example, which PMTs were operating during a particular data collection run) are read from a dedicated database to accommodate the evolutionary nature of these values.

2.2 Physics Program

SNO+ is a general-purpose detector that has the ability to fulfill a broad physics program, including measuring solar and supernova neutrinos, reactor and geo anti-neutrinos, and searching for other rare processes, like nucleon decay [66]. Its main focus, however, will be on investigating the $0\nu\beta\beta$ of ^{130}Te . As a large liquid scintillator experiment, SNO+ benefits from the ability to load a large active mass of the $2\nu\beta\beta$ isotope into the detector volume to obtain high sensitivity to observing a potential $0\nu\beta\beta$ signal. Its size enables fiducialization of the detector volume to limit surface and other external backgrounds. Spectral fitting of a large statistical sample of events (for the entire physics program) helps compensate for the relatively poor energy resolution from which liquid scintillator experiments inherently suffer.

2.2.1 Supernova and solar neutrinos

In the event a core-collapse supernova occurs within the galaxy, SNO+ will be in a position to measure a strong neutrino signature [67]. A supernova 10 kpc from earth releasing 3×10^{53} erg of energy in the form of neutrinos would result in nearly two hundred electron anti-neutrino interactions via inverse beta decay within SNO+. A similar number of events would occur from neutrino-proton elastic scattering above a low energy threshold of 200 keV. Neutral current interactions with ^{12}C nuclei may provide information regarding the total flux of all neutrino flavours. SNO+ plans to participate in the Supernova Early Warning System (SNEWS) [68] by providing a semi-automated response to a supernova-like event burst.

Like its forebear, SNO+ also has the ability to make significant contribu-

tions to solar neutrino research. The SNO+ low energy threshold may enable the study of neutrinos produced in the pp, pep and CNO solar fusion chains. Neutrinos from the pep reaction, with an energy of 1.44 MeV, fall in the transition region between vacuum- and matter-dominated neutrino oscillations where the MSW large mixing angle solution is generally assumed. A precise measurement of the pep neutrino flux could reveal new physics [69]. The Borexino experiment [70, 71] has measured the pep neutrino flux, but ^{11}C backgrounds have made it difficult for them to make a precise statement. SNO+ expects a two order of magnitude reduction in the ^{11}C rate relative to Borexino, which, assuming SNO+ achieves all other background levels at the levels Borexino achieved⁷, may allow SNO+ to measure the pep rate to better than 10% with 1 year of data.

Obtaining a precise measurement of the CNO neutrino fluxes would provide insight to the discrepancy between helioseismic and photospheric measurements of the solar metallicity [72, 73]. This measurement will be highly challenging because of the background rates at low energies. The same is true for measuring neutrinos from the pp chain. Assuming sufficiently low backgrounds, SNO+ may be able to obtain a 15% measurement of the CNO neutrino flux and a few percent measurement of the pp flux. SNO+ expects to match the Borexino experiment, which has also set upper limits on the CNO neutrino flux [71] and has also made the first direct measurement of pp neutrinos [74]. Figure 2.3 shows the expected MC energy spectrum of solar

⁷There is a possibility SNO+ may not initially achieve Borexino-level backgrounds in the low-energy region because of the radon daughters that have accumulated on the inner surface of the AV. This could lead to ^{210}Bi levels in the scintillator that are too high for a pep or CNO solar neutrino measurement without further mitigation, say from enhanced online scintillator purification. ^8B solar neutrino measurements are not affected by these backgrounds.

neutrino signals and various backgrounds in the energy range 0.3–6.5 MeV. ^8B and ^7Be neutrinos are also measurable to approximately 7% and 4% with 1 year of data.

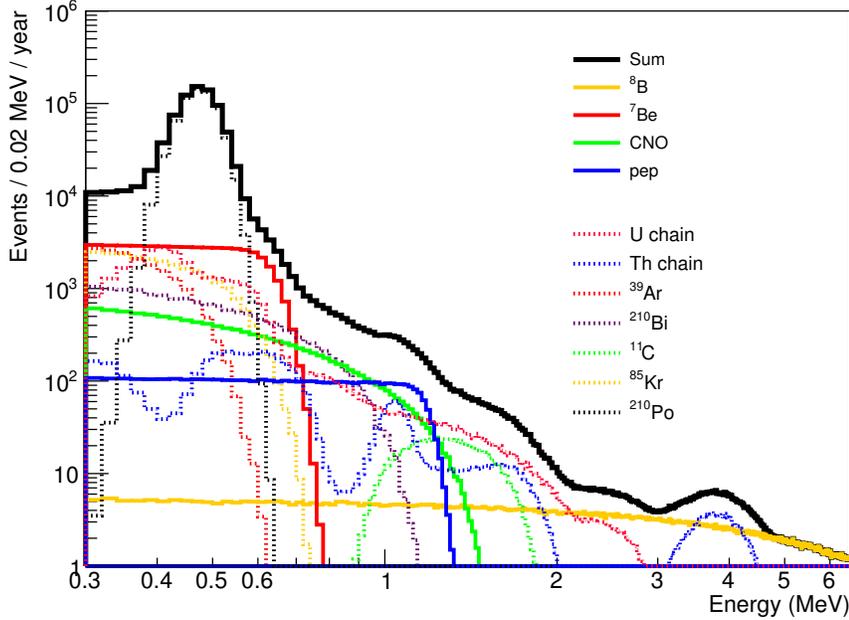


Figure 2.3: Expected MC solar neutrino (solid lines) and background (dashed lines) energy spectra in the energy range 0.3–6.5 MeV in SNO+, for events in LAB-PPO LS assuming 400 Nhits/MeV and a fiducial volume cut at 5.5 m. Solar neutrino fluxes are extracted from this data via an extended maximum likelihood fit.

2.2.2 Reactor and geo anti-neutrinos

SNO+, like other liquid scintillator detectors, is sensitive to anti-neutrinos through the process of inverse beta decay. The 200 μs time coincidence between prompt light from positron annihilation and light from gamma-ray interactions following the subsequent neutron capture on H will allow SNO+ to identify electron anti-neutrino events. This may even be possible above the large $2\nu\beta\beta$

background⁸ when the scintillator is loaded with ^{130}Te .

SNO+ sits in the vicinity of a sufficient population of nuclear reactors and expects to observe approximately 90 oscillated reactor anti-neutrino events per year of operation. A reactor at a 240 km baseline (Bruce) and a set of reactors at an approximately 340 km baseline (Darlington and Pickering) will give SNO+ a distinct oscillation-distorted reactor anti-neutrino spectrum. Other reactors in the area will also provide a sample of events, though with a less distinct spectral shape. Two oscillation-induced dips in the electron anti-neutrino energy spectrum, as shown in figure 2.4, will give SNO+ similar sensitivity to that of KamLAND [75] to the parameter Δm_{12}^2 , at $0.2 \times 10^{-5} \text{ eV}^2$ with seven years of data collection in the TeLS phase.

A main contribution to earth's internal heat generation is predicted to be the decay of long-lived radioisotopes in the ^{235}U , ^{238}U and ^{232}Th decay chains and the beta decay of ^{40}K . The abundances of these isotopes in the earth's crust and mantle are model-dependent, and estimates of the mantle composition are typically inferred by non-geophysical means, often using simplified assumptions [76]. SNO+, like Borexino [77] and KamLAND [78] before it, may directly measure the combined U/Th/K crust and mantle composition of the earth by observing the anti-neutrino signal emitted alongside these decays. Having extensively characterized local geology may allow SNO+ to infer the contribution to the geo anti-neutrino signal from the mantle by subtracting the expected crustal component. Also, as figure 2.4 shows, a large part of the geo anti-neutrino energy spectrum lies below the reactor anti-neutrino spectrum, giving SNO+ some power in distinguishing between them.

⁸Given 1.3 t of ^{130}Te with a half-life of 8.2×10^{20} years, we expect nearly 5×10^6 ^{130}Te decays within SNO+ per year.

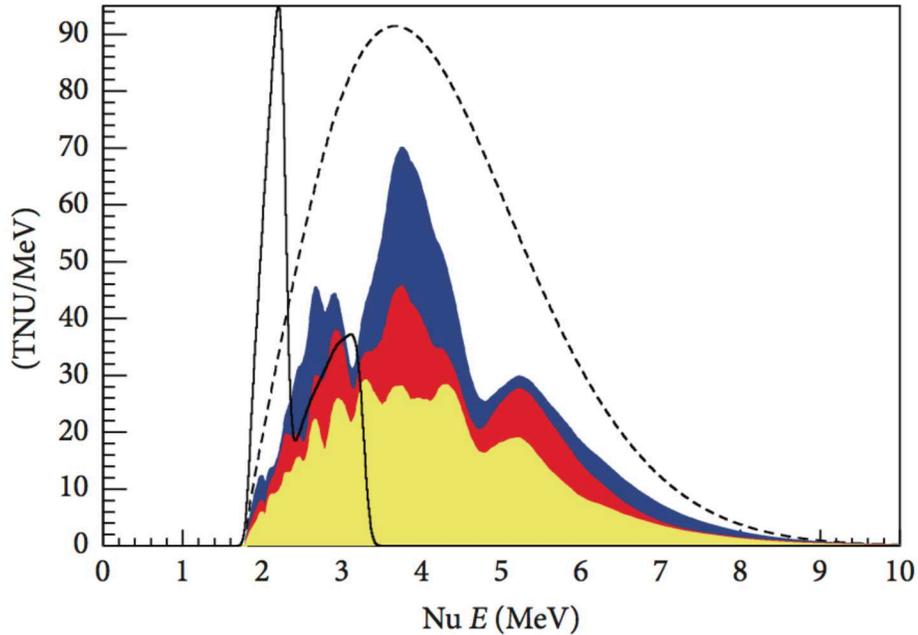


Figure 2.4: Shown is the expected reactor anti-neutrino spectrum for 10^{32} proton-years/MeV assuming a light yield of 300 Nhits/MeV for the TeLS phase and a fiducial volume cut of 5.5 m. SNO+ expects two distinct oscillation-induced dips in its observed reactor anti-neutrino spectrum, shown as the stacked blue, red and yellow regions (corresponding to the Bruce reactor, Pickering and Darlington reactors, and all other reactors contributions). The dashed line is the expectation assuming no oscillations. The solid line shows the expected geo anti-neutrino spectrum, which is clearly visible above the reactor anti-neutrino spectrum at energies below 2.5 MeV.

2.2.3 Neutrinoless double beta decay

Investigating the Majorana nature of the neutrino by searching for the $0\nu\beta\beta$ of ^{130}Te is the focus of the SNO+ experiment. ^{130}Te is a good choice of $2\nu\beta\beta$ isotope⁹ because it has a high natural abundance of 34.1%, so does not re-

⁹The original isotope chosen by SNO+ was ^{150}Nd . However, its low natural abundance of 5.6% and the inability to effectively enrich natural Nd prompted the SNO+ collaboration to investigate alternative isotopes, leading to the choice of ^{130}Te after nearly two years of investigation into the properties of Te-loaded scintillator and weighing its benefits relative to those of the already-established Nd-loaded scintillator.

quire enrichment to obtain it in large quantities, has a relatively high Q-value of 2.53 MeV, placing it above many naturally-present backgrounds, and is the longest-lived $2\nu\beta\beta$ isotope with a measured half-life of $(8.2 \pm 0.6) \times 10^{20}$ years [79], which reduces the background contribution from its $2\nu\beta\beta$ decay. Also, the SNO+ LS (see section 2.1.2) remains optically clear after Te-loading, which makes possible loadings of several percent.

SNO+ plans to use a Te loading of 0.5% by mass, corresponding to approximately 1.3 t of ^{130}Te in the detector. Inherent backgrounds in the TeLS will be reduced through a multi-pass purification process, while external backgrounds, including U/Th chain daughters from the PMTs, AV, water shielding and ropes may be reduced through coincidence tagging and cutting on the fiducial volume. The ^{130}Te $2\nu\beta\beta$ decay and ^8B solar neutrino backgrounds are irreducible. A two-pass purification process, once above and once below ground, will reduce backgrounds produced via cosmogenic activation of the natural Te during its exposure on surface to negligible levels, with a demonstrated reduction in the levels of ^{60}Co by a factor of 2.7×10^{-6} . Cosmogenic activation of the LS while underground is negligible.

Figure 2.5 shows the resulting expected background energy spectrum after 1 year of data collection and within a 3.5 m diameter fiducial volume. It also shows for reference the expected $0\nu\beta\beta$ signal assuming an effective Majorana neutrino mass, $m_{\beta\beta}$, of 200 meV/ c^2 with matrix element and phase-space factor from [37, 34]. The number of events, including a possible $0\nu\beta\beta$ signal, observed in an asymmetric energy region of interest (ROI) near the Q-value

of the ^{130}Te $2\nu\beta\beta$ sets the half-life, $\tau_{1/2}^{0\nu}$, of the ^{130}Te $0\nu\beta\beta$ decay as

$$(m_{\beta\beta})^{-2} \propto \tau_{1/2}^{0\nu} = \frac{\epsilon N_{^{130}\text{Te}} \ln(2) T}{\sigma \sqrt{(bM + C) \cdot T \delta E}}, \quad (2.1)$$

where ϵ is the detector efficiency, $N_{^{130}\text{Te}}$ is the number of atoms of ^{130}Te in the detector, T is the data collection time, M is the mass of ^{130}Te in the detector, b is a mass-dependent background factor (which scales with isotope loading level), C is a constant background term (backgrounds that do not scale with loading, like external backgrounds and ^8B neutrinos), δE is the width of the energy ROI, and σ is the Gaussian significance level being calculated.

For the scenario in figure 2.5, SNO+ will measure the ^{130}Te $0\nu\beta\beta$ half-life to be $> 8 \times 10^{25}$ years with one year of data collection, corresponding to $m_{\beta\beta} = 75 \text{ meV}/c^2$. This begins to probe the inverted hierarchy of neutrino masses (see section 1.3.1). With higher loadings and a longer data collection period, SNO+ may be able to reach well into the inverted hierarchy region.

2.3 Experiment phases and timeline

The SNO+ experiment will progress through several distinct phases leading up to the $0\nu\beta\beta$ measurement, each of which has specific goals. In the first phase, the AV will be filled with UPW, allowing for commissioning of various detector components, like the data acquisition system. Searches for exotic physics and studies of backgrounds external to the AV will also occur in this phase. The detector is being filled with UPW now and will be ready to take data before the end of 2016.

In the second phase, the AV will be filled with LS, the UPW acting as a

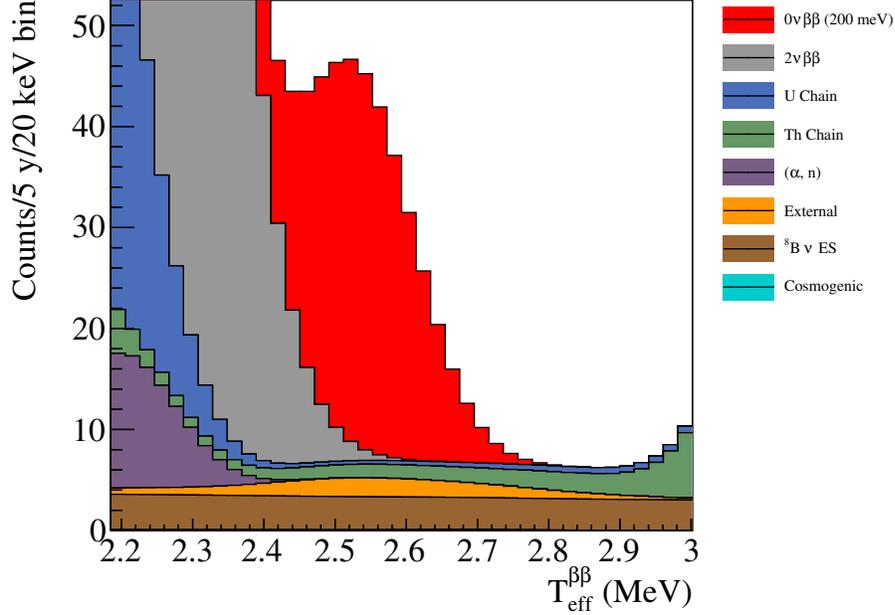


Figure 2.5: Averaged profiles of the MC background energy spectra for 1 year of data collection at 0.5% loading assuming 390 Nhit/MeV, a 20% fiducial volume cut, 99.99% efficient tag for ^{214}Bi , 97% efficient tag for ^{208}Tl and 98.8% efficient tag for ^{212}Bi . The energy spectrum of a $0\nu\beta\beta$ signal assuming an effective Majorana neutrino mass of 200 meV/ c^2 is shown for reference. The energy scale $T_{\text{eff}}^{\beta\beta}$ compensates for the effect of quenching (see section 5.1) by scaling all spectra so the $0\nu\beta\beta$ endpoint lies at its true value of 2.53 MeV.

buffer fluid during the fill. During this phase, calibration data from optical and radioactive sources will help to understand the detector response to the LS and verify the detector’s optical model. Both external and internal backgrounds, including those inherent in the LS, will be measured. Low energy solar neutrinos and reactor and geo-antineutrinos will be studied. LS will start flowing into the AV at the beginning of 2017, with data collection beginning during the summer of 2017.

In the last phase, the LS will be loaded with natural tellurium in the effort to observe the $0\nu\beta\beta$ decay of ^{130}Te . Reactor and geo-antineutrinos will

continue to be studied during this phase, though backgrounds are expected to be too high to do meaningful solar neutrino studies (except perhaps for high energy solar neutrinos, like ^8B). In all phases, the detector will observe for supernovae. Tellurium loading is expected to begin before the end of 2017, and five years of data collection are planned.

Chapter 3

SNO+ sensitivity systematics

They are ill discoverers that think there is no land, when they can see nothing but sea.

- Francis Bacon, *The Advancement of Learning, Book II, vii, 5*

While the focus of the SNO+ experiment has from the outset been to investigate the Majorana nature of the neutrino via an observation of $0\nu\beta\beta$ (see section 2.2.3), the collaboration did not make the definitive choice to load the liquid scintillator with ^{130}Te until 2013. Prior to this, the collaboration was well on its way toward selecting ^{150}Nd as its choice isotope. The properties of a ^{150}Nd -loaded liquid scintillator were already established, but the benefits to switching to ^{130}Te simply outweighed those of remaining on the ^{150}Nd track¹.

The initial experiment called for loading the LS with 0.1% (or, potentially, up to 0.3% [80]) natural neodymium by mass², corresponding to approximately 43.7 kg of ^{150}Nd . Because the Q-value of the $2\nu\beta\beta$ is rather high at 3.37 MeV,

¹Other slippages in the experiment's timeline afforded the collaboration the time to thoroughly weigh the benefits of this switch without impacting the overall experiment timeline.
² ^{150}Nd has a natural abundance of 5.6%.

there are few backgrounds that will act to shadow a $0\nu\beta\beta$ signal. Of the backgrounds, the largest is the $2\nu\beta\beta$ of ^{150}Nd itself, either to the ground state or an excited state of ^{150}Sm . Internal backgrounds include ^{214}Bi and ^{208}Tl , daughters in the decay chains of ^{238}U and ^{232}Th , respectively. While much of these stem from the period of time when the detector was in contact with radon-laden mine air, they are also present in trace amounts in natural neodymium (though we reduce these levels through a process of neodymium purification [81]). There is also an irreducible background contribution resulting from the elastic scattering of ^8B solar neutrinos. Event pileup, where two or more uncorrelated decays deposit energy in the detector in a short time window, may also lead to an energy response by the detector that mimics an event near the ^{150}Nd endpoint [82].

Although the SNO+ experiment has yet to begin its $2\nu\beta\beta$ physics data collection phase, it is important to estimate *a priori* the experiment's sensitivity to observing a $0\nu\beta\beta$ signal. There is an entire internal group dedicated to doing just this³; however, my interest in investigating the experimental sensitivity is in understanding the requirements for calibrating the SNO+ energy response. Unknown systematic errors in the energy scale or energy-dependent resolution (including a high-energy tail in the resolution⁴) may have a dramatic effect on the experimental sensitivity, either degrading it to a level where the experiment is no longer able to meet its physics goals, or artificially enhancing it by introducing a falsely positive signal. How large these effects are informs us as to the precision and accuracy any calibration source must provide.

³[83] outlines early work by this group relevant to ^{150}Nd .

⁴This could be physically manifest as a misunderstanding of the light production process in the liquid scintillator, for example, as a result of quenching of the scintillation process for low energy electrons that contribute to the $0\nu\beta\beta$ signal (see chapter 5).

To complete this study, I performed a MC simulation, borrowing from code already available from the $0\nu\beta\beta$ group⁵, to calculate the sensitivity to observing a ^{150}Nd $0\nu\beta\beta$ signal while including various systematic effects in the detector energy response. Section 3.1 outlines the method I used to estimate the sensitivity and compares my results with those of the $0\nu\beta\beta$ group⁶. The two methods are complementary and the results from each are consistent. I present the results of my systematics study in section 3.2. This extension to the work of the $0\nu\beta\beta$ group helps to inform us of the size of the systematic errors we may tolerate on a calibrated energy response.

As a final note, while investigating the sensitivity to a ^{150}Nd $0\nu\beta\beta$ signal may appear moot given the collaboration’s choice of ^{130}Te , the method and conclusions are applicable to *any* choice of $2\nu\beta\beta$ isotope (although the specific results may change owing to a different signal combined with an alternate set of backgrounds). In addition, because SNO+ anticipates the ability to remove the dissolved tellurium from the liquid scintillator, the possibility that ^{150}Nd could be deployed in a future phase of the experiment cannot be discounted completely. This becomes especially relevant in the event of an observation of $0\nu\beta\beta$ during the ^{130}Te phase, as verification of this signal using a different isotope could become paramount.

⁵In rewriting the code, I corrected several errors from which the original code suffered. The result was a MC simulation that behaved in a more predictable manner. Section A.1 outlines tests I performed to verify the rewrite performed appropriately and discusses in more detail how the original code failed.

⁶My method relies on using the shape of the log-likelihood function of a fit parameter to estimate the sensitivity [3, 84], whereas the $0\nu\beta\beta$ group uses a method based on the CL_S technique [85, 86].

3.1 ^{150}Nd $0\nu\beta\beta$ sensitivity

The method I use [87] to evaluate the sensitivity of SNO+ to detecting a $0\nu\beta\beta$ signal relies on fitting a pre-determined MC energy spectrum that includes all known detector effects to the energy spectrum of all observed events (the data). To begin, I create a bin-wise probability density function (PDF) that describes the total SNO+ energy response by summing the contribution from each background component and a $0\nu\beta\beta$ signal component. I then scale each background component to the number of events of that type expected over the data collection period and scale the signal component to one event. I also include the effect of the detector's energy-dependent resolution, assuming a scintillator response of 400 Nhits/MeV (corresponding to a 6.4% full-width half maximum (FWHM) at 3.37 MeV), by numerically convolving the energy response with a Gaussian distribution, where

$$G(E) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(E-E_o)^2/2\sigma^2}$$

$$\sigma(E_o) = \sqrt{E_o/400} + 0.001 , \quad (3.1)$$

with E the energy and E_o the energy at the centre of the bin I am convolving. I also include a 0.1% systematic broadening of the resolution σ .

The number of events per bin k in the PDF is

$$\nu_k = \sum_{i=1}^N \alpha_i \varphi_i^k , \quad (3.2)$$

where N is the total number of contributors to the PDF, φ_i is the bin-wise energy distribution of background or signal contribution i (with all resolution and systematic effects folded in) and α_i is a free parameter that allows the

contribution from each component to float. The total number of events in the PDF becomes

$$\nu_{tot} = \sum_{k=1}^n \nu_k , \quad (3.3)$$

where n is the total number of bins in the PDF.

The backgrounds I consider include the internal backgrounds ^{214}Bi and ^{208}Tl , the $2\nu\beta\beta$ of ^{150}Nd to both the ground state and an excited state of ^{150}Sm (which I denote $2\nu\beta\beta+2\nu\beta\beta^*$), and the elastic scattering interaction of ^8B solar neutrinos, assuming flavour oscillations. I consider both decays of ^{150}Nd together as a single background, since they will be indistinguishable (or very nearly so) in the data. I do not consider the effect of pileup ([82] studies the effect of pileup on the $0\nu\beta\beta$ sensitivity). Table 3.1 lists these backgrounds, along with the endpoints of their energy distributions, the total number of events of each we expect to see per kilotonne per year of running and the total number of events we expect to see each year. This last value assumes we cut on the fiducial volume at $0.8R$, where R is the radius of the SNO+ acrylic vessel, cut on the effective running time as $0.8T$, where T is the time since the beginning of the data collection period, and remove 90% of the ^{214}Bi events through tagging.

To calculate the number of expected events, I assume a contaminant level of $1.60 \times 10^{-17} \text{g/g}_{\text{LAB}}$ and $1.0 \times 10^{-15} \text{g/g}_{\text{Nd}}$ for ^{238}U , and a level of $6.8 \times 10^{-18} \text{g/g}_{\text{LAB}}$ and $1.0 \times 10^{-14} \text{g/g}_{\text{Nd}}$ for ^{232}Th [88]. Assuming we use 780 kt of LAB loaded with 0.1% by mass of natural neodymium, then the detector will contain approximately 43.7 kg of ^{150}Nd . I use a ^{238}U half-life equal to 4.468×10^9 years and a ^{232}Th half-life equal to 1.405×10^{10} years, and assume

the decays of ^{238}U and ^{214}Bi , and ^{232}Th and ^{208}Tl are in secular equilibrium to calculate the expected number of each of those respective backgrounds. To calculate the expected number of decays from ^{150}Nd , I use 9.11×10^{18} years [89] and 1.33×10^{20} years [90] for the half lives of the ^{150}Nd decay to the ground and excited states of ^{150}Sm , respectively. Using these lifetimes, the branching ratio for the decay to the ground state is 0.935895 and the branching ratio for the decay to the excited state is 0.064105. Table 3.1 lists the expected numbers of each type of event per year, while figure 3.1 shows the PDF, made using the values in the last column of table 3.1.

Background	Endpoint (MeV)	Events/year/kt	# Events in PDF
^{214}Bi	3.367	6671.67	213.15
^{208}Tl	4.99	774.4	247.4
^8B	15	1505.5	480.98
$2\nu\beta\beta+2\nu\beta\beta^*$	3.37	18287287.27	5842568.84
$0\nu\beta\beta$	3.37	0	1

Table 3.1: The backgrounds I consider with their endpoint energies and the number of events of each type we expect to see per kilotonne per year of running. The last column contains the number of events per year of running I used to create the PDFs, assuming a fiducial volume cut of $0.8R$ and effective run-time cut of $0.8T$, and assuming that we will be able to remove 90% of the ^{214}Bi background through tagging.

I next produce MC fake data by allowing the contribution of each background component in each bin to fluctuate as a Poisson random variable, while including zero contribution from the $0\nu\beta\beta$ component. I then perform an unconstrained fit of the PDF to the fake data in the energy range 2.5 MeV to 5.0 MeV by minimizing the negative log-likelihood, $-\ln(L)$, of the fit. Figure 3.2 shows an example of such a fit for one year of data collection. I expect each α_i to be one for background components and zero for the $0\nu\beta\beta$ component, call it $\alpha_{0\nu}$. By setting the signal to zero in the data, but not the PDF, I obtain

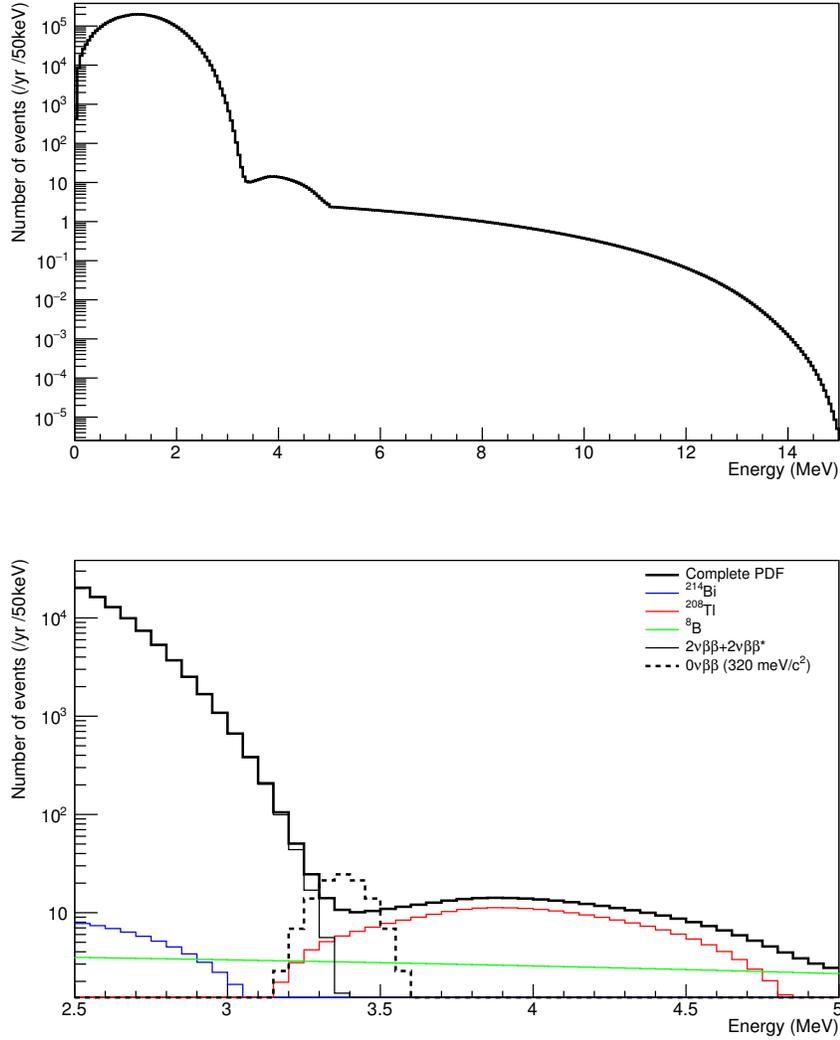


Figure 3.1: The complete PDF has contributions from $2\nu\beta\beta+2\nu\beta\beta^*$, ^{208}Tl , ^{214}Bi and ^8B , each scaled to its expected number of events, and the $0\nu\beta\beta$ signal, scaled to one event. The PDF includes the energy resolution given by equation 3.1. The top plot shows the PDF over the entire energy range (0–15 MeV), while the bottom plot shows the complete PDF and its components over the fitting range (2.5–5 MeV). The bottom plot shows a $0\nu\beta\beta$ signal assuming a $320 \text{ meV}/c^2$ neutrino mass (this is scaled to 1 in the PDF).

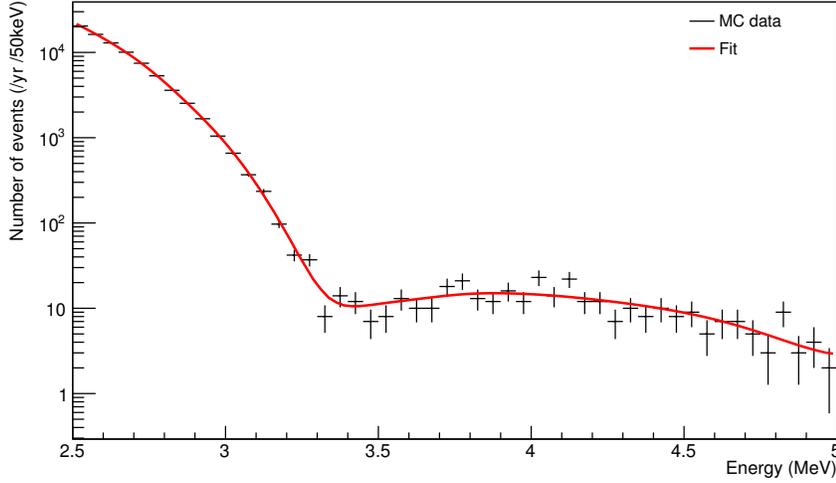


Figure 3.2: The fit to MC generated data in the energy range 2.5 MeV to 5.0 MeV after one year of data collection (for a single experiment) results in the value of α_i for ^{214}Bi of -10.21, ^{208}Tl of 0.9965, ^8B of 1.102, $2\nu\beta\beta+2\nu\beta\beta^*$ of 1.005, and $0\nu\beta\beta$ of 0.7873.

the best limit that may be set on the number of observed signal events (the value of $\alpha_{0\nu}$ the fit returns) if indeed SNO+ sees no $0\nu\beta\beta$ events.

To estimate the 90% confidence limit on the number of observed signal events, I add 1.645 times the upper 1σ error on the parameter $\alpha_{0\nu}$ (the point where $-\ln(L)$ increases by 0.5), call this $\Delta_{\alpha_{0\nu}}$, to $\alpha_{0\nu}$. To obtain the 90% confidence limit on the $0\nu\beta\beta$ half-life, $\tau_{1/2}^{0\nu}$, and effective neutrino mass, $m_{\beta\beta}$, I calculate the half-life as

$$\tau_{1/2}^{0\nu} = \frac{M_{150\text{Nd}} N_A \ln(2) T}{m_A (\alpha_{0\nu} + 1.645 \Delta_{\alpha_{0\nu}})}, \quad (3.4)$$

where $M_{150\text{Nd}}$ is the total mass of ^{150}Nd in the detector (43.7 kg), m_A is the atomic mass of ^{150}Nd , N_A is Avogadro's number, T is the total running time (in years, and not the effective running time) and $(\alpha_{0\nu} + 1.645 \Delta_{\alpha_{0\nu}})$ is the 90% upper confidence limit of the number of observed signal events. The effective

neutrino mass becomes

$$m_{\beta\beta} = \frac{m_e}{|M_{m_\nu^{eff}}| \sqrt{G_{01} \tau_{1/2}^{0\nu}}}, \quad (3.5)$$

where m_e is the electron mass, the nuclear matrix element $|M_{m_\nu^{eff}}|$ is 2.5 and the integrated kinematic factor G_{01} is $2.69 \times 10^{-13} \text{ yr}^{-1}$ (using the interacting boson model (IBM-2)) [91].

These confidence limits are the result of a single experiment and, as such, are just one measurement, subject to the random fluctuations inherent in any counting experiment, of the true confidence limit. To investigate the statistical spread in the 90% confidence limits on $\tau_{1/2}^{0\nu}$ and $m_{\beta\beta}$, I repeat this fitting process 50000 times, with each fake data set being distinct, but the PDF remaining the same each time. Figures 3.3 and 3.4 show the distributions of the fit results, α_i and Δ_{α_i} , for each of the backgrounds listed in table 3.1. If I perform this experiment a sufficiently large number of times, then I expect $\alpha_{0\nu}$ (figure 3.5) to be normally distributed about zero, barring any bias in the fit. Figure 3.6 shows there is a correlation between the fit number of $0\nu\beta\beta$ events ($\alpha_{0\nu}$) and the upper error in that number ($\Delta_{\alpha_{0\nu}}$). An upper fluctuation of $\alpha_{0\nu}$ causes $\Delta_{\alpha_{0\nu}}$ to increase, whereas a downward fluctuation in $\alpha_{0\nu}$ means a decrease in $\Delta_{\alpha_{0\nu}}$ (noting that I do not constrain the total number of events ν_{tot} in the fit). Because the correlation seems to be linear, with a correlation coefficient of 0.847, I also expect $\alpha_{0\nu} + \Delta_{\alpha_{0\nu}}$ to be normally distributed. Figure 3.7 shows this distribution of 90% confidence limits on the number of $0\nu\beta\beta$ events.

Transforming the fit number of $0\nu\beta\beta$ events⁷ to $\tau_{1/2}^{0\nu}$ and $m_{\beta\beta}$ using equa-

⁷Note that I do not disregard fits that return an unphysical (negative, say) number of events for any of the signals or backgrounds. Each fit returns one measurement of the true number of events of each type in the data. While this may be unphysical, on average we

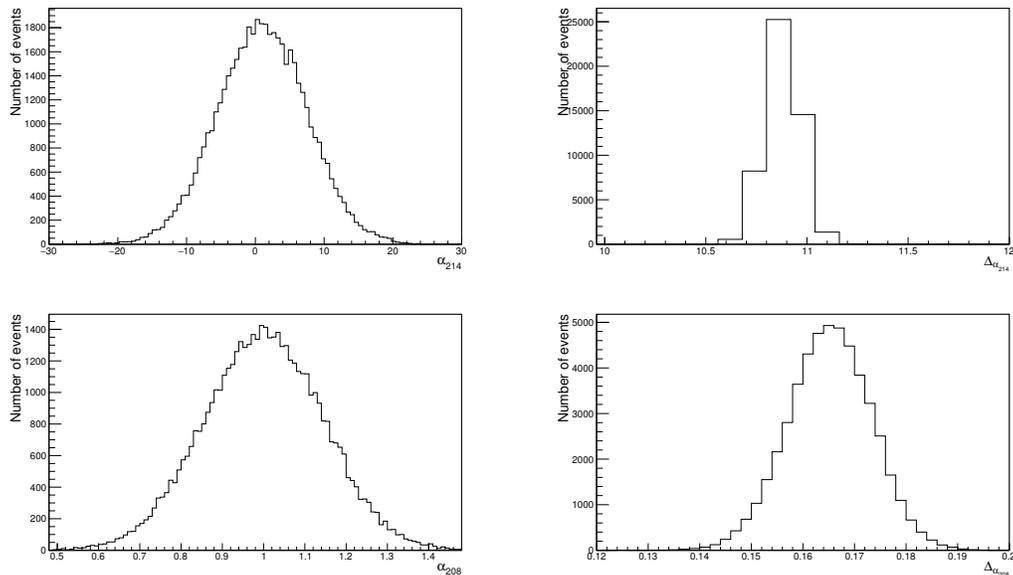


Figure 3.3: The fit values of the ^{214}Bi scaling parameter (α_{214}) and its upper 90% error ($\Delta\alpha_{214}$) (top), and the ^{208}Tl scaling parameter (α_{208}) and its upper 90% error ($\Delta\alpha_{208}$) (bottom), using one year of MC data. The mean value of α_{214} is 1.001 with the mean of $\Delta\alpha_{214}$ equal to 10.88. The mean value of α_{208} is 0.999 with the mean $\Delta\alpha_{208}$ equal to 0.165.

tions 3.4 and 3.5, respectively, results in the distributions in figures 3.8 and 3.9. Appendix A.2 contains a treatment of the mathematics of these transformations. Since the mean in figure 3.8 is approximately zero, the maxima are symmetric about zero and follow from equation A.5. Transforming the 90% confidence limit on the number of $0\nu\beta\beta$ events, as figure 3.7 shows, results in the distributions in figure 3.9, the top distribution showing the 90% confidence limits on $\tau_{1/2}^{0\nu}$ and the bottom distribution showing the 90% confidence limits on $m_{\beta\beta}$. Notice how the $\tau_{1/2}^{0\nu}$ distribution has quickly become asymmetric, as we expect from equation A.8. The $m_{\beta\beta}$ distribution has a shape matching that

expect the fit to return values for α_i that are physical and, after performing a sufficient number of experiments, we expect the mean value for each α_i to approach the true value. Thus, including negative values for $\tau_{1/2}^{0\nu}$ and $m_{\beta\beta}$ is not incorrect as long as, on average, they lie in a physical region.

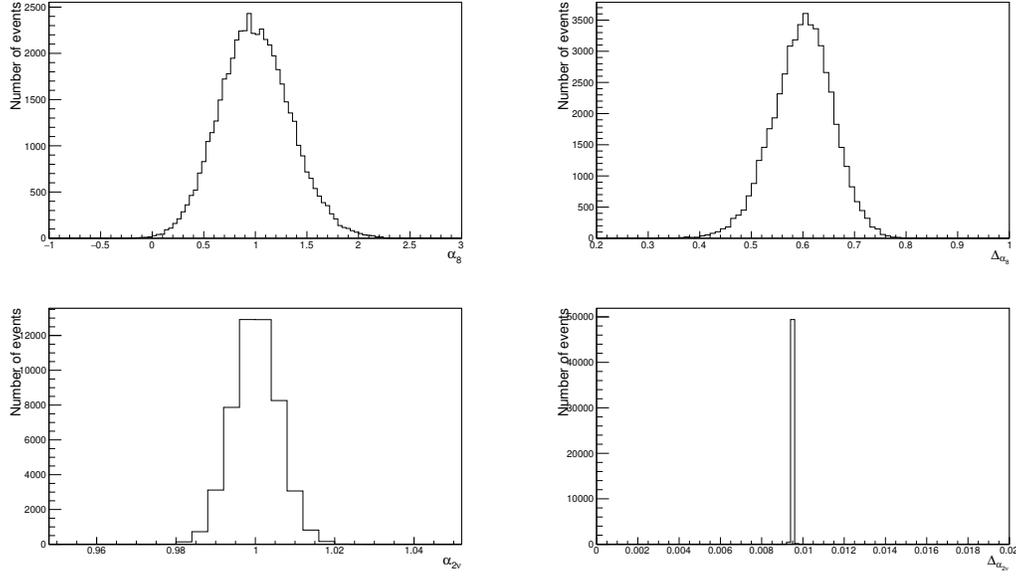


Figure 3.4: The fit values of the ${}^8\text{B}$ scaling parameter (α_8) and its upper 90% error (Δ_{α_8}) (top), and the $2\nu\beta\beta+2\nu\beta\beta^*$ scaling parameter ($\alpha_{2\nu}$) and its upper 90% error ($\Delta_{\alpha_{2\nu}}$) (bottom), using one year of MC data. The mean value of α_8 is 1.001 with the mean Δ_{α_8} equal to 0.599. The mean value of $\alpha_{2\nu}$ is 1.000 with the mean $\Delta_{\alpha_{2\nu}}$ equal to 0.009.

from equation A.19.

To obtain the 90% confidence limits on $\tau_{1/2}^{0\nu}$ and $m_{\beta\beta}$, I fit the spectra in figure 3.9 with equations A.3 and A.19, respectively, and choose the best value of each 90% confidence limit to be that at which the corresponding fitted function is maximal. To get the statistical 1σ and 2σ lower limit on $\tau_{1/2}^{0\nu}$ and upper limit on $m_{\beta\beta}$, I numerically solve equations A.18 and A.24. These limits are one-sided because the associated distributions are highly asymmetric.

Figures 3.10 and 3.11 show the results of this analysis, along with the official results of the $2\nu\beta\beta$ group, showing the relative agreement between the two methods (though the $2\nu\beta\beta$ group computes 2-sided statistical intervals).

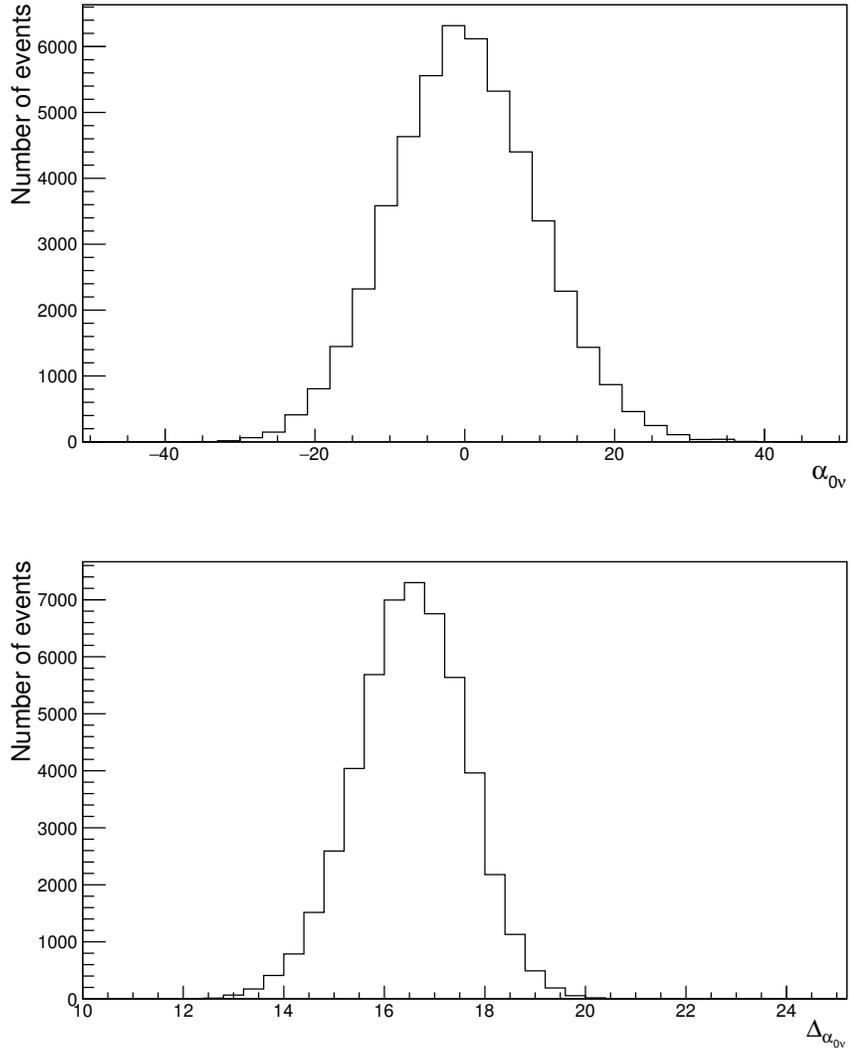


Figure 3.5: The fit values of the $0\nu\beta\beta$ scaling parameter ($\alpha_{0\nu}$) and its upper 90% error ($\Delta\alpha_{0\nu}$), using one year of MC data. The mean value of $\alpha_{0\nu}$ is 0.0434 and the mean value of $\Delta\alpha_{0\nu}$ is 16.53.

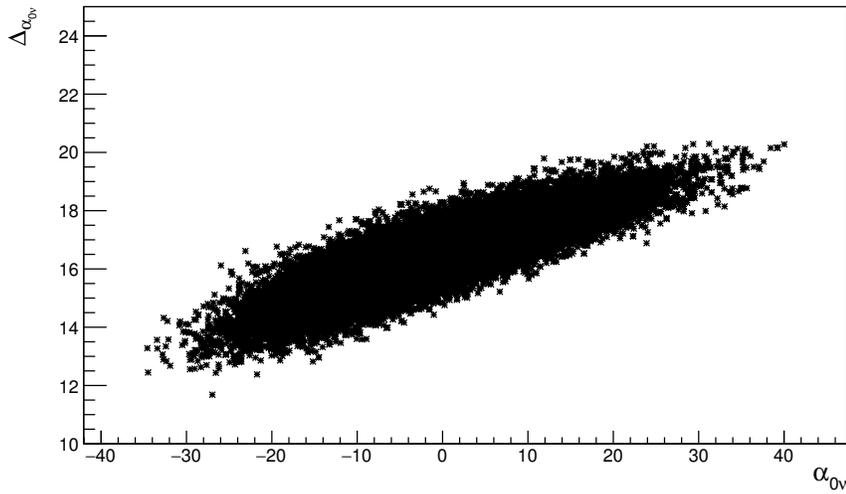


Figure 3.6: The correlation between $\alpha_{0\nu}$ and $\Delta_{\alpha_{0\nu}}$, from figure 3.5, with a correlation coefficient of 0.847.

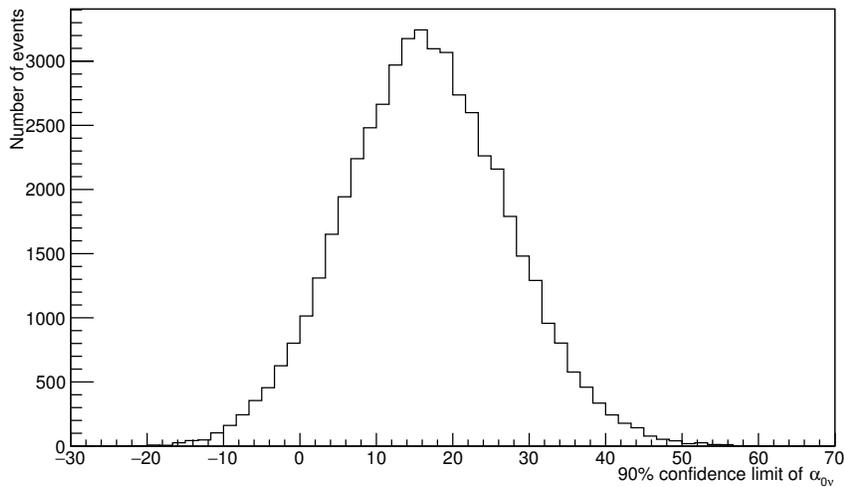


Figure 3.7: The 90% confidence limits on the number of fit $0\nu\beta\beta$ events maintains a Gaussian shape, but is shifted so the tail does not extend as far into the negative region as the distribution shown in figure 3.5.

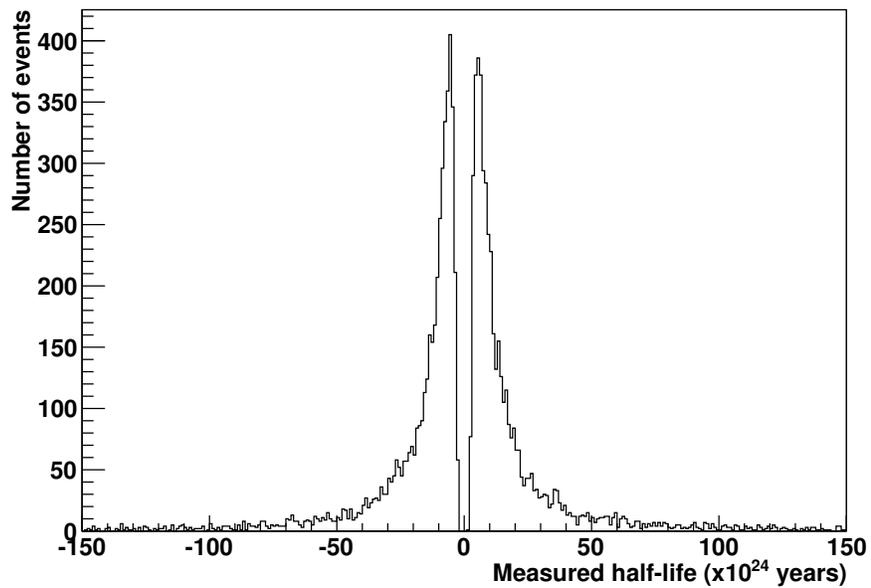


Figure 3.8: The distribution of half-lives from applying equation 3.4 to the fit values from figure 3.5. Notice the symmetric peaks about zero, which correspond to the value of $\alpha_{0\nu}$ where the distribution of $\alpha_{0\nu}$ falls to $1/e$ its maximum.

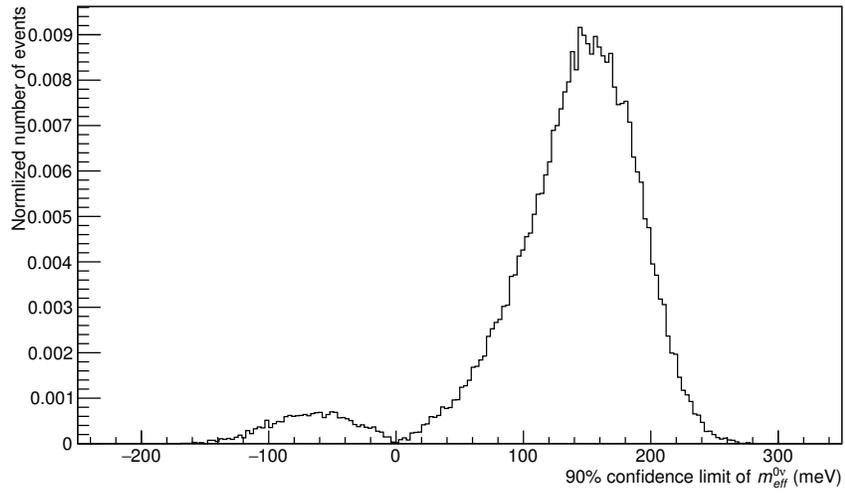
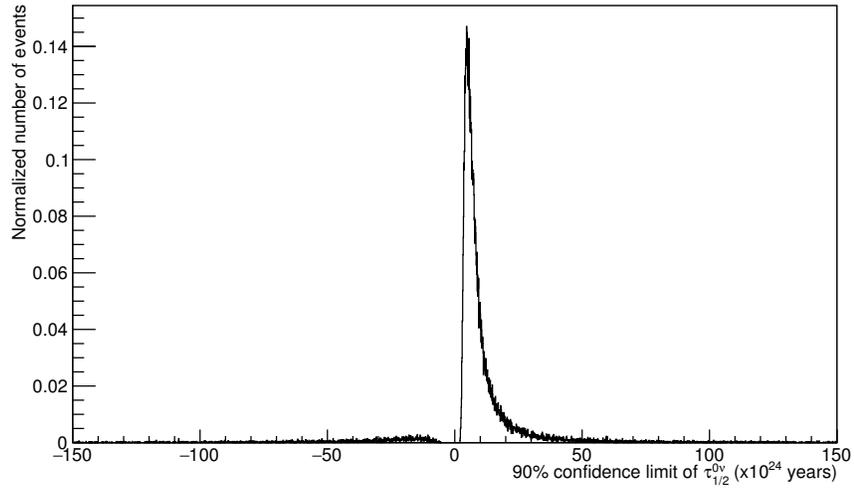


Figure 3.9: Taking the 90% confidence limits from figure 3.7 and converting them to 90% confidence limits on $\tau_{1/2}^{0\nu}$ (top) and $m_{eff}^{0\nu}$ (bottom) using equations 3.4 and 3.5, respectively.

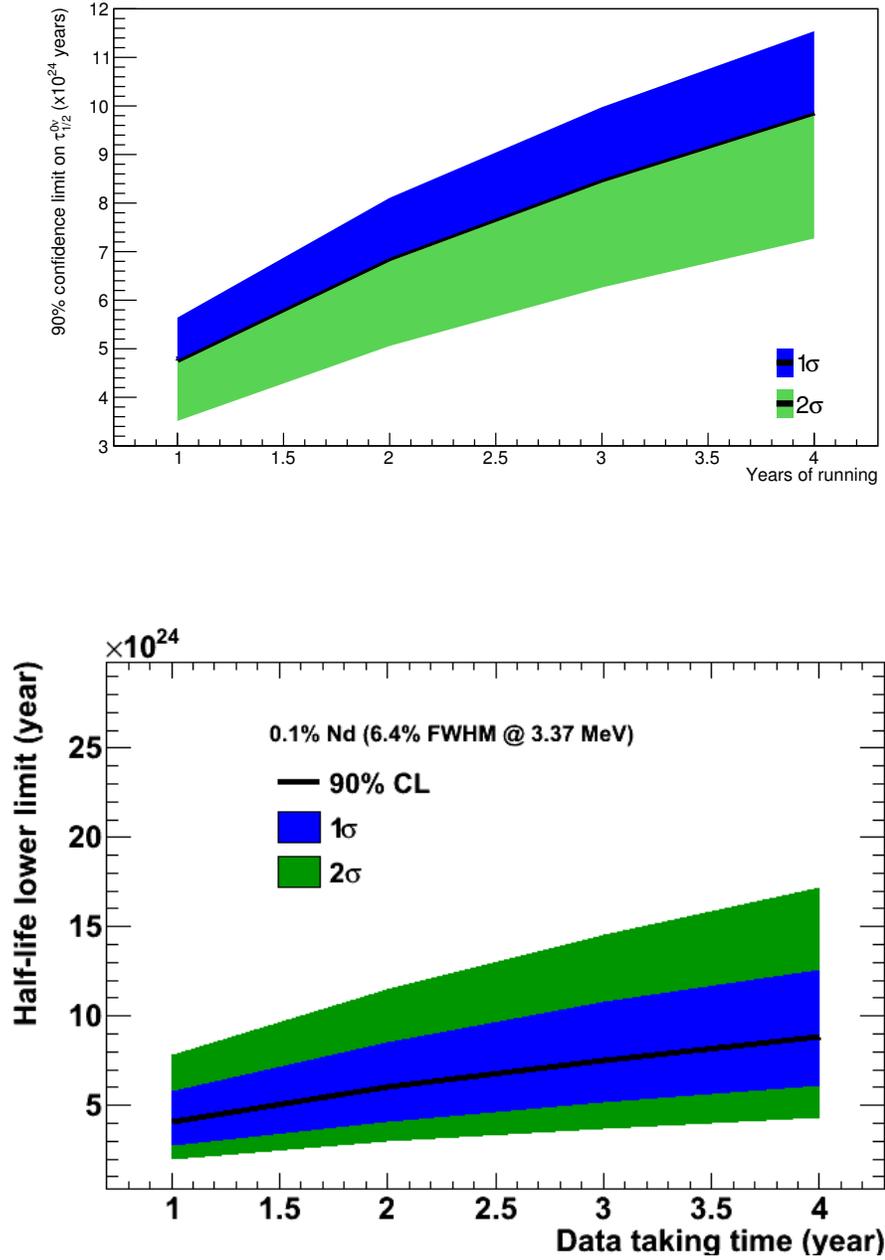


Figure 3.10: The 90% confidence limit on $\tau_{1/2}^{0\nu}$ (denoted by the thick black line) as a function of the number of years of data collection is shown along with the statistical 1σ and 2σ bands on the limit (top). The $2\nu\beta\beta$ group made the official SNO+ plot (bottom) using a likelihood ratio test based on the CL_S method [83].

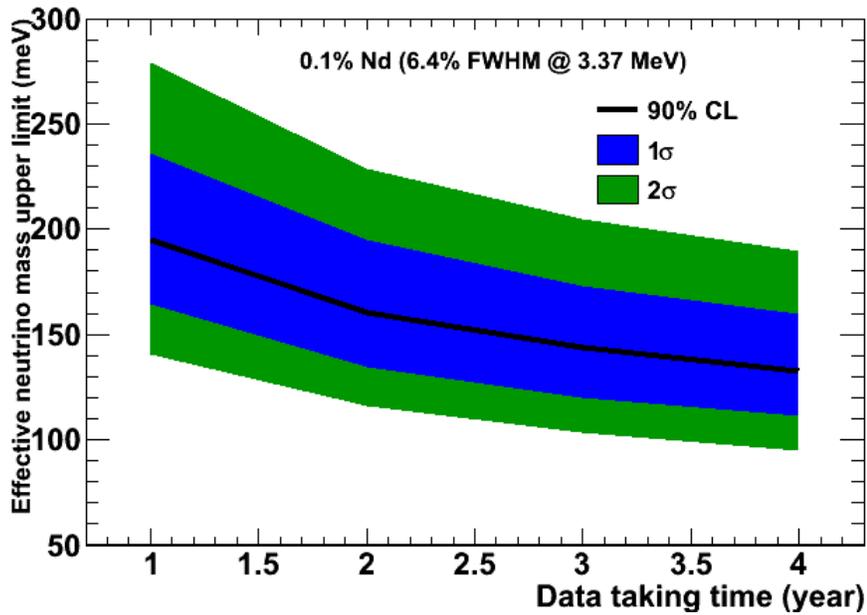
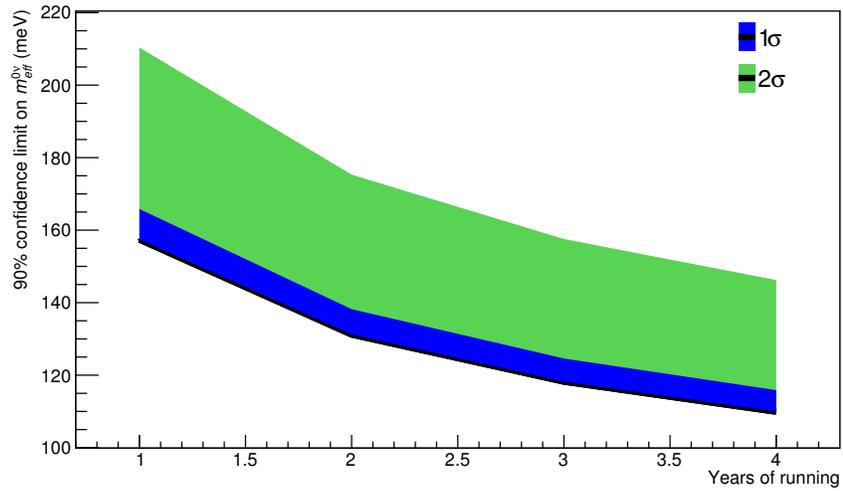


Figure 3.11: The 90% confidence limit on $m_{\beta\beta}$ (denoted by the thick black line) is shown along with the statistical 1σ and 2σ bands on the limit (top). The $2\nu\beta\beta$ group made the official SNO+ plot (bottom) using a likelihood ratio test based on the CL_s method [83].

3.2 Systematics

With the robust method of section 3.1 for determining the SNO+ sensitivity to observing $0\nu\beta\beta$ in ^{150}Nd , I investigate how the addition of systematic errors affects the limits the experiment is able to set. The systematics I consider relate to the energy resolution and energy scale. If I assume in the PDF that the energy resolution is better (or worse) than the actual resolution the data indicate, then how do the 90% confidence limits on $\tau_{1/2}^{0\nu}$ and $m_{\beta\beta}$ change to reflect this incompatibility? Also, if in the PDF I apply an incorrect energy scale that is too low or too high, then what is the effect on the extracted limits?

To study the effect of assuming an energy resolution that is different than the achieved resolution, I fit a series of PDFs that have resolutions varied in steps of 50 Nhit/MeV between 250 Nhit/MeV and 600 Nhit/MeV to a fake data set, which I create in the same way I described in section 3.1 and which has a nominal resolution of 400 Nhit/MeV. I fit each PDF to the same fake data set, and repeat this for 50000 distinct data sets. Figures 3.12 and 3.13 show the resulting 90% confidence limits on $\tau_{1/2}^{0\nu}$ and $m_{\beta\beta}$, respectively. I compare each limit to the limit at the nominal resolution of 400 Nhit/MeV, as both the PDF and the fake data have the same resolution here.

Systematically increasing the assumed energy resolution relative to the actual energy resolution artificially lowers (worsens) the half-life lower limit, whereas reducing the assumed energy resolution relative to the actual energy resolution has the opposite effect. Sharpening the energy resolution in the PDF leads to a relative drop in the number of background events in the region about the expected $0\nu\beta\beta$ peak (the region of interest (ROI)), so the fit compensates

by finding more signal events in that area, thus lowering the half-life lower limit. Widening the resolution relative to the data causes the fit to find fewer signal events in the ROI. The bottom plot in figure 3.12 shows the percent difference in the lower half-life limits per year of data taking compared to the standard limits. This plot indicates that, to have an error no higher than 5% on the lower half-life limit after five years of data taking, the assumed energy resolution must be within approximately +10% and -5% of the actual energy resolution. Figure 3.13 shows similar trends for the upper limit on the effective neutrino mass, where, to have an error no higher than 5% on the limit also requires the assumed energy resolution be within +10% and -5%.

I study the effect of systematically shifting the assumed energy scale in a similar way. Here, I shift the PDF relative to the data either higher or lower in energy by an integer number of bins, where each bin has a width of 50 keV. The PDFs I use have a systematic shift between -100 keV and 100 keV (two bins in either direction⁸). Again, I fit each PDF to the same unaltered fake data set, and repeat the series of fits for 50000 distinct data sets. The limit with a 0 keV shift becomes the standard to which I compare the other extracted limits. I hold the energy resolution of each PDF and the fake data at 400 Nhit/MeV.

From the bottom plots in figures 3.14 and 3.15, it is evident that shifts in the energy scale have a much more dramatic effect on the extracted half-life and effective mass limits than errors in the energy resolution. Even a small shift in the energy scale of 1.5% (50 keV) may lead to an error in the extracted limits of up to approximately 50% after five years of data collection. Also, the shape of these plots are more complicated than the equivalent plots in figures

⁸The fit quickly fails if I shift the PDF beyond this.

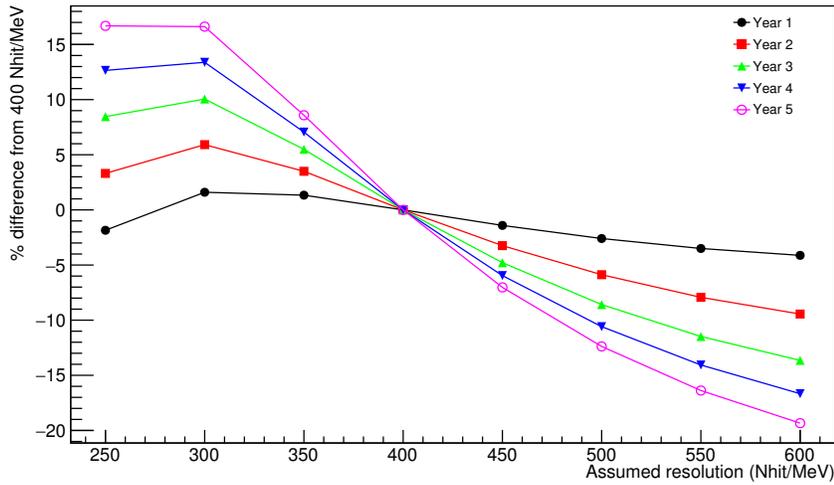
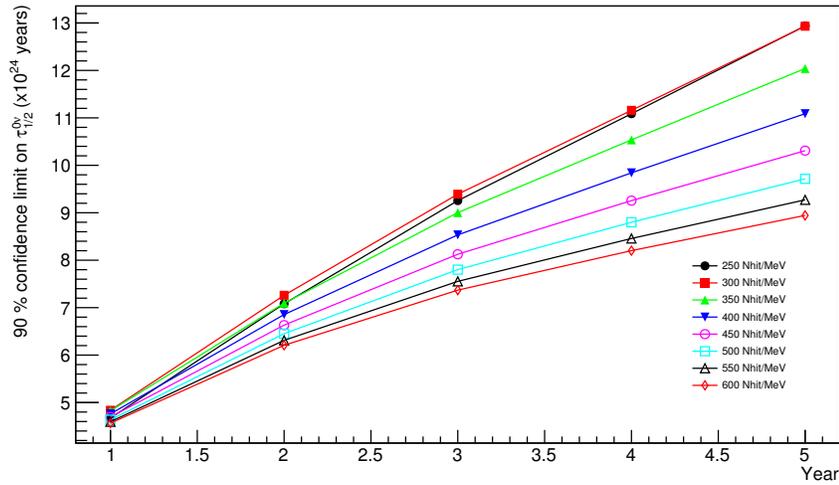


Figure 3.12: The lower limit on $\tau_{1/2}^{0\nu}$ (upper plot) becomes artificially worse (better) as the assumed energy resolution relative to a true energy resolution of 400 Nhit/MeV is increased (lowered). Calculating the percent difference (lower plot) in the lower limit to that where the assumed energy resolution matches the actual energy resolution (400 Nhit/MeV) shows that, after 5 years of data collection, the lower limit can shift by as much as approximately 10% with a 10% error in the assumed resolution. Lines are drawn to guide the eye.

3.12 and 3.13. Systematically shifting the energy scale upward causes many more background events to fall into the ROI, particularly as a result of the very sharp $2\nu\beta\beta$ background, thus drastically degrading the limits on the half-

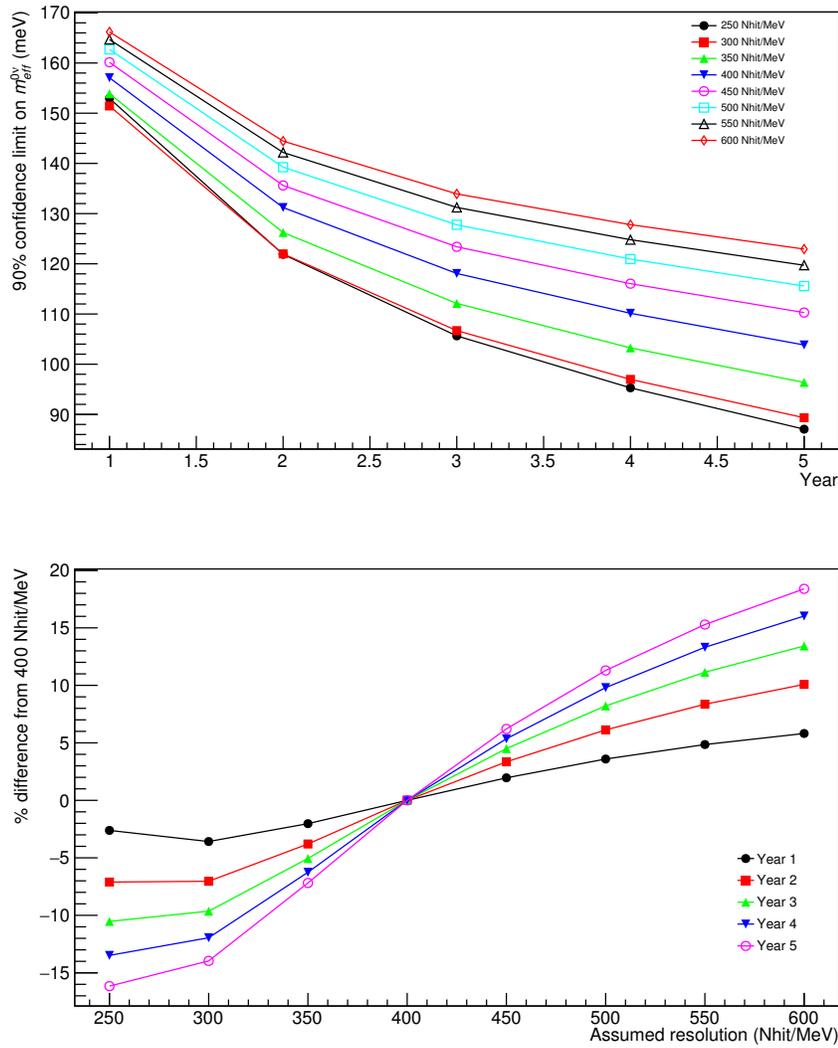


Figure 3.13: Similar to figure 3.12, the upper limit on $m_{\beta\beta}$ (upper plot) becomes artificially worse (better) as the assumed energy resolution relative to a true energy resolution of 400 Nhit/MeV is increased (lowered). The percent error in the upper limit (lower plot) relative to the limit where the assumed energy resolution matches the actual energy resolution again shows that, after 5 years of data collection, the upper limit can shift by as much as approximately 10% with a 10% error in the assumed resolution. Lines are drawn to guide the eye.

life and effective neutrino mass. Shifts of the energy scale downward have the opposite effect for small shifts, causing small artificial improvements in the

limits. However, larger downward shifts of the energy scale also cause major reductions in the set limits, as the fitter attempts to compensate for the large drop in the number of $2\nu\beta\beta$ events in the ROI. Thus, accurately calibrating the energy scale is crucial to extract the 90% confidence limits on both $\tau_{1/2}^{0\nu}$ and $m_{\beta\beta}$.

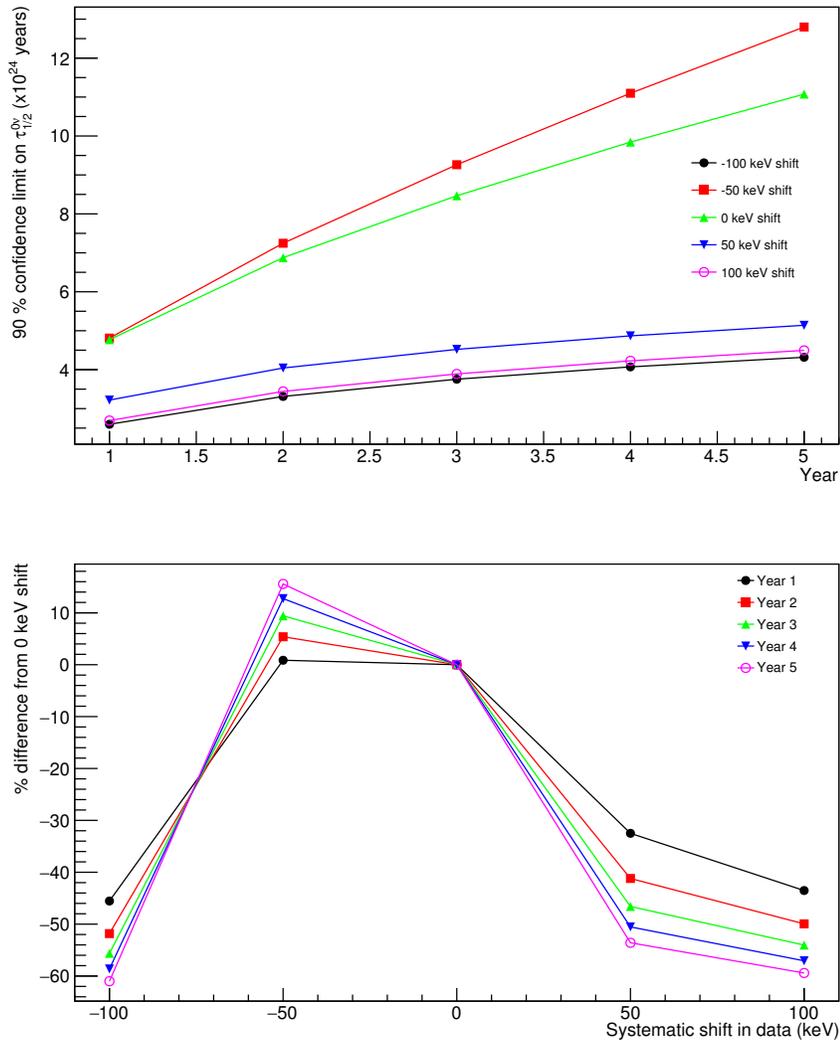


Figure 3.14: Small shifts in the assumed energy scale cause large differences in the extracted lower limit on $\tau_{1/2}^{0\nu}$ (upper plot). The error in the limit after five years of data taking may be as large as 60% for a 3% shift in the energy scale (lower plot). Lines are drawn to guide the eye.

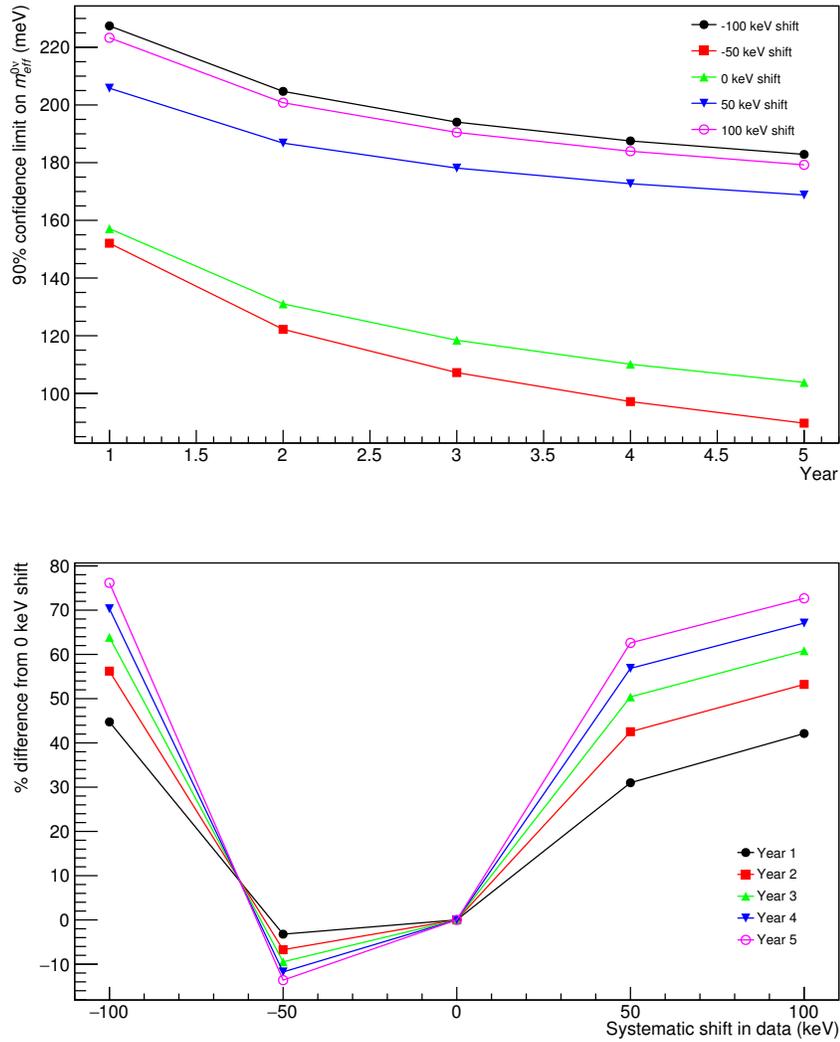


Figure 3.15: Like in figure 3.14, small shifts in the assumed energy scale lead to large differences in the extracted upper limit on $m_{\beta\beta}$ (upper plot), with the error growing as high as approximately 70% for a 3% shift in the energy scale after five years of data collection (lower plot). Lines are drawn to guide the eye.

3.2.1 Tail Study

The work in section 3.1 was driven by an interest to understand the effect a systematic high-energy tail on the SNO+ energy resolution function would

have on the experiment's sensitivity to observing a $0\nu\beta\beta$ signal. We typically assume the energy resolution, given by equation 3.1, is Gaussian with a width that depends explicitly on energy (that deposited in the liquid scintillator), or, more accurately, on the number of PMTs that observe a photon produced in any interaction (which we denote N_{hit}). If some physical process exists that causes a systematic upward fluctuation in the number of photons produced in either all or some subset of interactions, then we might expect this to add a high-energy tail onto the assumed Gaussian resolution⁹.

A suite of radioisotope calibration sources will enable the measurement of the SNO+ energy resolution (see section 2.1.2). One of these sources, which forms the subject of chapter 4, contains ^{60}Co . In brief, ^{60}Co beta decays, almost always emitting two gamma rays with a sum energy of 2.507 MeV in the process. The electron released in the decay has an endpoint energy of 314 keV, and is usually fully contained within the source. Nevertheless, an electron may sometimes escape the source (with only some fraction of its initial energy remaining at escape) and add its energy to the calibration line at 2.507 MeV. This mimics a high-energy tail on the energy resolution and, unless a high-energy tail resulting from physics in the scintillator is larger than this contribution from the ^{60}Co calibration source, the uncertainty on measuring the tail will be limited by our understanding of the ^{60}Co source behaviour.

To investigate the effect of a high-energy tail, I assume the calibration data from the ^{60}Co source contains a tail similar to that in figure 3.16, which is the energy distribution of escaped electrons (see section 4.2.3 for details). I fit this

⁹The number of photons produced in any interaction is a function of the energy deposited in the liquid scintillator, but effects like quenching make this function non-linear. Also, the process of creating photons for very low energy depositions is typically considered a Poisson process, but may, in fact, have larger tails than a Poisson process accounts for. Sections 2.1.2 and 5.1 offer more details.

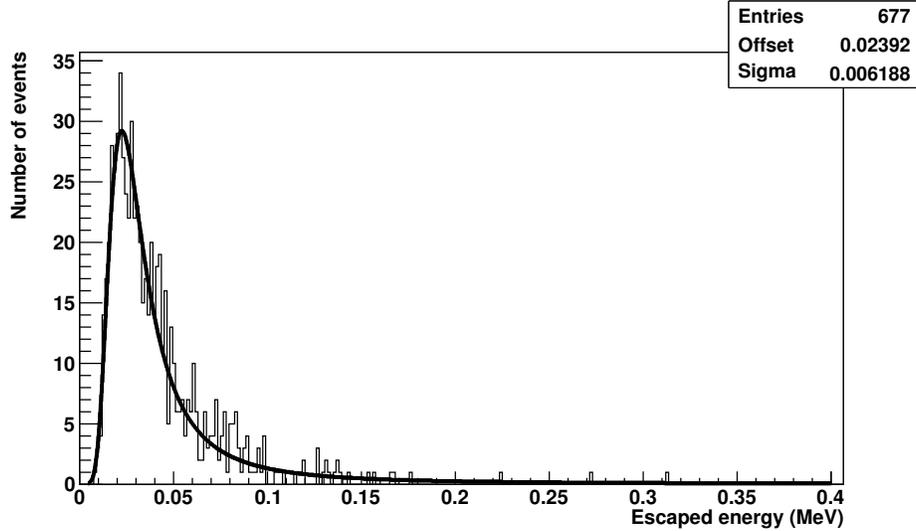


Figure 3.16: The energy of electrons that escape the ^{60}Co calibration source is fit with a Landau distribution (thick line). The resulting Landau distribution describes a high-energy tail on the SNO+ energy resolution. Here, 677 electrons out of 10^6 simulated escaped the source and deposited energy in the liquid scintillator.

to a Landau distribution (which describes energy loss processes), and scale its contribution to the energy resolution as a percentage of events that have a tail contribution. Then, using PDFs that contain this high-energy tail in the resolution to fit data that do not exhibit those tails (or that have a Gaussian energy resolution), I can investigate the effect of an unknown tail on the $0\nu\beta\beta$ sensitivity¹⁰.

To include this high-energy tail in the PDFs, I numerically convolve each PDF with a binned version of the Landau distribution, thus systematically upward-shifting the energies in the PDF. I control the contribution of this tail to the PDF by scaling the percentage of events that have a tail contribution.

Figure 3.17 shows the effect a high-energy tail has on the $2\nu\beta\beta+2\nu\beta\beta^*$ PDF

¹⁰This is really the inverse of the question of an unknown high-energy tail resulting from the scintillation process, but, nonetheless, it gives us some understanding of how large a tail may be before it significantly affects the experiment's sensitivity.

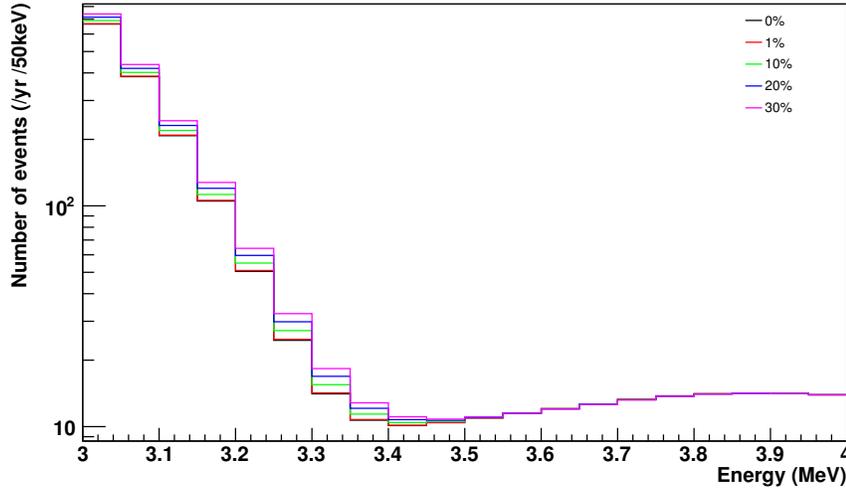


Figure 3.17: Convoluting the $2\nu\beta\beta+2\nu\beta\beta^*$ PDF with the Landau distribution systematically shifts the energy distribution to higher energies. The result is an increase in the number of events in the ROI relative to the number if no tail is present, which mimics a $0\nu\beta\beta$ signal.

near its endpoint.

I fit the PDFs containing the high-energy tail contribution, which I scale between 0% and 30%, to fake data sets generated using the tail-less energy resolution in equation 3.1. I repeat this for 50000 distinct fake data sets, which results in the 90% confidence limit distributions on the number of observed $0\nu\beta\beta$ events shown in figure 3.18. This figure suggests that adding a high-energy tail onto the energy resolution when data does not exhibit such behaviour seems to falsely improve (lower) the 90% confidence limit of the measured number of $0\nu\beta\beta$ events¹¹. This, in turn, artificially improves the corresponding limits on both $\tau_{1/2}^{0\nu}$ and $m_{\beta\beta}$, as figures 3.19 and 3.20 show. After three years of data collection, the error in the extracted 90% confidence

¹¹Based on these results, I expect the opposite trend in the extracted confidence limits for the inverse problem of an unknown tail in the data and a Gaussian energy resolution in the PDFs. That is, the limits would likely be artificially worsened as the fit attempts to compensate for data that extend to higher energies than the PDFs consider.

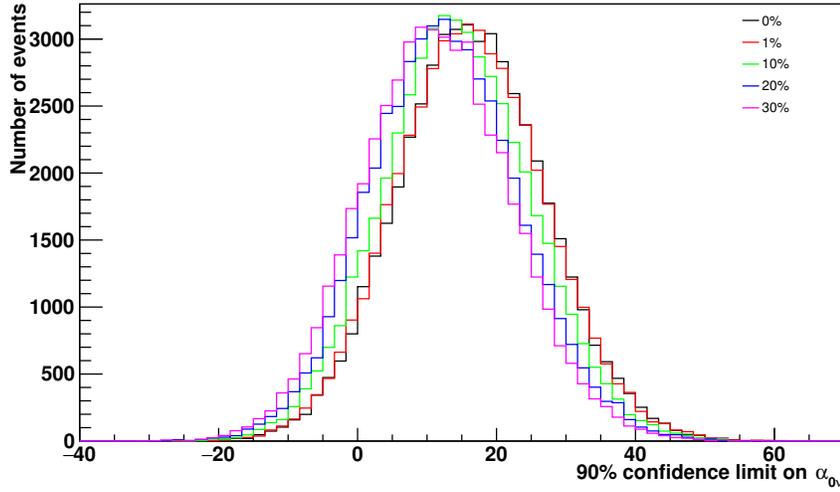


Figure 3.18: The 90% confidence limits of the number of observed $0\nu\beta\beta$ events after one year of data collection show how larger tail components cause more significant reductions in the 90% confidence limit.

limit of $\tau_{1/2}^{0\nu}$ may be as high as 6% if 5% of events contain a contribution from a high-energy tail. Similarly, the 90% confidence limit of $m_{\beta\beta}$ could suffer up to 5% systematic error resulting from a 5% high-energy tail contribution.

It is critical, then, to understand any tails on the energy resolution to prevent overstating our ability to observe a $0\nu\beta\beta$ signal. This is especially true in the event of a real $0\nu\beta\beta$ observation by SNO+. Thus, SNO+ will require a well-understood set of calibration sources in order to measure the Gaussian energy resolution and any deviance from it (and to peg down the energy scale). Chapter 4 discusses this idea further. Also, both through its calibration program and through *ex situ* laboratory measurements, SNO+ must be able to understand to high accuracy the response of its liquid scintillator to various interaction types. One aspect of this—quenching—is the subject of chapter 5.

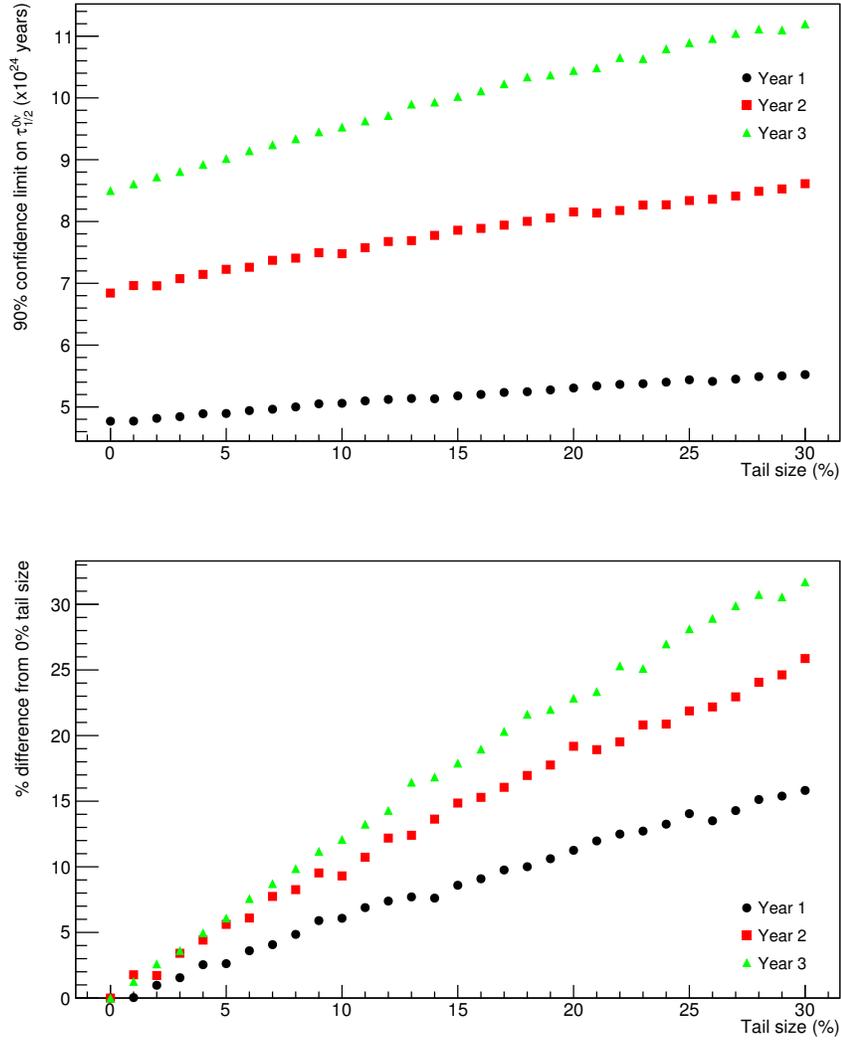


Figure 3.19: Adding a high-energy tail to the PDF when there is no tail in the data causes better extracted lower limits on $\tau_{1/2}^{0\nu}$ (upper plot). The error in the limit after three years of data taking may be as large as 6% for a 5% tail contribution (lower plot).

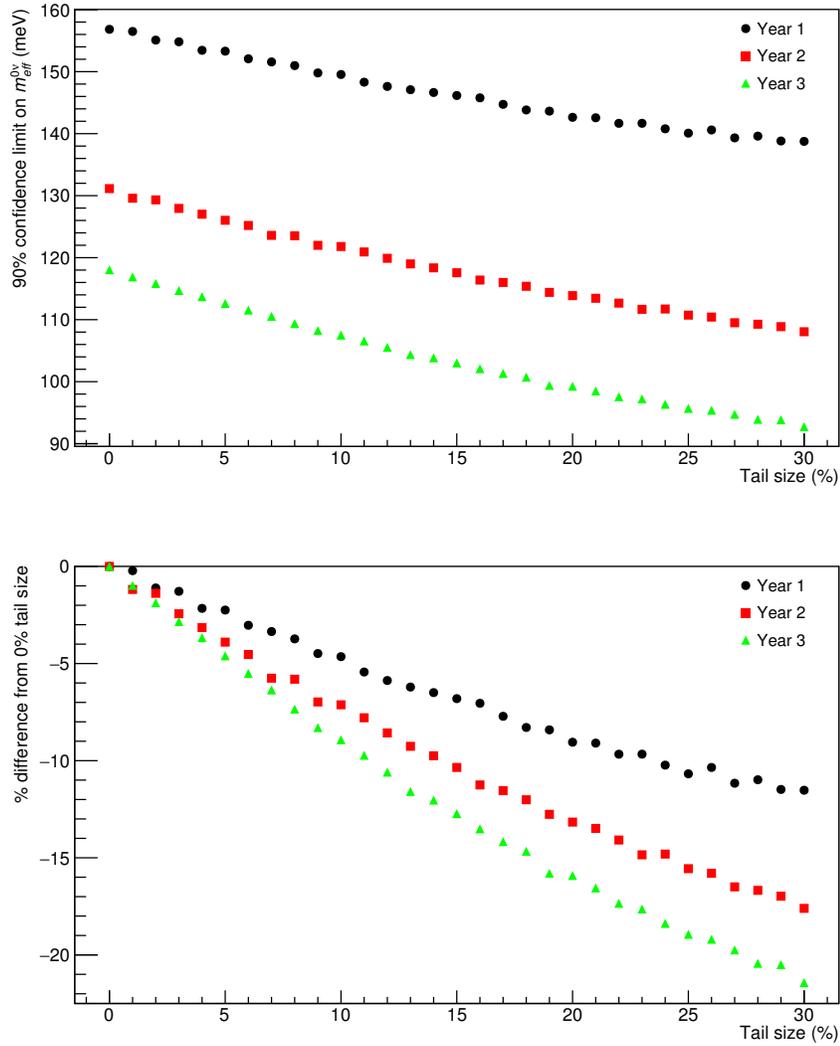


Figure 3.20: Like in figure 3.14, small shifts in the assumed energy scale lead to large differences in the extracted upper limit on $m_{\beta\beta}$ (upper plot), with the error growing as high as approximately 5% for a 5% tail contribution after three years of data collection (lower plot).

Chapter 4

The ^{60}Co Calibration Source

It is not by reading that science is sustained. It is by contact with the object—by systematic observation and experiment.

- Charles Singer, *A Short History of Scientific Ideas to 1900*

Because SNO+ will focus on investigating the $0\nu\beta\beta$ of ^{130}Te , the response of the detector near the decay endpoint of 2.53 MeV must be well understood. Of particular importance are understanding the detector energy scale and resolution, with emphasis on any non-Gaussian shape of the resolution function. As chapter 3 concludes, just a 1% contribution from a high-energy tail in the resolution function could lead to a 5% systematic error on the number of observed $0\nu\beta\beta$ events. A robust calibration program, outlined in section 2.1.2, is envisioned for SNO+, aiming to deliver a complete understanding of the LS energy response (and other optical properties) in each experimental phase.

Radioactive sources that we may consider using to calibrate the SNO+ detector should satisfy one or more of the criteria: the decay leads to the emission of gamma-rays with energies or sum energies near the ROI¹ for a

¹Electrons are the ideal particle for calibration, as that is what SNO+ will be observing,

detailed understanding of the detector response there; no more than a few gamma-rays are emitted, providing an uncomplicated signal; gamma-rays are emitted with a high intensity greater than 85% for a high-efficiency calibration source; and, the decaying isotope is long-lived with a half-life comparable to the experiment's duration. A non-requisite, though advantageous, trait of any calibration source would be that it supplies a secondary radiation to enable tagging of calibration events in conjunction with triggering by the SNO+ detector. This would provide a set of background-free calibration data to assist with studying any non-Gaussian tails on the detector resolution function.

Of the number of well-known and oft-used (and readily available) radioactive sources, ^{60}Co satisfies all of the above criteria and is ideally suited for studying the detector response near the ^{130}Te endpoint. ^{60}Co has a relatively simple decay scheme, leading 99.88% of the time to the emission of an electron with endpoint energy 0.318 MeV and two subsequent gamma-rays of 1.17 MeV and 1.33 MeV. SNO+ will observe the sum energy of the two gamma-rays at 2.51 MeV, giving a calibration point near the endpoint energy of the ^{130}Te decay, while the electron supplies an accompanying tag of any calibration event. In addition, ^{60}Co has a half-life of 5.27 y, so will remain sufficiently active over the course of the experiment.

I have spent a significant amount of time designing and constructing a ^{60}Co calibration source for the SNO+ experiment, specifically for use during the LS and TeLS experimental phases (see section 2.3). The SNO+ calibration plan [92] lists the ^{60}Co calibration source as a LS phase *need*, where the source

but designing an electron calibration source is challenging because of the need to encapsulate the radioactivity. Any encapsulation will have a strong effect on the electron spectrum, and very thin-walled encapsulations typically do not have the structural integrity to hold up to the stringent environmental conditions of SNO+.

plays a significant role in checking, among other things, the SNO+ energy scale and resolution, both globally and as a function of position. I discuss the construction and testing of the source in section 4.1 (detailed design and testing procedures are outlined in appendix B). I also created a RAT simulation of the calibration source, including a validation of the ^{60}Co nuclear decay, which section 4.2 describes in detail.

The SNO+ experiment schedule has the LS phase beginning in summer 2017. As a result, the ^{60}Co calibration source will not be used prior to the completion of this thesis. I have constructed a prototype of the calibration source, but it may happen that small changes to its design occur between now and when the source is required to be ready for deployment. In particular, work is now ongoing to investigate changing the radioisotope from ^{60}Co to ^{46}Sc . While ^{60}Co is a good choice because of its ability to probe the detector response at an energy near the ^{130}Te endpoint, that property also makes it a high-risk choice, as the threat to the experiment of *any* source of ^{60}Co contamination is significant². Any decision to change isotopes will be made neither lightly, nor in the immediate future.

Before it can be shipped to SNOLAB, the calibration source must undergo a rigorous approval process facilitated by both the SNO+ and SNOLAB source committees. This will not occur before the completion of this thesis; however, this chapter (and its appendices) documents most of the information required for that process. Much of the work for this calibration source is done, but it will not be me who sees the source used to calibrate the SNO+ detector.

²We chose to develop a ^{60}Co calibration source when SNO+ was investigating the $0\nu\beta\beta$ of ^{150}Nd , with an endpoint energy of 3.37 MeV, significantly away from the sum energy of the ^{60}Co gamma-rays. The ^{46}Sc decay mimics ^{60}Co closely, with a 0.357 MeV endpoint electron and two gamma-rays with a sum energy 2.01 MeV, which sits below the ^{130}Te endpoint and outside of the SNO+ ROI.

4.1 Source design, construction and testing

The ^{60}Co calibration source is a tagged source with an activity³ of approximately 200 Bq. In brief, the ^{60}Co is contained within a small disk of plastic scintillator that is optically coupled to a fast photomultiplier tube (PMT). This optical apparatus is contained in an internal metal and external plastic housing that encapsulate the radioactivity, preventing it from escaping the source and contaminating the SNO+ detector or experiment area. The electron emitted in the ^{60}Co decay creates scintillation light in the plastic scintillator that the source PMT observes and the SNO+ data acquisition records. The two gamma-rays escape the calibration source largely unimpeded, causing the detector to trigger on the light created by the gamma-rays interacting in the LS. Offline analysis of calibration data will allow SNO+ events to be temporally correlated with events in the calibration source, creating a background-free set of calibration data.

The Department of Physics at the University of Alberta provided an ideal environment in which to construct this calibration source. A new laboratory designed with the intent of working with radioisotopes was available for completing all open-source work. This was the first radioactive work to be completed in that laboratory, so there was no concern of contamination by other sources of (non-naturally occurring) radioactivity. I monitored the work environment for contamination using a Geiger counter on a daily basis. I also performed wipe tests of the laboratory once all work was completed. No laboratory surface exhibited measurable traces of contamination at any time.

The sections below describe in muted detail the design and construction

³This is high enough to provide a sufficient number of calibration events in a reasonable amount of time, but not so high that it challenges the upgraded SNO+ electronics and DAQ.

of the ^{60}Co calibration source, including contamination testing methods and results, and the measurement of the absolute activity of the calibration source. I created detailed procedures that I followed for each step of the process. These procedures are contained in appendix B. An assistant was available, when necessary, to help me complete all laboratory work where there was a risk of spreading contamination. The assistant never came into direct contact with the radioisotope.

4.1.1 ^{60}Co standard

Eckert & Ziegler Isotope Products, Inc. supplied the ^{60}Co liquid standard. The standard contained 37 kBq of $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ in 5 mL of 0.1 M HCl when received in July 2012. I wipe tested the packaging in which the source arrived prior to opening it and moving the source into storage. There was no evidence that the source packaging was contaminated.

According to the manufacturer, the liquid standard contained 10 μg of Co per mL, corresponding to 40.4 μg of $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ per mL. Before doing any work with the liquid $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ standard, I created a non-radioactive $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ standard to test all procedures and to help inform those procedures going forward. I made test samples containing 1 \times or 100 \times the concentration of $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ compared to the liquid standard by dissolving solid $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ into 0.1 M HCl. The high concentration sample provided an easier way to see the $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ as I was working with it.

4.1.2 Plastic scintillator

The $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ is heat-sealed⁴ between two small disks of polyvinyltoluene-based plastic scintillator BC-408 [94], which has a number of properties that make it well-suited for this project. BC-408 is a fast scintillator (2.5 ns pulse width) with a relatively high light output (64% that of anthracene) in the blue/violet part of the spectrum (maximal at 425 nm), allowing for accurate timing of its light emission by a fast PMT that is wavelength-matched to the scintillator's spectral output. The scintillator softens at 70 °C, well above the operating temperature of SNO+, and is relatively inexpensive⁵.

In total, I produced 19 of these so-called ‘button sources’, 16 of which range in activity from approximately 40 Bq to 750 Bq (see table 4.1), the other three being control samples with intended activities of 0 Bq. I constructed two different groups of buttons, which I label 1–5 and 1A–4C, C1, C2, respectively. Construction of the first set occurred in April 2013, with the second set following in November 2013, both using the same ^{60}Co liquid standard (see section 4.1.1).

In the sections below, I describe all of the work that went in to producing these buttons, including the button construction and contamination testing methods. Appendix B contains detailed work procedures pertaining to the button source construction.

⁴Heat-sealing radioactivity between layers of plastic scintillator was used in [93] to build an on-board calibration system for a satellite-based detector. That paper does not elucidate the details of actually doing the heat-sealing, but this is where I gleaned the idea for this technique.

⁵The particular BC-408 sheet I used is a remnant from the construction of the ALTA experiment [95], so was, in fact, free.

Machining and polishing

Starting with a 10 mm thick sheet of BC-408, I rough-cut a 10 mm face length square prism then turned this prism on a lathe to produce an 8 mm diameter, ~10 cm long cylinder. I polished this cylinder (see below) before parting off 2 mm thick disks of the plastic. I refer to these disks as buttons. During the machining process, I used reverse osmosis (RO) water as coolant and sharp carbide tools [96]. These help create a smoother machine finish and prevent overheating of the plastic, which could otherwise lead to crazing (the production of micro-cracks on the plastic surface) [97]. Figure 4.1 shows samples of the rough cut prisms and the machined and polished cylinders.

Each button requires an optically clear surface finish, the reason for which is twofold. First, a smooth surface increases the proportion of photons that will totally internally reflect, which in turn increases the number of photons striking the PMT face. Second, early tests of depositing liquid $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ onto the plastic (see below) indicated that a smoother surface caused the $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ residue to concentrate nearer the centre of the button⁶. Both are essential for an efficient calibration device, making a smoother button face desirable. To arrive at a surface with these properties, I hand-polished each button following a multi-step polishing process.

The polishing process consists of scouring the plastic surface with varying grit sanding paper, polishing paper and polishing paste, with each pass using a decreasing grit size [98]. Sanding papers with grit sizes 400, 600 and 1500, and polishing papers with grit spacings of 5 μm and 3 μm , used in that order,

⁶Plastic surfaces with significant post-machining scratches caused the $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ salt to become trapped in the scratch marks, so the residue would remain spread out over the surface.



Figure 4.1: Rough-cutting the plastic scintillator resulted in significant marking of the cut sides (top). After turning on the lathe, the scintillator cylinders had a somewhat rough surface finish (cylinders 6 and 3), which the polishing process made clear (cylinders 7 and 8) (bottom).

produce a near optical surface finish⁷. I wetted the papers with RO water prior to polishing, again to prevent overheating of the plastic surface. To polish the button faces, I placed each paper flat on a table, then moved the button face across the paper in a circular motion using a single finger. On each face, 15 counter-clockwise rotations followed 15 clockwise rotations, which I repeated per each paper until no large scratches in the surface were visible. I polished the cylinders prior to slicing them into the buttons to allow for a polished button edge (without the need to polish that small edge of the button afterward). I always polished the cylinders by scouring the surface along the cylinder length (never circumferentially) until I obtained a sufficiently clear optical finish.

Clear plastic polishing and cleaning pastes⁸ succeeded the paper polishing. I wetted a small area of a microfibre cloth with one of the pastes, then rubbed the button with the cloth between my thumb and index finger, including both faces and the button's side, until I obtained the desired optical surface finish (modified from the paper polishing process for ease of execution). Finally, I cleaned the plastic surface using a microfibre cloth wetted with RO water, which removed any remaining plastic cleaner.

Because I hand-polished each piece of plastic scintillator, each final button had slightly different diameters and thicknesses. I matched pairs of buttons that were closest in radii and whose thicknesses summed closest to 4 mm to use for the remainder of the button source construction process. Table 4.1 lists the dimensions of each button.

⁷The 400 and 1500 grit paper is from Carborundum Abrasives, the 600 grit paper is 3MTM WetordryTM. The 3 μm polishing paper is aluminum oxide LFG3P and the 5 μm polishing paper is silicon carbide LFG5P, both from Thor Labs.

⁸Both products are Meguiar's Mirror Glaze Professional, number 10 Clear Plastic Polish and number 17 Clear Plastic Cleaner.

Button Number	Thickness (± 0.01 mm)		Diameter (± 0.01 mm)		Drops	Expected Activity (Bq)	Measured Activity (Bq)		
	Top	Bottom	Top	Bottom					
1	2.00	1.96	7.88	7.88	1	34	42.3	\pm	0.5
2	1.97	1.91	7.90	7.97	1.5	50	57.7	\pm	0.6
3	2.05	1.86	7.91	7.99	2	67	79.8	\pm	0.9
4	1.85	2.04	7.94	8.01	2.5	84	117.8	\pm	1.2
5	—	—	—	—	0	0	$< 2.48 \times 10^{-3}$		
1A	1.99	1.98	7.93	8.02	10	300	326.5	\pm	5.5
1B	1.88	1.84	7.87	8.01	11	330	344.3	\pm	5.8
1C	1.85	1.90	7.86	8.03	12	360	382.3	\pm	6.4
2A	2.01	1.93	7.88	8.02	13	390	423.1	\pm	7.0
2B	1.86	1.91	7.93	7.96	15	450	478.7	\pm	8.0
2C	1.94	1.82	7.98	8.01	16	480	510.9	\pm	8.5
3A	1.90	1.82	7.92	8.01	18	540	553.6	\pm	9.2
3B	1.85	1.79	7.91	8.01	20	600	627.2	\pm	10.5
3C	2.00	1.89	7.86	8.07	22	660	693.9	\pm	11.5
4A	1.90	1.86	7.96	8.01	7	210	229.5	\pm	3.9
4B	1.91	1.80	7.95	8.02	9	270	293.8	\pm	4.9
4C	1.88	1.99	7.95	8.01	24	720	769.6	\pm	12.8
C1	1.55	1.51	8.01	8.03	0	0	$< 5.29 \times 10^{-3}$		
C2	1.66	1.62	8.04	8.00	0	0	$< 3.40 \times 10^{-3}$		

Table 4.1: This table lists the dimensions, and expected and measured activities of each button source. Top and bottom refer to the buttons I paired prior to bonding. All radioactive depositions were on the bottom buttons. Buttons 5, C1 and C2 were controls receiving no $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$. I measured the absolute activity of each button after performing the initial dry wipe test. The activity is the average of the activity measured using the 1172 keV and 1333 keV gamma-ray peaks for buttons 1–4 and 1A–4C. The activities of buttons 1–4 are the weighted average over three separate counts. The activity limits for buttons 5, C1 and C2 are the 90% Feldman-Cousins limits on the activity from the 1333 keV gamma ray peak. The limit for button 5 is the combined Feldman-Cousins limit from three separate counts. The errors in the activities incorporate a 0.5% systematic error (see below). I measured buttons 1–5 on May 9, 2013, and 1A–4C, C1, C2 on December 9, 2013.

Liquid source deposition

To control the amount of ^{60}Co each button source contains, I deposited the liquid $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ on the bottom buttons using a pre-calibrated $5 \mu\text{L}$ micropipette consisting of a thin glass tube and steel plunger. The activity of the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ liquid standard was approximately 30 kBq in October 2013 when most of the deposition took place. As such, each $5 \mu\text{L}$ deposition contained approximately 30 Bq of activity, meaning each button required multiple drops of solution to make up the intended activity. Table 4.1 lists the number of drops I deposited on each button, and the intended activity of each.

I performed the deposition inside a fume hood, which had a constant flow of air running into it from the room. Figure 4.2 shows the fume hood setup. With the fume hood sash left open⁹ approximately 10 cm, the HCl from a single $5 \mu\text{L}$ droplet would evaporate completely within approximately 1 h. To restrict the overall size of the liquid drop on a button, I allowed approximately 10 min to elapse between depositing individual droplets on any button. This was enough time to allow for some evaporation of the HCl, so the liquid drop size never grew to more than approximately 4 mm in diameter¹⁰. Thus, I deemed the 4 mm droplet size to be sufficient to not have any $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ too near the edges of the buttons. After depositing all of the droplets, I left the buttons for a period of 24 h to allow for the HCl to completely evaporate.

Appendix B.2.1 contains the complete procedure I followed for the deposi-

⁹This created a modest flow of air over the button surface. I noted in early tests with large droplets from a pipette that too high an air flow actually smeared the $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ salt residue along the direction of the airflow, rather than allowing it to concentrate at the centre of a button. I did not observe this effect with the 10 cm opening.

¹⁰Early tests showed that the $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ residue would retract toward the centre of the liquid droplet as it evaporated, causing its radius to gradually decrease and restricting the area over which it would spread.



Figure 4.2: I deposited the liquid $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ onto the buttons in a fume hood, set up as shown (left). The liquid collected at the centre of each button (right). The size of the liquid drop depended on the number of droplets required to obtain the necessary activity.

tion process. An assistant was present in the laboratory throughout, ensuring I followed each step of the procedure in turn.

Heat-sealing

Once the HCl completely evaporated from the button surface, a thin layer of $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ residue remained, confined mostly to the central region of the button. To encase the residue inside the plastic scintillator, I heat-sealed (bonded) the corresponding top and bottom button halves together. In creating a bonding procedure (the procedure I eventually adopted is in Appendix B.2.2), I did many tests. Using a heat gun to increase the button temperature caused deformation of the plastic, largely a result of overheating the plastic in a short period of time. A dedicated polystyrene flood lamp-heated annealing oven proved to be too difficult to maintain at a particular temperature. However, using this oven we learned that simply heating two buttons that had been placed one atop the other was insufficient to bond them. We needed to apply some pressure on the buttons to cause them to

bond. Attempting to compress them between glass microscope slides¹¹ was challenging as it was difficult to maintain a uniform pressure across the button face, leading to air gaps in the resulting bonds. We also encountered difficulty in aligning the centres of the two button halves. Attempting to encase them in an o-ring only resulted in deformation of the sides of the buttons in the heating process.

To circumvent these problems, I machined an aluminum bonding jig, as in figure 4.3, whose purpose was to align the button centres and to apply uniform pressure to the button face during the heating. The jig consisted of an aluminum tube with windows cut out of the sides (to see the alignment) that was set on a base with an 8 mm diameter, 1 mm deep inset machined out of the centre. A plug with a similar groove and a mass of 215.0 g slid into the cylinder from the top. I polished both grooves to remove machining scratches and to ensure the buttons would not stick to the aluminum during the bonding process. To prepare the buttons for bonding, I set the button holding the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ residue into the groove in the jig's base, residue-side up. I pressed the top button half into the groove in the plug (fitting snugly inside so as to ensure it would not fall out while assembling the jig), then lowered the plug down the cylinder to rest on top of the base, thus aligning the two buttons.

Next, I moved the jig into an oven, which figure 4.3 also shows. I chose this oven specifically because it had the option of heating under vacuum, which we thought may decrease the chance of air bubbles being trapped in the bond. This never proved to be a useful feature; however, the heating and cooling

¹¹Glass does not stick to the plastic scintillator.



Figure 4.3: I constructed an aluminum bonding jig (left) to hold the scintillator buttons in the oven (right) during the heat-sealing process.

of this oven¹² matched the manufacturer-recommended [99] annealing cycle of the BC-408 plastic scintillator: heat the plastic to 85–90 °C, soak for a period of 15 minutes (for the button dimensions), then allow to cool by 10 °C/h to room temperature. According to the manufacturer, this annealing procedure does not affect the light yield of the scintillator.

To allow sufficient time for the buttons to bond, I modified the annealing cycle by increasing the soak time to 14 h. I monitored the oven temperature using the temperature indicator strip on the aluminum jig, as well as a glass thermometer inside the oven (the temperature indicator always lagged the thermometer owing to the time lag inherent in the heat transfer between the jig, buttons and air in the oven). Figure 4.4 shows the temperature profile of

¹²It was not possible to tightly control the oven temperature and rate of temperature change. Nevertheless, I did find a setting on the oven that enabled me to mimic the annealing process very well. Because I had done all of the testing using this oven, I chose to continue to use it, regardless of the lack of finesse in controlling it.

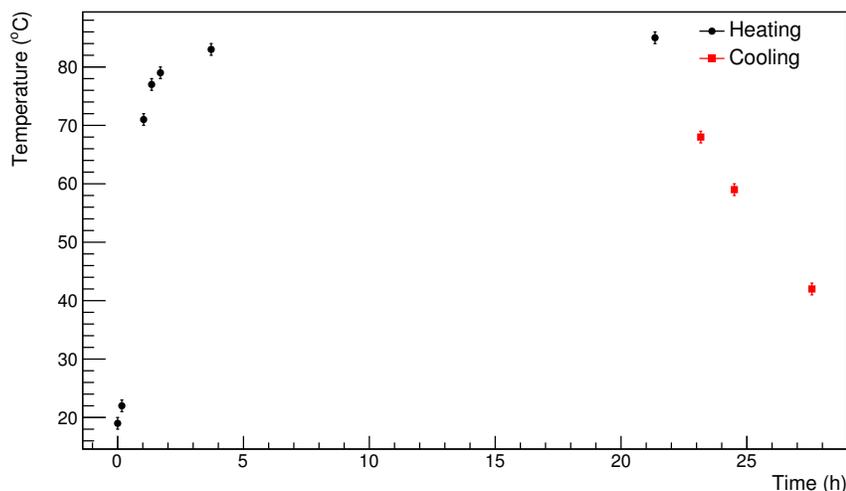


Figure 4.4: The temperature profile of a successful bond includes a fairly quick increase in button temperature, followed by a significant soak period, with a cool down rate near 10 °C/h. During this particular cycle, the oven was under 28 mmHg vacuum. Whether or not the oven was under vacuum did not affect the temperature profile in a significant way.

a bonding cycle. Once the bonding cycle was complete, I removed the jig from the oven to the fume hood, removed the cylinder and plug, then moved the bonded button source into a clean storage container to await further testing.

To test the strength of the bond, I bonded test buttons holding non-radioactive $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$. Afterward, I attempted to pry the buttons apart by hammering a flat screwdriver into the bond. This proved ineffective. I also cooled the plastic down to approximately -10 °C and repeated that test. The bond did not fail, nor did the plastic shatter or crack.

I also examined the bonds using a microscope. To begin, I viewed the deposition of non-radioactive $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ on the plastic surface before and after bonding. These tests used either 1× or 100× the nominal concentration of $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ of the ^{60}Co standard. As figure 4.5 shows, the crystal structure of the $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ is visible in the residue of the 1× solution after the HCl

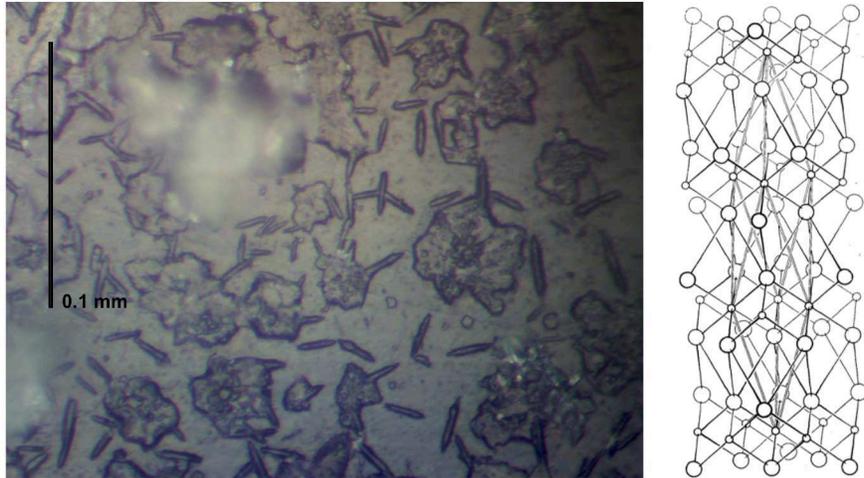


Figure 4.5: The crystal structure of $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ (right) is clearly visible in the residue remaining on the button under $40\times$ magnification (left). In the drawing of the crystal structure, small circles represent Co atoms, large circles represent Cl atoms and the double lines indicate the unit cell [100].

has completely evaporated.

Figure 4.6 shows the same deposition under $4\times$ magnification. It is clear how the residue retained a somewhat circular shape as the droplet evaporated. After bonding, the residue largely retains its shape, though the extent to which it is spread across the button surface is somewhat increased. Prior to bonding, the major axis of the ellipse encompassing the residue was approximately $1150 \mu\text{m}$, whereas this axis increases to nearly $1260 \mu\text{m}$ after bonding. I observed this increase in size in all tests. I also observed how, prior to bonding, the surface structure of the residue spanned several focal planes in the microscope's field of view, which the bonding process seemed to flatten into a single focal plane. Assuming the residue maintains its volume, this flattening explains the increase in radial extent.

Figure 4.7 indicates how a small air gap remains surrounding the residue after bonding. This air gap, using the $100\times$ solution, ranges in size from ap-

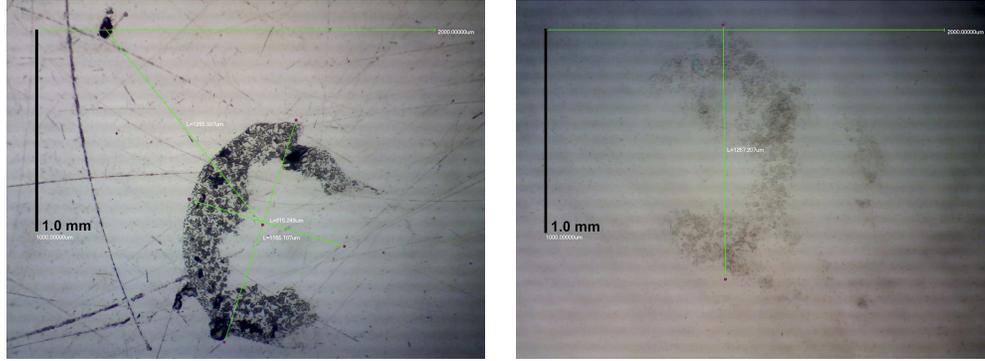


Figure 4.6: The $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ residue prior to bonding (left) under $4\times$ magnification shows how the residue maintains the approximate shape of the droplet after evaporation. After bonding (right), the same residue largely retains its shape, though is slightly flattened. The image on the right is flipped about the vertical axis relative to the image on the left.

proximately $10\text{--}30\ \mu\text{m}$. The existence of this air gap confirms the expectation that the heat-sealing process fuses the two buttons together in a seamless way outside of the area containing the residue.

I inspected the bonds for each button source in table 4.1. The main goal of this inspection was to search for defects in any of the bonds that could have produced a direct line from the residue to the exterior of the button source. Focussing on the bond plane of the button source, I observed what appeared to be bits of dust that had become encased in the scintillator during the bonding. This flotsam likely settled on the bottom button while the HCl was evaporating in the fume hood. Several button sources showed pieces of dust at the edge, where the bond between the two button halves is imperfect¹³. Nevertheless, button sources did not typically show evidence of dust pieces connecting the

¹³The microscope showed that the extreme edges of the button did not bond. This region appeared very dark, implying light from the microscope was not penetrating the entire button source, but being reflected back in the air gap between the two buttons. This did not affect the overall integrity of the bond, however, as the mechanical stress test on the bond, which I described earlier, shows. The poor bond near the edge is likely a result of the polishing process, which tended to round the edges of the buttons, so the two button faces were not touching at the edge during bonding.

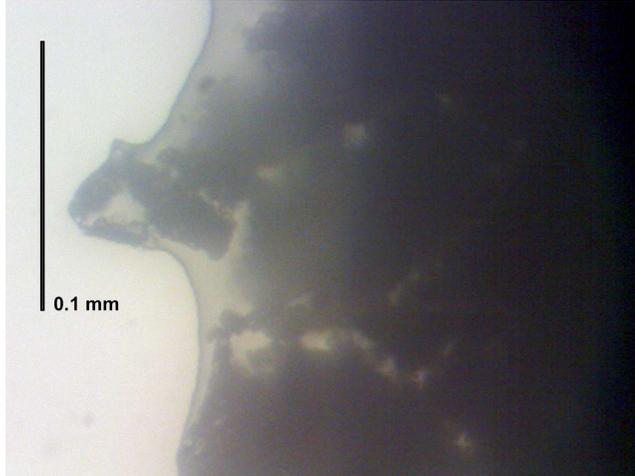


Figure 4.7: An air gap of 10–30 μm surrounds the residue after the bonding process is complete. I observed this in most tests. This image shows the 100 \times solution at 40 \times magnification.

residue to the exterior. Figure 4.8 shows the bond plane for button source 4C, which exhibits the two characteristics I just described. Lastly, looking at the polished button surfaces through the microscope showed that the seemingly clear button sources still suffered from some surface scratching. However, I consider this level of surface scratching negligible in terms of the overall source efficiency.

Activity determination

I used a high purity germanium (HPGe) detector¹⁴ to determine the activity of each button source, both before and after performing each leak test (see below). I placed one button source on a clean filter paper at the centre of the HPGe end cap, counting until the 1172 keV (for buttons 1-4) or the 1333 keV (for buttons 1A to 4C) gamma-ray peak contained 10^4 counts. Given the

¹⁴The detector is a GC12023 HPGe detector from Canberra Industries with a Canberra Lynx digital signal analyzer acquiring the data. I calibrated the detector response using a standard suite of gamma-ray calibration sources.



Figure 4.8: The $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ residue is clearly visible at $4\times$ magnification in this composite microscope image of button source 4C. The dashed black line is there only to guide the eye. Some features to note are the $\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ crystal structure toward the right and the dust fragments throughout (the long, smooth structures). The residue has largely maintained the liquid drop circular shape. Also visible toward the top left is the button edge, which appears darker, indicating a weaker bond at the edge. This button source is also one where a piece of dust reaches from the area of the residue to the edge (two dust pieces on the left side form a sort of bridge). The liquid deposition on this button was also off-centre, which is why the residue is nearer the button edge. Also, because this is the highest activity button source, the area over which the residue is spread is quite large.

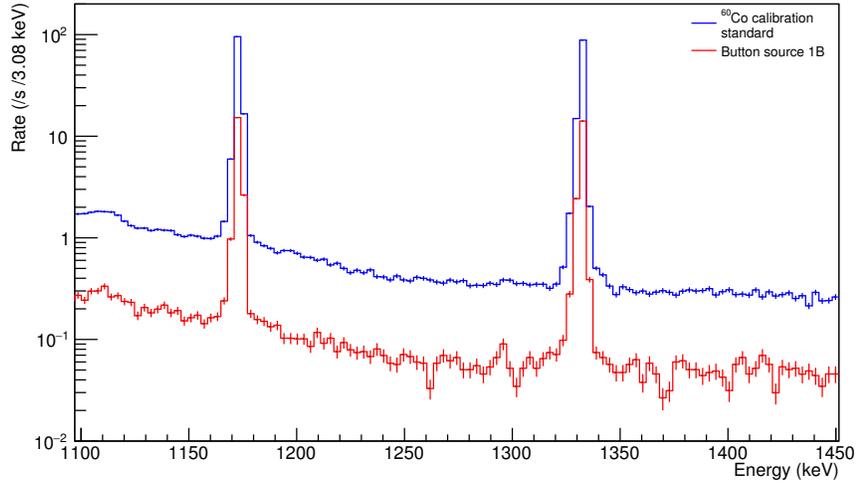


Figure 4.9: The button source 1B ^{60}Co rate is significantly less than that of the ^{60}Co calibration standard, which had an activity of approximately 2150 Bq on December 9, 2013, when this data was collected. Each set of data are background subtracted, using the same background data for the subtraction.

expectation that buttons 5, C1 and C2 would contain no activity, I counted these for a period of approximately 24 h. Figure 4.9 compares the background-subtracted ^{60}Co rates of button source 1B and the ^{60}Co calibration standard. Appendix B.2.3 outlines the exact procedure I followed to measure the activity of each button source.

The activity A of a source is

$$A = \frac{1}{I\epsilon} \frac{N - N_b}{\tau}, \quad (4.1)$$

where I is the intensity of the particular gamma ray line, ϵ is the detector efficiency, N is the total number of observed events, N_b is the number of expected background events and τ is the detector live time. This requires knowledge of the detector efficiency, which may be difficult to determine. In lieu of measuring the efficiency, I used a ^{60}Co calibration standard with known

activity ($10.73 \mu\text{Ci} \pm 1.9\%$, as specified by the manufacturer, in April 1974) to determine the unknown activity of the button sources. Taking a ratio between the activity of a source of unknown strength A_u and one of known strength A_c results in

$$A_u = A_c \frac{N_u \tau_c}{N_c \tau_u}, \quad (4.2)$$

where N_u and N_c are the number of observed counts (background subtracted) in the unknown and calibrated sources, respectively, and τ_u and τ_c are the corresponding measurement live times.

Using the ratio eliminates both the efficiency, assuming the efficiency of each measurement is identical, and line intensity. To ensure the efficiency was the same for each measurement, I counted the calibration standard placed atop an un-bonded half button. The ^{60}Co was then sitting at nearly the same position relative to the HPGe end cap as that in the button sources. There were small deviations in the exact source positions between measurements because I placed the source (either button or calibration standard) at the end cap centre by hand and because the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ within the button is physically extended. This introduced a systematic error on the determined activity owing to the position dependence of the detector efficiency.

To measure this systematic, I compiled four sets of counts, moving the source standard radially away from the nominal end cap centre in four different directions in steps of 10 mm. Figure 4.10 shows the resulting data, clearly indicating how the HPGe efficiency falls off toward the outer edges of the end cap. Taking the systematic error in the ^{60}Co position to be 8 mm (accounting for both the spatial extent of the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ and the hand-positioning of each button) results in a systematic error of 0.5%. This is added in quadrature

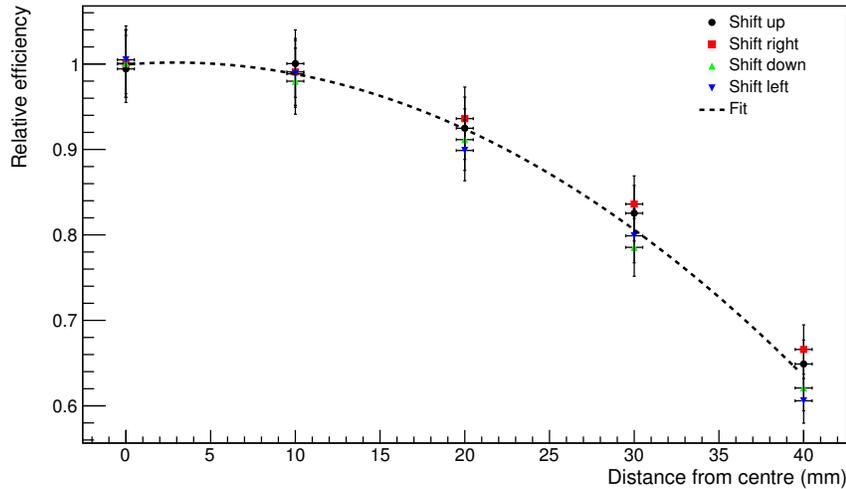


Figure 4.10: The efficiency of the HPGe degrades as the source moves farther from the end cap centre. I made four sets of measurements moving the ^{60}Co source standard outward from the end cap centre in four different directions (separated by $\sim 90^\circ$). The scale of the efficiency is relative to the average efficiency at zero distance from centre. The fit function is a quadratic, resulting in $\epsilon = (-2.7 \pm 0.5) \times 10^{-5}r^2 + (1.6 \pm 2.0) \times 10^{-3}r + (1.00 \pm 0.02)$.

to the statistical error for each activity measurement. Table 4.1 lists the measured activity of each button source on their measurement dates.

Leak testing

Upon completing the bonding of all the buttons, and as part of the laboratory decontamination process, I wiped each button by hand using a dry 4.25 cm diameter filter paper. I wiped buttons 1 through 5 with #5 grade filters and buttons 1A through C2 with #1 grade filters. The wipe test process consisted of repeatedly rubbing each surface of the button over the filter paper, ensuring I wiped each button surface with approximately equal vigour. Table 4.2 tabulates the results of the wipe tests of the buttons¹⁵. I calculate the

¹⁵Staff in the Radiation Safety Division of the University of Alberta's Environment, Health & Safety Office (EHS) counted all filters that were part of the decontamination process in

contamination level using

$$\frac{\text{Bq}}{\text{cm}^2} = \frac{R_s - R_b}{\epsilon_c \epsilon_w A \cdot 60}, \quad (4.3)$$

where R_s is the sample rate in counts per minute (CPM), R_b is the background rate in CPM, ϵ_c is the counter efficiency (assumed to be 50%), ϵ_w is the wipe efficiency (assumed to be 10%), and A is the wiped area in cm^2 (approximately 2.0 cm^2 for the buttons) [101]. Most wipes exhibit activities consistent with background rates, though the limited statistics make this a poor test of the true level of surface contamination present.

To more rigorously measure the level of surface contamination on each button and test the integrity of the encapsulation, I performed two additional tests¹⁶. The first was another wipe test, but using a filter paper moistened with RO water. The second (subsequent) was a soak test, where I soaked each button in a small volume of RO water at $45 \text{ }^\circ\text{C}$ for approximately 4 hours. I subsequently counted the test media (the filters or water) using the HPGe and compared each count to a background count consisting of an unused sample of the corresponding test medium. Appendix B.2.3 outlines the detailed leak test procedure.

Because I expect the amount of contamination to be zero, this value becomes difficult to measure on top of the background naturally present in the HPGe system¹⁷. Figure 4.13 shows an example set of data for which the

a scintillation Packard Cobra II Auto-Gamma Counter. The machine counted each filter for 2.0 min and supplied an integrated number of counts in the energy window 1050 keV to 1550 keV, suitable for searching for ^{60}Co contamination.

¹⁶I based these tests on the standard tests laid out in [102] and [103], with small modifications to accommodate this particular setup. Filters dried before being counted in all cases.

¹⁷There is some ^{60}Co present in the background HPGe spectrum, with a measured activity of $(2.58 \pm 0.06) \times 10^{-2} \text{ Bq}$, originating from the cosmogenic activation of the detector system's

Button Number	Number of Counts	Contamination (Bq/cm ²)
1	32 ± 5.7	0.33 ± 0.91
2	31 ± 5.6	0.25 ± 0.91
3	30 ± 5.5	0.17 ± 0.90
4	27 ± 5.2	-0.033 ± 0.870
5	25 ± 5.0	-0.25 ± 0.86
Background	28 ± 5.3	—
1A	23 ± 4.8	-0.25 ± 0.83
1B	24 ± 4.9	-0.17 ± 0.83
1C	32 ± 5.7	0.50 ± 0.90
2A	43 ± 6.6	1.42 ± 0.99
2B	24 ± 4.9	-0.17 ± 0.83
2C	32 ± 5.7	0.50 ± 0.90
3A	22 ± 4.7	-0.33 ± 0.82
3B	32 ± 5.7	0.50 ± 0.90
3C	17 ± 4.1	-0.75 ± 0.78
4A	39 ± 6.2	1.08 ± 0.96
4B	29 ± 5.4	0.19 ± 0.88
4C	19 ± 4.4	-0.58 ± 0.79
C1	33 ± 5.7	0.58 ± 0.91
C2	26 ± 5.1	0.00 ± 0.85
Background	26 ± 5.1	—

Table 4.2: Results of wipe tests using dry filter papers show the level of surface contamination on the buttons is consistent with background. EHS staff counted the filters for buttons 1 to 5 on April 23, 2013, and on December 3, 2013 for buttons 1A to C2. The table lists the corresponding background for each counting period below the button results.

contamination level is consistent with the background. I use two methods to determine the contribution from the contamination. The first method is to determine the upper limit on the activity of the contamination using the unified approach to construct a Feldman-Cousins interval. I use the number of background counts in a $\pm 3\sigma_E$ region about the expected gamma ray peak as the expectation, and the number of data counts in the identical region as the number of counts observed. One standard deviation, σ_E , is equal to the energy-dependent HPGe energy resolution, which we infer from calibration data, with σ_E and E in units of keV, as

$$\sigma_E = 0.00504\sqrt{E} + 0.000175E + 0.482. \quad (4.4)$$

Alternatively, for cases where the number of observed counts is significantly higher than the number of expected background counts (say, where the unified approach produces a two-sided interval), I may also calculate the activity under the gamma-ray peak by subtracting the background rate from the data rate, correcting for an overall baseline rate shift. I again determined the rates using the $\pm 3\sigma_E$ region. I estimate the baseline contribution using a linear sideband model, fitting the region outside of the peak region (from $\pm 9-36\sigma_E$) with a straight line, then subtracting the contribution of that baseline in the peak region. In tables 4.3 and 4.4 I calculated all one- and two-sided intervals using the unified approach and all values with \pm errors using the background subtraction method.

The result from the count of filters 1A-4C showed some signs of contamination, so I performed additional wipe tests on those buttons I suspected of

inner copper shielding layer.

exhibiting surface contamination (after working through a bifurcated search of the original filters). Table 4.3 reports the contamination levels I measured on the corresponding wetted filter for each button. For most filters, the detector's sensitivity limit is reached (\sim mBq), although filters 2A–2C and 4A–4C exhibit non-negligible amounts of contamination after the first wipe. Two additional wipe tests were done on all buttons to measure the reduction in the surface contamination the filters were able to remove. After three wipes, the measured contamination on buttons 2A–2C and 4A–4C is \lesssim mBq and at the detector's sensitivity limit. I re-measured the activity of each button between each wet wipe test, and they remained consistent with those I measured prior to performing the wet wipe tests.

Following the wipe test(s), I completed individual soak tests for each of buttons 1–5, C1 and C2, and a collective soak test for buttons 1A–4C, where I soaked all of the buttons together in the same volume of water¹⁸. For buttons 1–4, these tests used separate petri dishes filled with approximately 25 mL of RO water. The soak tests of buttons 1A–4C, 5, C1 and C2 occurred in approximately 15 mL of RO water in two 60 mL Nalgene containers with twist-sealing lids. Counting the water after removing the buttons indicated no contamination removal by the water for buttons 1–5, C1 and C2. However, the water that soaked buttons 1A–4C showed signs of contamination. We suspected the small amount of contamination in the water sample was likely surface contamination removed by the water.

To attempt to remove this supposed residual surface contamination, I subsequently washed the buttons¹⁹ by holding each button under a stream of run-

¹⁸My decision to do the collective soak test was based on the final wet wipe test results, which indicated that no contamination was being removed from the button source surfaces.

¹⁹Because the soak test showed higher levels of contamination removal than the wipe

Button number(s)	Activity (Bq)		
	Filter 1	Filter 2	Filter 3
1	$< 6.63 \times 10^{-3}$	—	—
2	$< 7.97 \times 10^{-3}$	—	—
3	$< 1.72 \times 10^{-3}$	—	—
4	$< 6.98 \times 10^{-3}$	—	—
5	$< 8.00 \times 10^{-3}$	—	—
1A	$< 7.83 \times 10^{-3}$	—	—
1B	$< 2.07 \times 10^{-2}$	—	—
1C	$< 4.47 \times 10^{-3}$	—	—
2A	$(0.418-1.48) \times 10^{-2}$	—	—
2B	$(0.524-4.68) \times 10^{-2}$	—	—
2C	$< 2.02 \times 10^{-2}$	—	—
3A	$< 2.87 \times 10^{-2}$	—	—
3B	$< 2.64 \times 10^{-2}$	—	—
3C	$< 2.11 \times 10^{-2}$	—	—
4A	$(0.512-3.93) \times 10^{-2}$	$< 1.67 \times 10^{-2}$	$< 5.17 \times 10^{-3}$
4B	$(1.80-4.20) \times 10^{-2}$	$< 1.94 \times 10^{-2}$	$< 1.38 \times 10^{-2}$
4C	$(3.61-5.30) \times 10^{-2}$	$< 2.73 \times 10^{-3}$	$(0.278-2.62) \times 10^{-3}$
C1	$< 4.16 \times 10^{-3}$	—	—
C2	$< 4.50 \times 10^{-3}$	—	—
1A-4C	$(1.10-1.21) \times 10^{-1}$	$(2.23-4.79) \times 10^{-3}$	$(0.545-3.92) \times 10^{-3}$
1A-2C	$(1.53-5.51) \times 10^{-2}$	—	—
3A-4C	$(0.567-1.16) \times 10^{-1}$	—	—
1A-1C	$< 3.57 \times 10^{-2}$	—	—
2A-2C	$(0.898-2.03) \times 10^{-2}$	$(0.012-8.29) \times 10^{-3}$	—
3A-3C	$< 2.29 \times 10^{-2}$	—	—
4A-4C	$(0.679-1.38) \times 10^{-1}$	$(0.291-1.26) \times 10^{-2}$	$(0.421-4.22) \times 10^{-3}$

Table 4.3: I counted the wetted filters used to wipe each source button for a period of several hours to several days, resulting in the 90% FC confidence limits on the residual activity on the filter papers shown. Button ranges correspond to a set of filters that I counted together.

Button number(s)	Activity (Bq)			
	Water 1	Water 2	Water 3	Water 4
1	$< 5.04 \times 10^{-3}$	—	—	—
2	$< 1.66 \times 10^{-2}$	—	—	—
3	$< 8.26 \times 10^{-3}$	—	—	—
4	$< 4.22 \times 10^{-3}$	—	—	—
5,C1,C2	$< 1.07 \times 10^{-3}$	—	—	—
1A–4C	$(4.67–10.64) \times 10^{-3}$	0.969 ± 0.023	—	—
3A–3C	—	—	1.067 ± 0.026	—
3A	—	—	—	$(0.39–1.80) \times 10^{-2}$
3B	—	—	—	0.082 ± 0.006
3C	—	—	—	0.822 ± 0.029

Table 4.4: I counted the water in which each source button was soaked for a period of several days. The 90% FC confidence limits on the residual activity in the water are shown, except for the individual tests of buttons 3A–3C, 3B and 3C, in which case I computed the activities by comparing to the ^{60}Co calibration source. To correct for assuming a detector efficiency relative to measuring a point source (like the calibration standards), we completed a Monte Carlo simulation of the water container within the HPGe detector system using Geant4. The efficiency of the detector for that geometry is 64.5% relative to a point source offset from the HPGe surface by 2 mm for both of the ^{60}Co gamma-rays.

ning water for 2 min, then drying with a clean Kimwipe. I repeated this three times for each of buttons 1A–4C. A second soak of buttons 1A–4C showed a significant amount of contamination remaining in the water, which led to a series of soak tests on various combinations of buttons in an attempt to identify which button was leaking. Table 4.4 lists the results of some of these tests.

It became clear that the soak tests were removing significant amounts of contamination from the buttons; however, the mechanism was unknown. It did not seem likely that it was surface contamination, because I had performed the wash under running water, and because the contamination level increased with each soak. This pointed to one of two scenarios: either the buttons tests, the thought was that washing the buttons in running water would be more effective at removing any surface contamination than by simply wiping them with a wet filter.

were leaking, with the water acting as a transport mechanism to transfer the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ out from inside the button along leakage channels, or the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ was diffusing through the plastic scintillator, with the water acting as a mobility aid. Visual inspections (including the use of the microscope) of the buttons after soaking did not show any indication of water being transported to the button centres. Water diffusion through BC-408 plastic scintillator is a known phenomenon [104]. Given this and the solubility of $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ in water, we suspected diffusion as the transport mechanism.

To test the diffusion hypothesis, I performed a series of soak tests at varying temperatures, where we expect higher diffusion rates at higher temperatures of the water [105, 106]. For this test I used button 4, as it had not undergone the same rigorous soaking process as buttons 1A–4C (so we would not expect significant diffusion to have occurred), but had a sufficiently high activity to determine if diffusion were occurring. Figure 4.11 shows the result of the temperature-dependent soak tests. I performed four consecutive soaks at each temperature (a total of 16 soaks), with each soak lasting approximately 4 h. An increase in the contamination level with temperature would be indicative of a diffusion mechanism. Because the level of contamination in the water does not increase with temperature, but seems to remain constant (and consistent with zero in the final tests), we conclude that diffusion is not the contamination transport mechanism.

To test the idea that enough washing may remove contamination from the button sources, I performed an additional nine soak tests on button 1B. Button 1B is the best candidate to be used in the final calibration source construction. Microscope images of this button, like figure 4.12, show no evidence of edge defects and the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ residue is clearly visible near the button's cen-

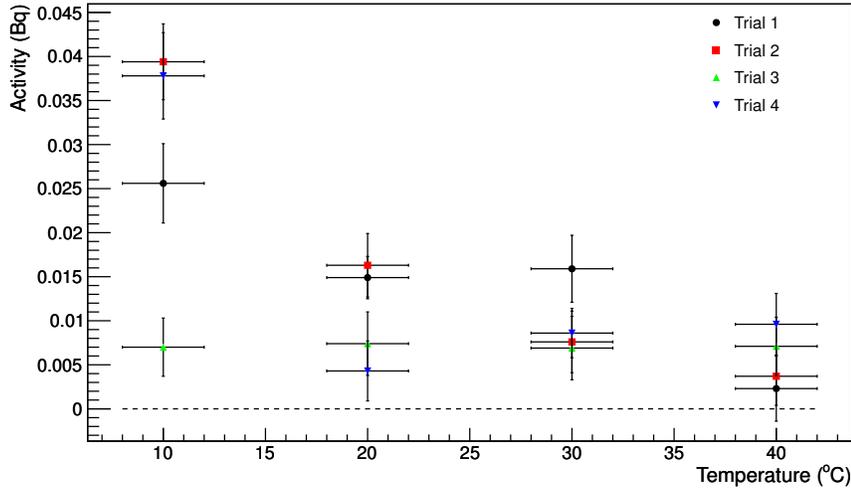


Figure 4.11: To test if diffusion was the mechanism responsible for transporting contamination out of the the button sources, I performed four soak tests of button 4 at each of four different temperatures. The data do not support a diffusion hypothesis, however, and seem to indicate that the contamination is washed away over time. The dashed line at an activity of 0 mBq is for reference only.

tre. In addition, the activity of this button will be approximately 180 Bq in 2018, which is near the desired activity of the source for the scintillator phase of SNO+. The first two soak tests (the tenth and eleventh soak test of button 1B) resulted in contamination levels only slightly above background, with activities in the ranges $(3.46\text{--}5.51)\times 10^{-2}$ Bq and $(1.06\text{--}2.80)\times 10^{-2}$ Bq, respectively, indicating a downward trend in the post-soak contamination level. The contamination level of the water from the ninth soak test, however, exhibits an activity of 7.95 ± 0.13 Bq, indicating that at some point, the integrity of the sealing process failed and allowed $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ to pass from the interior of the button into the water.

One final test I performed was to ‘soak’ button source 4 in the compound (BC-630) that will optically couple the button to the face of the PMT (see

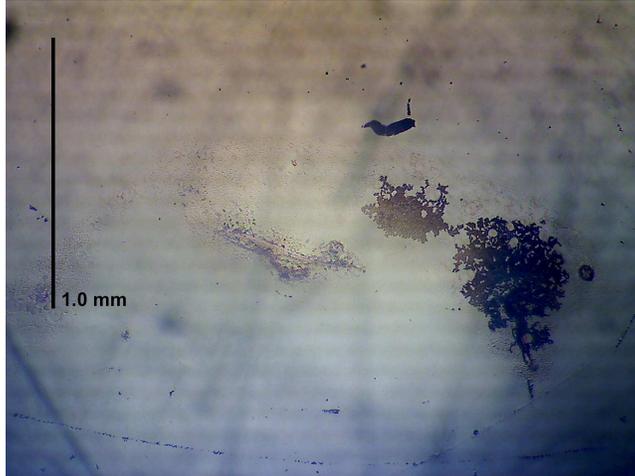


Figure 4.12: The $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ residue is clearly visible at $4\times$ magnification in button source 1B.

section 4.1.3)²⁰. The button source soaked in the compound for approximately 3 days in August 2014. I measured the contamination present in the compound to be $(2.20 \pm 2.69) \times 10^{-3}$ Bq, which is both consistent with no contamination being present and consistent with the results of this button source's last water soak test. Figure 4.13 shows the result of this test. I followed this with a soak of button source 1B in the same compound. (0.96 ± 0.02) Bq of contamination remained in the compound after removing the button source. While this is significantly lower than the contamination remaining in the water after this button source's last water soak test, the test clearly shows that contamination is transferable from the button source to the BC-630; however, it is not obvious from this test whether the compound only removed surface contamination from the button source, or if further leaking occurred.

I performed no further tests on any button source, and conclude that the plastic scintillator bonding method, while able to contain the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$

²⁰The button source will be in contact with BC-630, not water, when in place within the ^{60}Co calibration source, making this a useful test to determine whether contamination from the button source is transferable through the coupling compound.

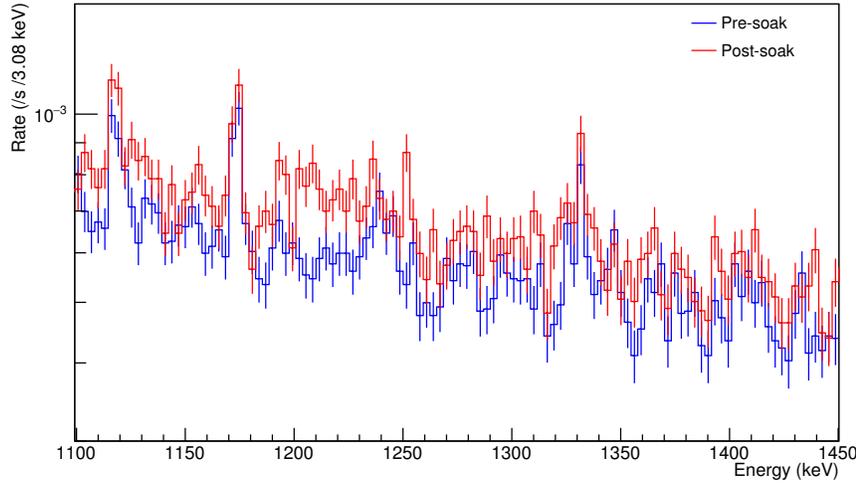


Figure 4.13: The HPGe spectra of the BC-630 prior to (blue) and after (red) soaking button source 4 in the compound show no indication that ^{60}Co contamination moved out of the plastic scintillator into the compound. The residual contamination in the BC-630 after removing button source 4 was consistent with 0 Bq. These data are not background subtracted.

within the plastic scintillator, is not effective as an encapsulation level for the ^{60}Co . Soaking the buttons in water leads to $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ escaping the plastic scintillator. The BC-630 optical coupling compound appears to allow contamination to escape the plastic scintillator, as well, though it is unclear whether the compound only removed surface contamination, or if it facilitated further leaking of the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ from within the plastic scintillator. I do not consider the plastic scintillator to be a level of encapsulation of the ^{60}Co .

4.1.3 Photomultiplier tube

When a ^{60}Co nucleus embedded in a plastic scintillator button (section 4.1.2) decays, the emitted electron will almost always deposit all of its energy within the button²¹ and result in the creation of scintillation light. To use this scin-

²¹I chose the dimensions of the button sources to ensure this.

tillation light as an indicator of a ^{60}Co decay, the button source is optically coupled to a Hamamatsu H10721-110P PMT [107]. This PMT is well suited to the spectral output of the BC-408 plastic scintillator, with a peak sensitivity at 400 nm and a fast pulse rise time of ~ 0.5 ns, making it suitable for photon counting. This PMT also requires only a low-voltage input power supply of +5 V (having a high-power supply circuit built within the module's metal housing), so will not require the SNO+ umbilical to carry a high-voltage power supply from the instrumentation deck.

The compound BC-630 [108] optically couples the button source to the PMT face. The face has a 10 mm diameter, with an active diameter of 8 mm, matching the button source diameter. To couple the two parts, I put a small amount of the compound on the PMT face using a wooden spatula, then press the button source onto the PMT face (holding the PMT face-up) until no air bubbles are visible in the interface between the plastic and PMT glass.

To increase the number of photons that reach the active part of the PMT face (thus, the source efficiency), I wrapped the face-end of the PMT in an (8×8) cm² sheet of 2 μm thick aluminized Mylar. The high reflectivity of the sheet helps direct photons toward the PMT face, while the air gap between the button source and sheet maintains the amount of total internal reflection within the button source. A piece of Kapton tape secures the Mylar to the PMT body. When complete, the Mylar is taut across the button source. This small amount of pressure keeps the button in place on the PMT face when the PMT is moved²², so I assume changes in the source efficiency resulting from displacement of the button source relative to the PMT face are negligible.

²²Without this pressure, the button is able to slide around in the coupling compound, though it does not separate from the PMT face because of the lack of air between the two.

Efficiency

While we accurately know the ^{60}Co decay rate of each button, that rate does not explicitly become the rate the source PMT will record. For example, the electronics threshold of the PMT readout may be too high to observe low numbers of photoelectrons (low charge pulses) produced at the photocathode. In particular, some ^{60}Co decays will produce an electron near the low-end of the electron energy spectrum, which will in turn produce only a few photons in the plastic scintillator that have a finite probability of striking the PMT face and subsequently producing a photoelectron. The efficiency of the ^{60}Co calibration source is equal to the percentage of ^{60}Co decays for which the source PMT observes the electron and SNO+ observes the corresponding gamma-rays. Knowing this efficiency will enable a sanity check of the SNO+ event rate during calibration and informs us as to how long a calibration run will need to last in order to achieve a particular level of statistical uncertainty.

While measuring this efficiency must ultimately occur in SNO+ (as it depends on the optics of the detector and the specific SNO+ electronics), it is useful to measure an approximate efficiency in the laboratory environment, particularly as a check that the calibration source functions as expected. To try and get a handle on the source efficiency, I measured the rate of PMT triggers coupling each button source to the PMT face with the PMT wrapped in its reflective Mylar layer. I supplied the PMT with a control voltage of +1.1 V (the maximum operating voltage) and used a custom-designed 4096 channel MCA²³ to read out the PMT trigger rate above a threshold of 175 ADC units (which I set high enough to cut out a low pulse height noise peak in the PMT

²³This MCA was designed in the Department of Physics at the University of Alberta.

data). I used a 7 dB attenuator on the PMT signal to ensure the PMT response spectrum fit within the voltage range of the MCA.

Figure 4.14 shows the rates the MCA measured for each button source versus the button source activities determined using the HPGe (corrected to the measurement date) in table 4.1. Fitting these data to a straight line gives

$$R_{\text{MCA}} = (0.703 \pm 0.004)A - (6.761 \pm 0.388) , \quad (4.5)$$

where R_{MCA} is the rate of the button source the MCA measures in units of Hz. This indicates an efficiency of the PMT of approximately 70%; however, this value is highly dependent on the MCA threshold, which I did not explicitly tune. In another test, we optically coupled button 3C directly to two PMTs²⁴ and measured a total PMT rate of approximately 650 Hz, which is close to the HPGe-measured activity of 657.9 Bq.

Thus, I conclude the source PMT is at least 70% efficient at observing the ^{60}Co decays inside the button source. Again, the true efficiency will depend on the SNO+ electronics, and also the efficiency with which it detects the gamma-rays from the ^{60}Co source.

PMT DAQ and trigger

The plan for the ^{60}Co source PMT, as with other calibration sources containing PMTs, is to connect it to a dedicated channel on the SNO+ FECD (front end card diagnostics) card. This card has logical channel numbers (LCNs) inclusively ranging from 9184 to 9215 (corresponding to card 15 in crate 17,

²⁴Section 5.2.3 describes these PMTs in more detail.

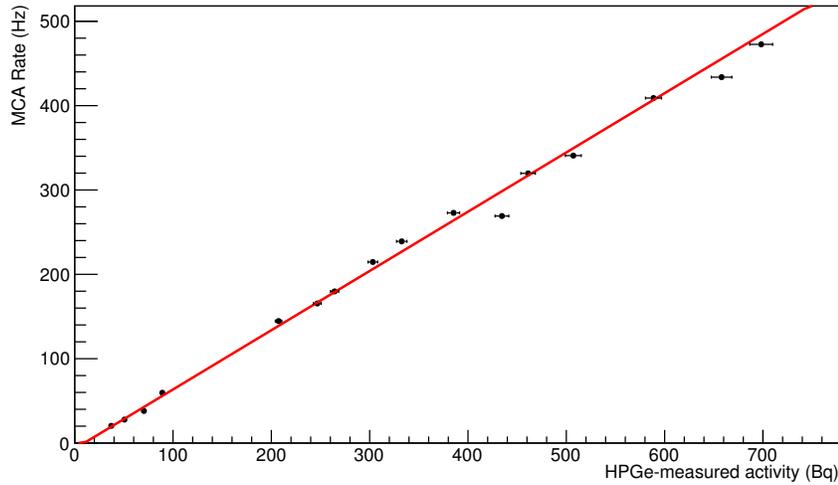


Figure 4.14: The PMT trigger rate is linearly proportional to the HPGe-measured activity of each button source. The MCA I used in this measurement leads to a source efficiency of approximately 70%, which comes from the slope of a linear fit to the data (red).

channels 1–32). Each calibration source will have a dedicated LCN²⁵, so signals from the source PMT will stream through the SNO+ DAQ in the same manner as the SNO+ PMTs.

Although the possibility exists to use the signal from a calibration source PMT to provide a trigger for SNO+, the ^{60}Co source will not use this feature. Rather, events in SNO+ for which the ^{60}Co source PMT also records an event will be identified offline by temporally correlating events (tagging) between the two detectors.

²⁵Channel 6 (LCN 9190) is dedicated to reading the unlatched raw trigger signal, which is then used to correct PMT hit times for the 20 ns time jitter associated with the latching of the global trigger signal to a 50 MHz clock. Thus, this channel does not correspond to any physical calibration source.

4.1.4 Encapsulation

Owing to its size, radioactive calibration sources, like the ^{60}Co source, must be deployed directly within the SNO+ LS volume. As a result, it is critical that all calibration sources are built in a robust way that guarantees there is no pathway for radioactive contamination to make its way into the LS. To ensure this, the SNO+ Calibration Source Review Committee²⁶ requires that each calibration source must utilize at least two independent encapsulation mechanisms. The job of the encapsulation is two-fold: it must prevent the spread of contamination outward from the calibration source and it must prevent the ingress of LS into the calibration source, as the LS could provide a path for the transfer of contamination into SNO+. These encapsulations should be both liquid and helium leak-tight.

The ^{60}Co calibration source satisfies this double encapsulation requirement by placing the PMT assembly within a pair of independent, nested containers²⁷. The inner container is formed from a thin sheet of copper and is intended to prevent any $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ from making its way out of the metal container. The outer container (and stem) is made from Delrin[®]. This plastic is compatible with the LS, easily machinable, and durable enough to not break under mechanical upset. Its purpose is to prevent the ingress of the LS into the calibration source. The subsections below describe these components in detail.

²⁶This committee ultimately reviews the calibration source motivation and design and decides whether the source is acceptable for deployment within SNO+. A second committee, the SNOLAB Radiological Source Committee (distinct from the SNO+ committee), approves the use of radiological sources within the SNOLAB environment, focussing on risks the source presents to other SNOLAB experiments and personnel.

²⁷I do not consider the plastic scintillator to be a layer of encapsulation.

Inner encapsulation

The inner encapsulation consists of a thin-walled copper sleeve, sized to snugly house the PMT assembly, attached to a copper flange that incorporates a feedthrough for the PMT electronics. The sleeve is constructed from $\frac{5}{1000}$ " (0.13 mm) thick copper sheet (about the thickness of a soda can), 67.5 mm in length. It is bent into an open-ended square rectangular prism with side length 23.75 mm. A 3 mm overhang enables soldering of the two unjoined edges of the box. A square piece of copper with 23.75 mm long edges, with an additional 3 mm overhang on each, is soldered to one open end of the copper sleeve. Figure B.1 contains the sleeve technical drawing.

The solder, used for all solder joints on the copper encapsulation, is lead-free (for lower radioactive backgrounds) Castolin 157 Eutectic solder and liquid flux²⁸. This low-temperature (activating near 200 °C) solder is highly corrosion resistant and has high mechanical strength, and is often used for ultra-high vacuum (UHV) applications.

A 43.23 mm diameter copper flange is soldered to the open end of the copper sleeve. The bottom half of the flange is soldered directly to the copper box, the top half being secured to the bottom half with eight screws. Each half of the flange is 2.5 mm thick. A 1.0 mm wide, 0.75 mm deep, 30.76 mm inner diameter groove is cut into the top side of the bottom flange, and a matching 0.5 mm deep, 0.8 mm wide lip protrudes from the bottom of the upper flange. On the underside of the bottom flange is a 1.0 mm wide, 2.5 mm deep lip onto which the copper sleeve is soldered. A 1.2 mm wide, 3.0 mm tall, 13.3 mm inner diameter lip on the top side of the upper flange acts as a soldering surface to

²⁸Available from <https://www.castolin.com/product/157-0>.

connect a six-pin glass feedthrough for the PMT electronics²⁹. A set of eight holes enable 4-40 stainless steel screws and nuts to connect the two halves. Four additional holes allow us to secure the copper encapsulation in place within the outer Delrin[®] encapsulation. Figure 4.15 shows the assembly of the copper encapsulation. Figures B.2 and B.3 contain the technical drawings for the bottom and top half of the flange, respectively.

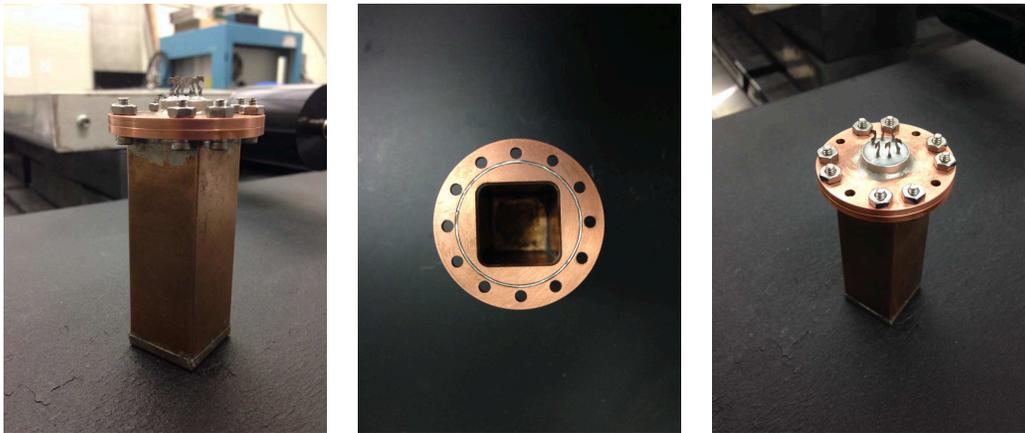


Figure 4.15: The copper encapsulation consists of a thin-walled copper sleeve and copper flange that is sealed using an indium wire. The sleeve (and its solder joints) are visible in the left image. The right image shows how the screws are used to connect the two flange halves and also shows the glass electronics feedthrough. The middle image shows the indium wire in place before attaching the upper flange.

To ensure a helium leak-tight seal at the copper flange, we set a 0.025” diameter 99.99% pure indium wire³⁰ into the groove in the bottom half of the copper flange. The wire is long enough so its two ends overlap slightly (by about 1 mm). When we join the two halves of the flange, the lip on the upper flange crushes the soft indium wire, causing it to fill the gap between the

²⁹These feedthroughs are legacy pieces I obtained from the University of Alberta’s Department of Physics electronics shop. Nothing is known of their origin, except that they are likely around 30 years old (though never before used).

³⁰Available from <http://www.indium.com/>.

lip and groove. This creates an effective seal, completely isolating the PMT assembly within the copper encapsulation.

We performed two tests to investigate the effectiveness of the indium wire sealing technique. First, we assembled the copper encapsulation in the laboratory environment. Opposing screws were tightened in succession using half-turn increments until they were wrench-tight. We then placed the copper in a beaker of water and placed this inside an acrylic vacuum chamber, as figure 4.16 shows. We did not observe any air bubbles escaping from within the encapsulation while pumping on the copper (which lasted approximately 10 minutes).

The mass of the copper prior to submerging it in water was 73.80 ± 0.01 g. After removing it from the beaker of water, we patted the copper dry with a Kimwipe and allowed a further 2 h of air drying time to allow residual water sitting on the copper surface to evaporate. The mass of the copper after drying was 73.85 ± 0.01 g. We repeated this test, allowing the copper to sit under water in the vacuum overnight. After removing the copper from the water and drying, the copper mass was 73.82 ± 0.01 g, implying no water entered the copper encapsulation. Upon taking the flange apart, we noted that no water was present on the inside of the encapsulation. Thus, we deemed the indium wire sealing technique to provide a water-tight seal.

The second test we performed was to investigate whether the indium wire provides a helium leak-tight seal. For this test, we assembled the copper encapsulation inside a helium atmosphere³¹, as figure 4.16 shows, so the inside of the encapsulation would be filled with helium. After removing the sealed

³¹We used an overturned plastic container and flooded it with helium until we were able to detect helium outside of the container using our helium leak-checker.

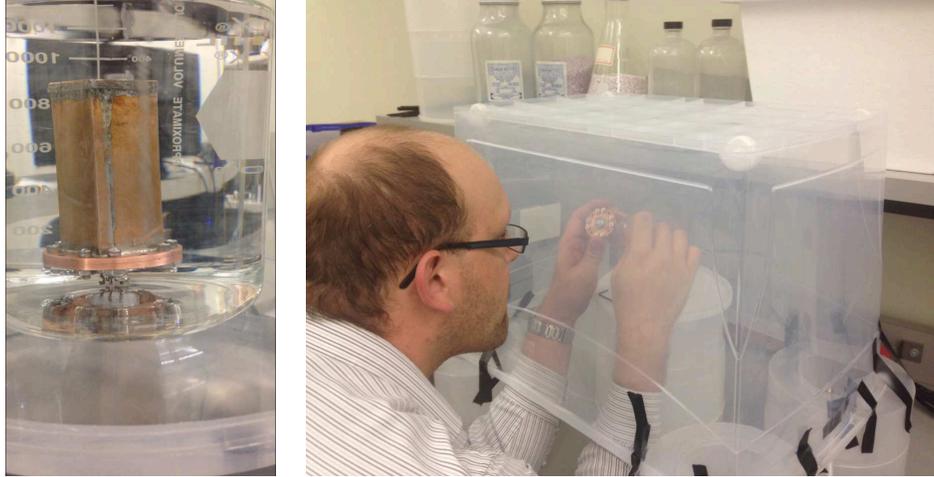


Figure 4.16: The copper encapsulation was submerged in water to see if we would observe bubbles escaping it when putting it under vacuum (left). This was done both by assembling the copper in the regular laboratory environment and by assembling it in a helium atmosphere (right), whereby we were able to determine that the indium wire provides a helium leak-tight seal. Note that gloves were not worn while handling the copper during these tests, but are required for the actual source assembly in appendix B.2.4.

copper from the helium, we used a helium leak-checker³² to sniff around the indium seal. There was no indication of helium escaping the encapsulation. We repeated the water test, as above, and again did not observe bubbles escaping the copper. After drying, we opened the flange while using the sniffing near the indium seal with the helium leak-checker. In uncalibrated units of mbar L/s, the helium level went from 1.5×10^{-8} when the flange was sealed, to 9.9×10^{-8} as the screws were loosened (only changing once all screws had been loosened), to $\sim 10^{-2}$ as we opened the flange. The conclusion from this test is that the indium wire sealing technique provides a robust, helium leak-tight seal and will thus fully encapsulate the $^{60}\text{CoCl}_2 \cdot 6\text{H}_2\text{O}$ within the copper.

Lastly, the PMT assembly will be encapsulated using this indium wire technique *in situ*, so is connected to the glass feedthrough before the flange

³²The leak-checker is an AdixenTM ASI 20 MD from Alcatel Vacuum Products, Inc.

halves are connected. An electronics technician solders each PMT wire to one of the six feedthrough pins using InterTAN 64-027 0.05" lead-free, rosin core solder³³. Shrink wrap and teflon tubing protect the wires and ensure shorts between wires do not occur.

Appendix B.2.4 contains the detailed assembly procedure for the copper encapsulation, including the PMT soldering procedure and leak-testing procedures.

Outer encapsulation

A second encapsulating layer made of Delrin[®], pictures of which are in figure 4.17, fully contains the copper encapsulation. Somewhat like the copper, the outer encapsulation layer is formed by a container and a stem, connected together at a flange. The container is an 88.17 mm tall cylinder with a 2 mm wall thickness. The top portion of the container is one-half of a flange with outer diameter 68.78 mm. Eight screw holes are evenly spaced about its circumference and encircle an o-ring groove with inner diameter 50.52 mm, sized for an O33 o-ring. The main body of the container has an inner diameter of 43.4 mm, just large enough to accommodate the copper encapsulation. Inside the container, 58.0 mm up from the bottom is a 10.0 mm thick collar containing four screw holes that allow the copper encapsulation to be secured.

The other half of the flange is part of the stem that connects the source to the SNO+ umbilical (the umbilical carries the electronics and other lines that allow calibration sources to communicate with the instrumentation deck). The stem, also made from Delrin[®], is 309.74 mm in length, with a 5.0 mm thick

³³This solder is no longer available for purchase, but is typically used for electronics work, particularly where signal quality is important.

flange at the end connecting to the container and a 9.53 mm thick flange at the other end that will allow the stem to attach to the SNO+ source connector (see below). The purpose of the stem is to move the ^{60}Co decays farther from the hardware that carries the calibration sources in the detector. This bulky hardware blocks light produced in the LS by the calibration sources, so moving the radiation production point away from this hardware helps minimize any shadowing effects.

The stem has a 5.0 mm radius central bore to accommodate the wiring for the PMT. It has an outer diameter of 31.76 mm, but this tapers at an angle of 3° as it nears the source container, to an outer diameter³⁴ of ~ 22.0 mm. This is to reduce the solid angle of the stem with respect to the ^{60}Co decay point.



Figure 4.17: The Delrin[®] encapsulation consists of a container with 2 mm thick walls that houses the copper encapsulation (left). Eight screws connect this container to a stem at a flange (middle), which compresses an internal o-ring to provide a leak-tight seal. When connected (right), the Delrin[®] layer is nearly 398 mm in height. This does not include the source connector that will attach to the top of the stem.

³⁴As figure 4.17 shows, the transition from the stem to the container-end flange is rounded as a byproduct of the machining process. While it adds a small amount of extra material onto the stem, it will also improve its structural integrity, as this transition region is the thinnest area of the encapsulation.

The top end of the stem will connect to one half of the SNO+ source connector shown in figure 4.18. The source connector consists of two stainless steel cylinders that screw onto a central nut, each threaded with opposite handedness. O-rings on the nut are compressed once the screws are tightened, providing an effective leak-tight seal when the source connector is closed. Inside the source connector are stainless steel plates that hold the various hardware connections the calibration sources require (like electronics cables, gas lines and optical fibre bundles). Each calibration source will be capped with a dedicated connector, with the nut and other half of the connector belonging to the umbilical. The purpose of the source connector is to make source connection to the umbilical easy, in particular as sources will be interchanged inside a glove-box where visibility is limited and tactile feedback is not ideal.

Currently, prototypes of the source connector exist; however, exactly how the ^{60}Co calibration source stem will connect to its dedicated connector part is not yet defined. Since we are relying on the connector part to provide, along with the stem and container, the outer level of encapsulation for the source, it must be sealed even when it is not connected to the nut. My current idea is that, for the ^{60}Co source, a stainless steel plate may be welded to the connector and the stem connected to it using screws, compressing a single or double o-ring in the process. Soldering the same type of glass electronics feedthrough as the copper encapsulation uses into the centre of this steel plate would enable the passage of the PMT electronics into the source connector, while that feedthrough and the o-rings will provide a leak-tight seal. I have not physically tested this design, though.

Appendix B.2.4 contains the detailed assembly procedure for the Delrin[®] encapsulation, including proposed leak-testing procedures. Because the con-

nection of the stem to the source connector is not completely defined, I have not leak-tested the outer encapsulation. The source connector is an integral part of that encapsulation, so only once it is in place on the stem will the leak-tests be completed.

As a final note, the Delrin[®] used to form the outer encapsulation is black to ensure it is opaque and will not allow light to enter or exit the source. This choice is mainly due to earlier designs of the ⁶⁰Co source that used a transparent potting compound to seal the PMT within the Delrin[®] directly, which may have provided an optical pathway for scintillation photons produced in the button to reach the outer encapsulation. With the change to the copper internal encapsulation, it is no longer strictly necessary to make the outer encapsulation and stem out of the black Delrin[®], but Delrin[®] is relatively easy to machine and is compatible with the LS (see below), so there is really no incentive to change it to a different material.



Figure 4.18: The source connector consists of two cylinders that screw on to a central nut (left). An internal plate hosts connections each source requires to communicate to the instrumentation deck (visible right), like electronics and gas connections.

Material compatibility

Because the ^{60}Co source will be bathed in the SNO+ LS during calibration data collection, it is critical that the outer encapsulation be chemically compatible with the LS. Technicians at SNOLAB tested samples of the Delrin[®], Viton[®] o-rings and stainless steel nuts and screws to investigate the compatibility of these materials with a sample of the LS they prepared. Figure 4.19 shows a picture of these materials soaking in LS samples. Materials comprising the inner encapsulation have not been tested³⁵.



Figure 4.19: Samples of the Delrin[®], Viton[®] o-rings and stainless steel nylon-insert lock nuts soaked in LS for seven months. SNOLAB technicians performed UV/VIS scans at various times to check for changes in the LS absorbance, which would indicate material incompatibility.

The materials soaked in 100 mL samples of the LS for seven months, with periodic testing during that time. The test performed was ultraviolet-visible absorption spectrometry (UV/VIS) using a Perkin-Elmer Lambda 800 UV/VIS Spectrometer. The LS samples were scanned inside quartz cells, with

³⁵The purpose of the outer encapsulation is to prevent the ingress of the LS into the calibration source, so internal materials should not come into contact with the LS unless the outer encapsulation fails.

a control sample of pristine LS scanned for comparison. Figure 4.20 shows how the wavelength-dependent absorbance of each sample changed as a function of time. The conclusion of these scans is that the Delrin[®] appears compatible with the LS, though a small change in absorbance was observed after more than seven months of soak time. The Viton[®] appeared to remain compatible after two months of soak time, but dramatically changes at the seven month test, indicating a different material o-ring may be required³⁶. The nylon in the lock nuts causes the absorbance in that LS sample to continually worsen, indicating that nylon is incompatible with the LS. A change to pure stainless steel lock nuts is warranted.

4.2 Simulation

To simulate the SNO+ response to ⁶⁰Co calibration events, I have created a simulation of the ⁶⁰Co source within the SNO+ simulation and analysis software framework, RAT (see section 2.1.2). The simulation includes both a description of the ⁶⁰Co calibration source geometry, based on the as-built structure, and a custom event generator that simulates ⁶⁰Co decay from within the plastic scintillator button source. Section 4.2.1 describes in detail the simulated source geometry, including how RAT simulates the PMT electronics. This section also gives reference to early simulation results using a simpler geometry and which were ultimately used to inform the current geometry. Section 4.2.4 details the ⁶⁰Co source generator, including validation of the nuclear decay simulation. I present simulation results using this source geometry and generator in section 4.2.5.

³⁶Viton[®] is FKM type fluoroelastomer. Similar tests of Perlast[®] FFKM perfluoroelastomer show compatibility with the LS over periods longer than seven months.

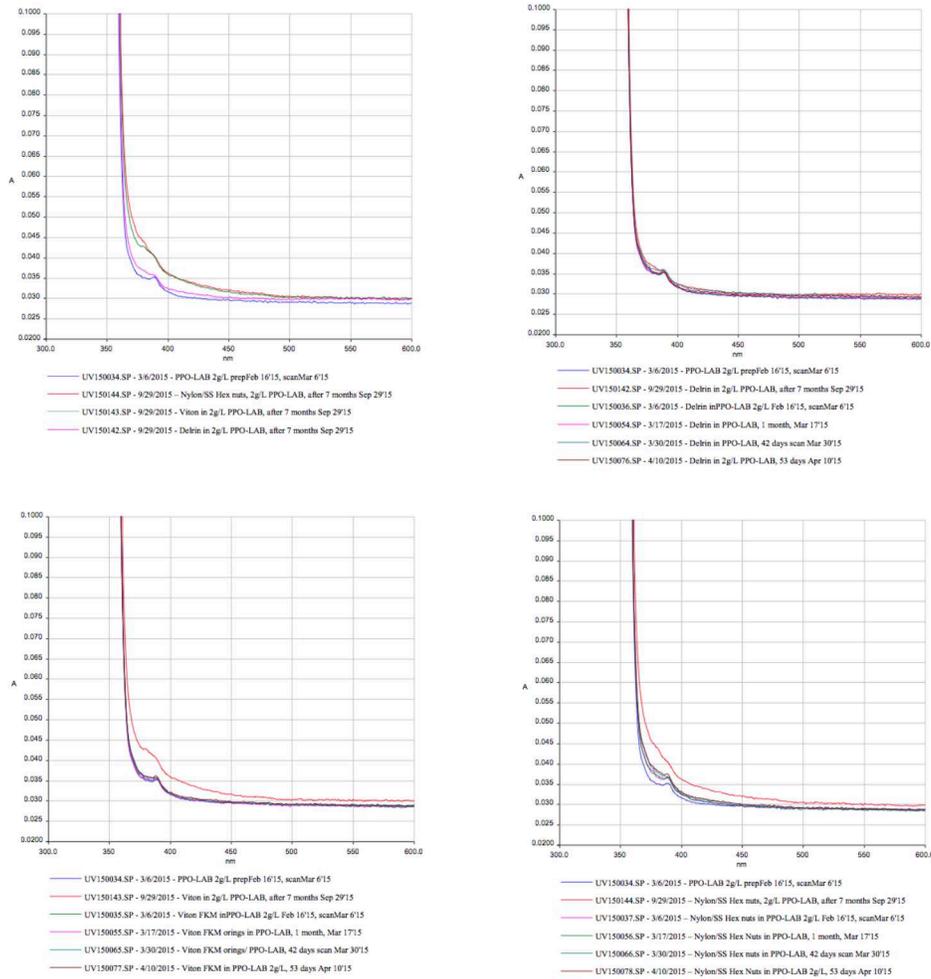


Figure 4.20: Results from the UV/VIS scans show the Delrin[®] (top, right) remains compatible with the LS, whereas the Viton[®] o-rings (bottom, left) exhibit compatibility until the seven month mark. The nylon-insert lock nuts (bottom, right) show a steady degradation of the absorbance over time. Data from the seven month soak test comparing those three samples to the pristine LAB sample (top, left) show these trends clearly. The wavelength in nm increases to the right on the x-axis, while the absorbance is on the y-axis.

4.2.1 The simulated ⁶⁰Co source geometry

The geometry of the ⁶⁰Co calibration source in RAT mimics its design drawing (see appendix B.1) in most respects³⁷. It consists of the source container,

³⁷Some minor differences are inevitable, sacrificing exactness for simulation efficiency.

stem, copper enclosure, PMT, plastic scintillator button, o-rings, and screws and nuts. All of the source dimensions and other required parameters, such as the source identification and version numbers, location &c., are accessed via the SNO+ database (`GEO` with index `Co60Source`). This allows for flexibility in specifying the source geometry (with minimal hard-coding of values). The geometry factory class `GeoCo60SourceFactory` builds the source geometry in its entirety and places it within the simulated SNO+ detector. Figure 4.21 shows a cutaway view of the ^{60}Co source geometry, as produced by Geant4.

The source geometry is constructed relative to the centre of the plastic scintillator button, which is the local geometry origin. This is the reference point for placing the source geometry within the larger SNO+ geometry and is where the ^{60}Co source generator (see section 4.2.4) generates ^{60}Co decays. The cylindrical plastic scintillator volume, 4.0 mm in radius, 4.0 mm thick and made from `G4_PLASTIC_SC_VINYLTOLUENE`³⁸, is centred flush to the active face of the PMT. The PMT face is a cylinder of glass (material `glass`) 5.0 mm in diameter and 0.1 mm thick (an arbitrary thickness), and is inset 1.5 mm into the main body of the PMT. The PMT is a square face aluminum (`aluminum`) box 50.0 mm long with 22.0 mm face side length.

The PMT and scintillator button are set inside a 0.0127 mm thick copper (`G4.Cu`) box. The box is 67.5 mm long and has a face side length of 23.75 mm. A cylindrical copper flange, 2.5 mm thick and 21.615 mm in radius, caps the top of the box. The underside of this flange has a 1 mm thick, 5 mm tall square lip with side length 23.49 mm (so it fits within the copper box, which is flush with the bottom of the flange). A 3 mm thick, 6.65 mm radius

³⁸In this section, simulated Geant4 materials are in `type` face, where Geant4 pre-defined materials begin with `G4`. Materials in regular face are physical materials.

glass cylinder is set in the top of the flange to model the glass electronics feedthrough. Eight screw holes are located along the flange circumference at a radius of 23.49 mm. These accommodate stainless steel (`stainless_steel`) screws. A groove containing indium (`G4_In`) is set at the centre of the copper flange, with radius 15.83 mm and depth 0.75 mm.

The copper enclosure is mounted inside a Delrin[®] (`G4_POLYOXYMETHYLENE`) outer container. It is a tube with a uniform wall thickness of 2.0 mm, length 73.0 mm and outer radius of 23.7 mm. At the top of the container is a 9.0 mm thick flange that extends 10.69 mm beyond the container outer radius and is flush with the container inner radius. This flange connects to the tube via a tapered section that is 6.17 mm in height (making the entire container 88.17 mm in height). A 10 mm thick internal collar 58.0 mm up from the insider of the container base acts as an attachment point for the copper enclosure, and contains a matching set of screw holes. It has a square hole cut from its centre to accommodate the copper box. A groove in the top flange with inner radius 25.26 mm, width 2.26 mm and depth 1.27 mm exists to hold an o-ring. The o-ring is comprised of Viton[®] (`G4_VITON`) and fills the entire o-ring groove. The flange also contains a circumferential groove to accommodate nuts that aid in attaching the container to the source stem. It is 7.35 mm deep, 4 mm tall and begins 3.5 mm below the top of the flange. By default, the flange contains eight equally-spaced screw holes. The screws, with cylindrical bodies and a cylindrical cap, and nuts, also cylinders, are comprised of stainless steel.

The stem is a Delrin[®] cylinder with a flange on both ends. The flange that connects to the source container is 5.0 mm thick and has the same radius as the flange on the container. It has a 5.0 mm radius opening in its centre that extends as a bore hole the entire length of the stem to allow for the passage of

wiring from the PMT (the wires are not simulated). Holes in this flange at the same locations as those in the container accommodate the screws. The stem at this flange has a radius of 11.0 mm for a length of 8.74 mm, at which point it tapers at an angle of 3.0° until the stem reaches a radius of 15.88 mm. It then extends to a second flange with thickness 9.53 mm and radius 29.0 mm that attaches to the source connector. The total stem length is 309.74 mm.

The source connector is modelled as a stainless steel cylinder, 46.0 mm in radius and 123.5 mm in height, with a wall thickness of 9.5 mm. A 5.0 mm thick plate sits at a height of 50.0 mm from its base. The source connector sits flush at the top of the stem.

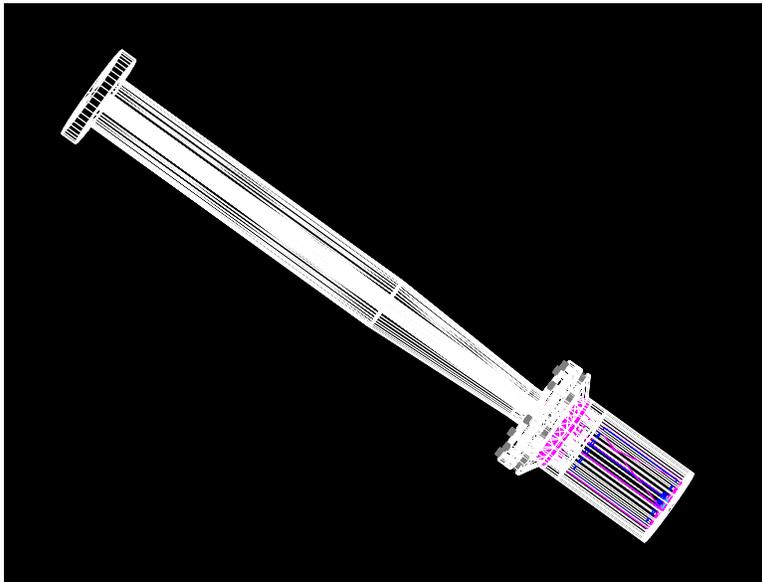


Figure 4.21: This wireframe image of the ^{60}Co source built by Geant4 shows the Delrin[®] stem and container as white, the copper as magenta, the PMT as blue, the scintillator button and glass feedthrough as cyan, the o-ring as red (just visible) and the screws and nuts as grey. The source connector, while simulated, is not drawn here.

All of the parameters that characterize the ^{60}Co source are read from the SNO+ database (GEO with index Co60Source). Most of these describe the

source geometry, so will not typically be altered; however, there are a handful of parameters that a user may need to change. The source version number (which accommodates the possibility of having more than one version of the ^{60}Co source in the future) is controlled through the database entry `version`. The boolean entry `check_overlaps`, if true (default), causes Geant4 to check whether any of the source volumes overlap when constructing the source³⁹. The volume into which the ^{60}Co source geometry is placed (called its mother volume) is controlled by setting the string `mother` (by default, it is set to `inner_av`, the SNO+ active volume). The array `sample_position` sets the location of the source (the position of the centre of the plastic scintillator button) relative to the coordinate system of its mother volume (set to `[0., 0., 0.]` by default). Lastly, the boolean entry `screws_enable` allows the user to decide whether the screws, nuts and screw holes are included in the geometry. This level of detail may not be required by all users, though these volumes are included in the source by default. Table B.1 contains a description of every database value used to build and/or characterize the ^{60}Co source.

4.2.2 The PMT simulation

A key feature of the ^{60}Co source is an internal PMT that enables calibration events to be tagged when analyzing data, thus producing a nearly background-free calibration data set. The simulation of the ^{60}Co source PMT, while important, need not be as complex as that of the SNO+ PMTs, meaning that a full optical simulation of the source and its PMT is not necessary. Simulating the PMT at this level would require detailed studies of the PMT response

³⁹This does not, however, cause Geant4 to check whether the source overlaps with other SNO+ volumes.

which are simply not practical, nor which would add to any additional understanding of the source response. What is measured in the laboratory is the overall efficiency of the ^{60}Co source in detecting light from a single ^{60}Co decay. Thus, the simulation should decide whether or not the source PMT records a hit based strictly on the knowledge of this measured PMT efficiency.

To accomplish this, a sensitive detector class, `CaLibPMTSD`, feeds a response from the source PMT to the simulation of the data acquisition (DAQ). Unlike the SNO+ PMTs, the plastic scintillator, rather than the source PMT, is the sensitive detector. If more than some threshold energy is deposited in the plastic scintillator in an event, then a uniform random number on $[0, 1]$ is checked against an efficiency value (also on $[0, 1]$). By default, the energy threshold is set to 0 MeV and the efficiency of the source⁴⁰ is set at 0.90, but users may alter these via the database values

```
/rat/db/set GEO[Co60Source] source_efficiency 0.95
/rat/db/set GEO[Co60Source] threshold_energy 0.
```

If the random number thrown is less than the efficiency, then, by default, 10 photoelectrons are detected on the channel corresponding to the source PMT. Creating 10 photoelectrons essentially guarantees that each ^{60}Co event will be above the channel threshold and the event recorded. Physically, the number of photoelectrons created in the PMT will be a function of the number of photons incident on the PMT's photocathode, which is itself a function of the true energy deposit in the plastic scintillator, as well as the efficiency with which the generated photons are transmitted to the photocathode. As these efficiencies are folded into a single measured value, creating many photoelectrons simply

⁴⁰This will become the value measured in the laboratory, once available.

ensures events pass the channel threshold⁴¹.

The DAQ records the signal from a calibration source PMT on a dedicated FECD channel, the default of which is 9191 for the ⁶⁰Co source. Users may change this by setting the database value

```
/rat/db/set GEO[Co60Source] lcn 9191
```

Attempting to set the LCN outside of the allowed FECD channel range causes the simulation to abort. It is important to note that, in the simulation, only the fact that a FECD channel is hit is useful because of the arbitrary choice to create 10 photoelectrons per hit (which is not representative of the true charge collected on that channel).

4.2.3 Early geometry optimization studies

To determine the optimal design for the ⁶⁰Co calibration source, I used a simple calibration source geometry, the so-called “can source”, to inform the final geometry, particularly in terms of building materials and the wall thickness of the source container. The can source is a simple nested geometry consisting of a PMT with base, a conical light guide, a cylindrical plastic scintillator button and a pair of nested containers enclosing the pmt, light guide and button. I modelled the PMT as a Hamamatsu R1635 PMT and base assembly⁴² [109], where the PMT has a radius of 10.0 mm and height of 45.0 mm, and the

⁴¹This only becomes an issue when a ⁶⁰Co decay deposits very little energy in the plastic scintillator, so that the number of photoelectrons created is small compared to the default value created. Nevertheless, this effect is expected to be negligible as the number of photons generated in the plastic scintillator should be above 300, on average.

⁴²Given its small size and relatively fast response, I originally chose the Hamamatsu R1635 PMT assembly for the ⁶⁰Co source; however, I subsequently changed it to the H10721P-110 low-voltage PMT option to be consistent with other calibration sources in development.

base has the same radius and a height of 50.0 mm. To ensure the PMT and base have the correct masses (according to data for the R1635), the PMT is comprised of glass (material type `cansource_pmtglass`) with a density of 1.69 g/cm³ and the base is comprised of polypropylene (`cansource_pmtbase`) with a density of 2.26 g/cm³.

The plastic scintillator button (material `G4_PLASTIC_SC_VINYLTOLUENE`) has a radius of 10.0 mm and a thickness of 4.0 mm, by default. The inner can is a solid cylinder that contains the PMT, PMT base, light guide⁴³ and button. It has a default thickness of 2.0 mm and radius of 12.0 mm. The outer can is a solid cylinder containing the inner can and has a thickness of 2.0 mm and radius 14.0 mm. Both cans are comprised of `G4_POLYOXYMETHYLENE` (Delrin[®]). Figure 4.22 shows a picture of the can source as modelled in Geant4.

The optimal design of the ⁶⁰Co source is that in which, by definition, the electron emitted by the ⁶⁰Co decay is completely contained (ideally within the plastic scintillator button) and where the gamma-rays emitted pass through the source into the main SNO+ volume altogether unhindered. This choice is not arbitrary, but is based on the ability to use the source to study in detail the SNO+ energy resolution function and any non-Gaussian tails it may have. To do this, a strong total absorption peak in SNO+ of the ⁶⁰Co gamma-rays is desirable, as this leads to less low-energy tailing on the absorption peak, and complete containment of the electron is necessary to ensure it does not escape into the SNO+ detector volume and add its energy to that of the escaping gamma-rays⁴⁴, which would otherwise lead to high-energy tails on the

⁴³Throughout these simulations, the light guide's height was set to 0.0 mm, which caused it to be explicitly excluded from the simulated source geometry.

⁴⁴It could also be that the electron produces Bremsstrahlung radiation that escapes the source rather than the electron itself physically escaping.

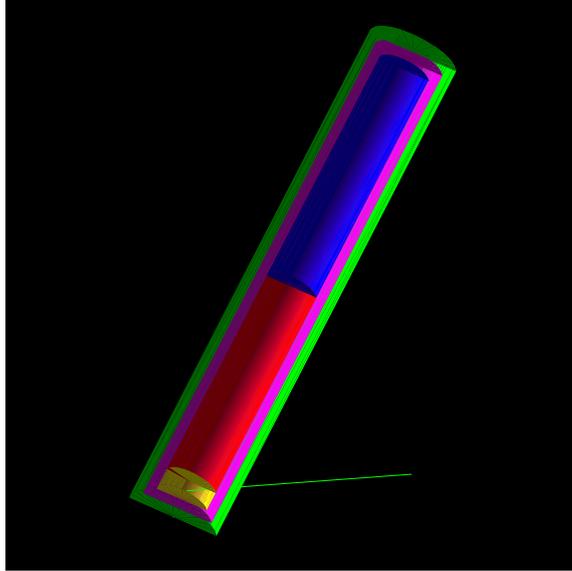


Figure 4.22: This image shows a representation of the can source produced by Geant4. The outer can (green), inner can (magenta), PMT (red), PMT base (blue) and scintillator button (yellow) are simple nested volumes (the light guide is not shown). The scintillator button is not a sensitive volume. The track shown is that of a gamma-ray with its origin at the centre of the button.

detector resolution function. This high-energy tail could, in particular, mask any unknown or higher-order effects inherent in the scintillation process [110] that may lead to a similar high-energy tail, making these effects difficult to study unless the electron is fully contained.

I used the total energy deposited by electrons and gamma-rays in the source volumes to study the optimal ^{60}Co source geometry, where I deemed the geometry that minimized the energy loss of gamma-rays while maximizing the number of electrons that were fully contained to be optimal. Switching the materials comprising the inner and outer source cans between `G4.POLYOXYMETHYLENE` and `stainless_steel`, I varied the cans' wall thicknesses from 1.0–50.0 mm and the thickness of the plastic scintillator button⁴⁵ from 1.0–24.0 mm to

⁴⁵I never altered the PMT geometry. As a result, the radius of the plastic scintillator button remained constant, set to match the active area on PMT face.

investigate how these changes affected the energy depositions.

Rather than simulate the complete decay of ^{60}Co for each event, I simulated the electron, and 1.17 MeV and 1.33 MeV gamma-rays independently from the centre of the plastic scintillator button. Figures 4.23 and 4.24 show examples of the resulting deposited and escape energy distributions for the electron and 1.33 MeV gamma-rays, respectively, in each can source volume. The electron deposits most of its energy in the plastic scintillator, and very few electrons escape the source ($< 0.1\%$). More than 90% of the gamma-rays escape the source unhindered, while the rest deposit some energy in the source volumes, notably the PMT, with an energy deposition spectrum showing a characteristic Compton scattering shape.

To determine the optimal geometry, I compare the average energy of those electrons that escape the source to the sum of the average deposited energies of the 1.17 MeV and 1.33 MeV gamma-rays in the source. Figure 4.25 shows a plot of these energies for each source material and geometry I considered. Ideal geometries lie in the lower left-hand side of the plot. Geometries utilizing only Delrin[®] for the source containers perform the best. The addition of stainless steel quickly increases the average electron escape energy, the result of increased Bremsstrahlung radiation production in the stainless steel. The magenta circle in figure 4.25 indicates the geometry I chose as best. While it is not strictly the optimal geometry, those with lower electron escape energies have either very thin (1 mm) container walls or have a very thick plastic scintillator button, both of which are impractical to construct. The optimal geometry has 2 mm thick walls for each Delrin[®] container and a 4 mm thick plastic scintillator button.

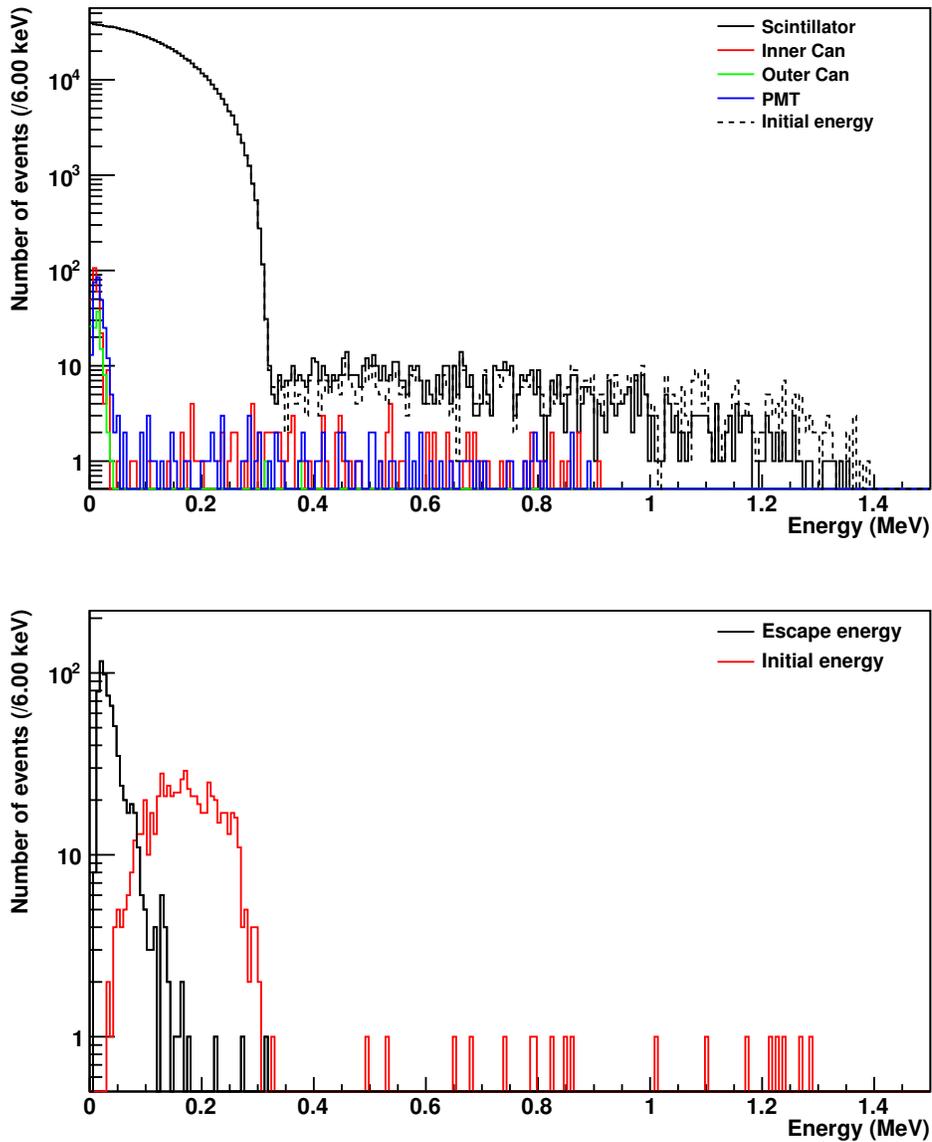


Figure 4.23: Electrons generated at the centre of the plastic scintillator button deposit nearly all of their energy in the can source volumes (top), mainly in the scintillator itself. Relatively few electrons escape the geometry (bottom), and those that do escape typically have higher initial energies. The geometry simulated consisted of two 2 mm thick Delrin[®] containers and a 4 mm thick plastic scintillator button.

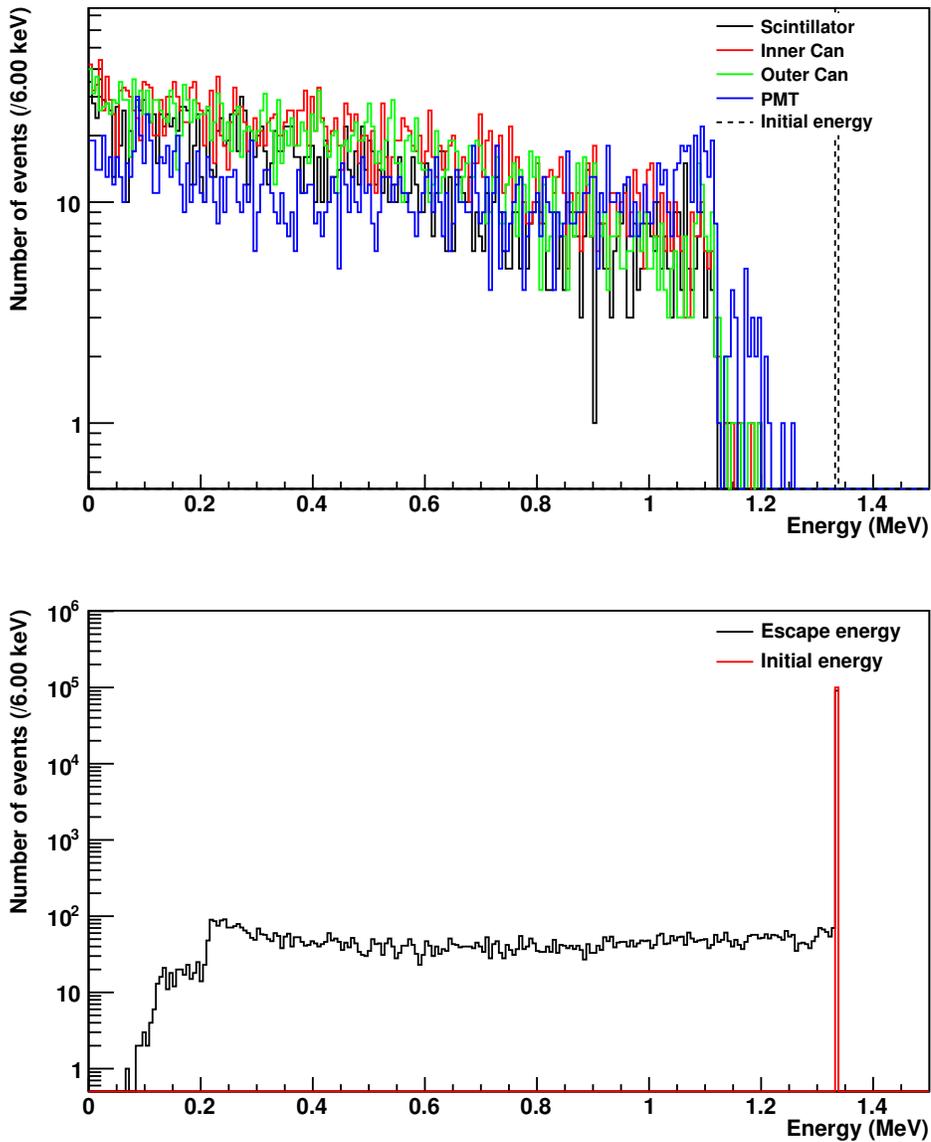


Figure 4.24: Some 1.33 MeV gamma-rays produced at the centre of the plastic scintillator button deposit energy in the can source volumes (top). The escape energy spectrum of the gamma-rays (bottom) has a reverse Compton scatter structure, corresponding to gamma-rays that deposit some energy in the source. The geometry simulated consisted of two 2 mm thick Delrin[®] containers and a 4 mm thick plastic scintillator button.

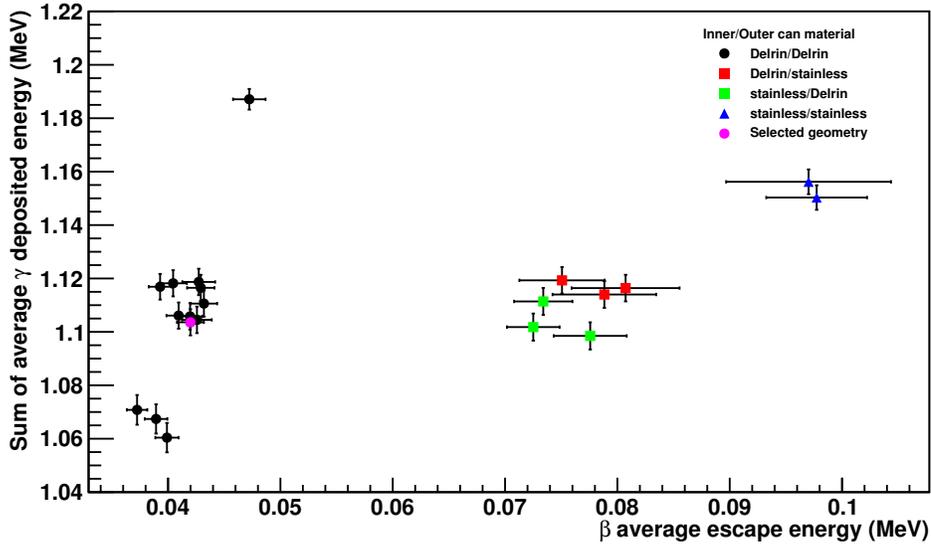


Figure 4.25: Plotting the average electron escape energy versus the sum of the average deposited energies of the 1.17 MeV and 1.33 MeV gamma-rays in the can source places preferred geometries near the lower end of both axes. Error bars are the root mean square of the respective energy distributions. The geometry I ultimately selected is shown as a magenta circle.

4.2.4 The ^{60}Co source generator

A top-level RAT event generator, `co60source` simulates ^{60}Co calibration source events in SNO+. This top-level generator creates ^{60}Co decays in the centre of plastic scintillator button of the calibration source geometry for any position of the source within SNO+.

To run this generator (an example RAT macro is found in appendix B.3.1), the ^{60}Co source geometry is first loaded in a RAT macro via

```
/rat/db/load geo/Co60Source.geo
```

Next, the user sets any database values of interest (a complete list of these is given in appendix B.3.2). Most of the database values are related to the

source dimensions and should not require user changes, particularly as doing so may lead to failures when Geant4 builds the geometry. There are, however, database values the user may wish to alter. These are:

```
/rat/db/set GEO[Co60Source] mother "inner_av"
```

Change the mother volume in which the ^{60}Co source is placed. For example, it may be of interest to place the source in the light water shield region.

```
/rat/db/set GEO[Co60Source] sample_position [0., 0., 0.]
```

Change the position of the source geometry relative to the `mother` volume. The geometry is constructed in such a way that this position is the centre of the plastic scintillator button.

```
/rat/db/set GEO[Co60Source] ref_date "01 Jan 2014 12:00:00"
```

This is the date on which the activity of the source was calibrated. The rate of the source on the run date is calculated based on the time elapsed from this reference date, along with the activity of the source on the reference date.

```
/rat/db/set GEO[Co60Source] ref_activity 200.
```

This is the calibrated source activity on the reference date `ref_date`. The rate of the source on the run date is calculated based on the elapsed time and this reference activity. If the `ref_date` field is set to an empty string, then the activity of the source during the run is taken to be the reference activity specified here.

```
/rat/db/set GEO[Co60Source] check_overlaps 1
```

This boolean field, if true, forces Geant4 to check for overlapping volumes

when constructing the source. It is set to 1 by default.

```
/rat/db/set GEO[Co60Source] screws_enable 1
```

This boolean field, if true, causes Geant4 to build the screws, nuts and screw holes in the ^{60}Co source geometry. Not all users will necessarily require this level of detail, though the value is set to 1 by default.

```
/rat/db/set DATE date_day 1722
```

Set the current date for the MC run. This sets the day as the number of days since 00:00:00 on January 1, 2010.

```
/rat/db/set DATE date_sec 2100
```

Set the current date for the MC run. This sets the number of seconds since 00:00:00 on `date_day`.

```
/rat/db/set DATE date_nsec 73002.
```

Set the current date for the MC run. This sets the number of nanoseconds since the start of second `date_sec`.

The generator is run using

```
/generator/add co60source
```

The generator state is completely specified internally, so users need not specify any state (vertex, position or time) in the RAT macro. The ^{60}Co decays are simulated using the `decay0` vertex generator (there is no option to change this). The position generator is set to `point`, with the generation point equal to `sample_position`, corresponding to the centre of the plastic scintillator button. The `poisson` time generator handles the event timing, with the rate

R calculated internally as

$$R = R_o e^{-\ln(2)\Delta t/\tau_{1/2}} , \quad (4.6)$$

where R_o is the reference activity specified as `ref_activity`, $\tau_{1/2}$ is the ^{60}Co half-life (read from `Decay0Backg.ratdb`), and Δt is the difference in time from the run date specified by `DATE` and the source reference date `ref_date` (the time subtraction is handled internally via the `EventTime` class). To specify a particular rate rather than have the generator calculate it based on the time elapsed from the reference date, users may simply set the reference date to be an empty string. In this case, the rate users specify in `ref_activity` is the simulated rate.

RAT may be run in either of two modes to generate MC events. The first is the static mode, where users specify a number of events to simulate, `NEVENTS`, through

```
/rat/run/start NEVENTS
```

This mode ignores the reference activity and generates exactly the number of events requested. The second mode, or dynamic mode, allows users to set the duration of a run through

```
/rat/run/duration TIME UNIT
```

```
/rat/run/start
```

which generates ^{60}Co events until the MC time of a generated event surpasses the run duration `TIME` (in units `UNIT`). This mode is well defined for the ^{60}Co source generator because the generator always has a defined rate. Note that,

since the ^{60}Co generator uses the `poisson` time generator, the time between events is not constant and the total number of events simulated over the run duration may change each time the simulation is run (with the average number of events given by the duration multiplied by the calculated (or specified) activity).

Validation of the ^{60}Co nuclear decay

A critical aspect of the ^{60}Co calibration source simulation is the simulation of the decay of ^{60}Co itself. The decay of ^{60}Co is somewhat straightforward⁴⁶, decaying almost always (99.88% of the time) to the 2.51 MeV excited state of ^{60}Ni , emitting a 0.318 MeV endpoint energy electron followed by 1.17 MeV and 1.33 MeV gamma-rays (or a combination of other gamma-rays, with much smaller probability). There is also a small chance (0.12%) it will decay to a lower excited state, emitting an electron with endpoint 1.491 MeV and only a single 1.33 MeV gamma-ray. The transition to this lower excited state requires a change in spin of 3 units with no change in parity, implying that this branch is a second-order forbidden (unique) decay, whereas the main branch is an allowed decay. Table 4.5 lists the electron endpoint and gamma-ray energies for each branch in the ^{60}Co decay, as well as the spin transitions for the gamma-ray cascades (the parity of each ^{60}Ni energy level is positive).

The attributes of the ^{60}Co decay that require particular attention are the decay half-life, the number of branches through which the decay may proceed along with the branching ratio and forbiddenness of each, the number and energies of the gamma-rays and the energy spectrum of the electron emitted in each branch, and the spin/parity of each energy level through which a particu-

⁴⁶This is one reason why it is a good candidate for a calibration source.

Branching Ratio	Electron endpoint (MeV)	Gamma-ray energies (MeV)			Spin-parity transition
0.9985	0.3182	1.1732	1.3325		5→4→2→0
0.0012	1.4914	1.3325			5→2→0
2.8798×10^{-4}	0.3182	0.3471	0.8261	1.3325	5→4→2→2→0
1.2×10^{-5}	0.3182	0.3471	2.1568		5→4→2→0
2×10^{-8}	0.3182	2.5057			5→0

Table 4.5: ^{60}Co decays predominantly to two gamma-rays of 1.1732 MeV and 1.3325 MeV, and an electron with endpoint energy 0.3182 MeV. Several other branches are accessible and contribute to the overall ^{60}Co decay scheme in a minor way. This table lists for each decay branch the energy of each gamma-ray emitted, along with the endpoint of the spectrum of electron energies and the branching ratio [64]. The branches are listed in order of precedence, the main branch listed first, the second branch following, &c.

lar branch proceeds (for consideration of correlations between the directions of gamma-rays from consecutive energy levels in a cascade decay). RAT uses two software routines to simulate radioactive decays⁴⁷, `decay0` and `decaychain`, both of which take these properties into account, albeit in different ways.

`decay0` is a port into RAT of the Fortran program [111, 112] of the same name⁴⁸. This software simulates both $2\nu\beta\beta$ and $0\nu\beta\beta$ processes for the known $2\nu\beta\beta$ isotopes, as well as beta and alpha decays of various other isotopes. Each of these latter isotopes has its own routine to handle its specific decay, where spectral shapes of emitted electrons are computed semi-empirically (for the non- $2\nu\beta\beta$ isotopes) using experimental spectral shape correction factors. `decay0` also computes the correlated opening angle between cascade gamma-rays. While this makes `decay0` a powerful tool (indeed, it is used by other $2\nu\beta\beta$ experiments other than SNO+), it is challenging to add additional iso-

⁴⁷While Geant4 has the ability to simulate radioactive decays, this feature is not explicitly implemented in RAT.

⁴⁸The original Fortran code was painstakingly translated, with approval of the author, into C++ in order to smoothly integrate into RAT. Details are contained in [113].

topes of interest into the software, as this requires an acute knowledge of both the software and the data comprising the decay. As such, `decay0`'s use is practically limited to the isotopes included by the author, which are few.

`decaychain` is a routine included in RAT⁴⁹ that simulates the beta and alpha decays of various isotopes (it does not simulate $2\nu\beta\beta$ or $0\nu\beta\beta$ decays). It calculates electron spectra for allowed and unique forbidden decays based on the method in [114] (it contains the functionality to use the same experimental shape correction factors as `decay0`, but these are never used, in practice). Because of this, it is a simple matter for users to add additional isotopes into the list of isotopes `decaychain` is able to simulate. Based on modifications I made to the routine, it now handles β^- , β^+ , electron capture, isomeric transition and alpha decays, as well as takes into account Auger and conversion electrons. A nice feature of `decaychain` is that it enables the simulation of decay chains (as its name would suggest), where a sequence of decays follows from the decay of some mother isotope (like the ^{238}U decay chain ending in stable ^{206}Pb). I have also modified the routine to make use of this functionality to simulate the decays of metastable states. `decaychain` does not have a mechanism to compute the correlated opening angle between cascade gamma-rays, as this is isotope specific and `decaychain` treats all isotopes in a general way.

At the onset of my work with the ^{60}Co calibration source, only `decaychain` existed in RAT. It was clear that it was not accurately simulating electron energy spectra, particularly at low energies. The electron energy spectrum, as a function of energy W and atomic number Z of the daughter nucleus, is

⁴⁹The author of the original routine seems to be J. Formaggio, though it is not clear when he wrote it or how it made its way into RAT.

given by

$$N(Z, W) = pW(W_0 - W)^2 F(Z, W) R(Z, W) Q(Z, W) , \quad (4.7)$$

where N is the emission probability, p and W are the electron momentum and total energy (in units of the electron mass m_e), W_0 is the maximum electron energy (the endpoint), F is the Fermi function, R is a nuclear size correction factor, and Q is a screening correction⁵⁰ that considers the effect of atomic electrons on the escaping electron. We may write the screening correction factor Q as

$$Q(Z, W) = \frac{p'W'F(Z, W')}{pWF(Z, W)} , \quad (4.8)$$

where $W' = W - V_0$ is the energy of the electron above the effective potential introduced by the atomic electrons (which is model-dependent). Details are in [114] and [115].

As W' approaches V_0 , the screening correction in equation 4.8 causes inaccuracies when computing the electron energy spectrum using equation 4.7. To overcome this difficulty, I introduced a minimum screening correction $Q_{\min}(Z)$ to prevent $Q(Z, W)$ from becoming too low. Figure 4.26 shows the values of the ratio of the screened to unscreened Fermi function ($F(Z, W')/F(Z, W)$) from table III of [116] as a function of Z . Fitting that data with a straight line (a simplified model) gives

$$Q_{\min}(Z) = 0.9958 - 0.0004439Z . \quad (4.9)$$

⁵⁰The experimental shape correction factor mentioned above takes the place of Q in the semi-empirical models.

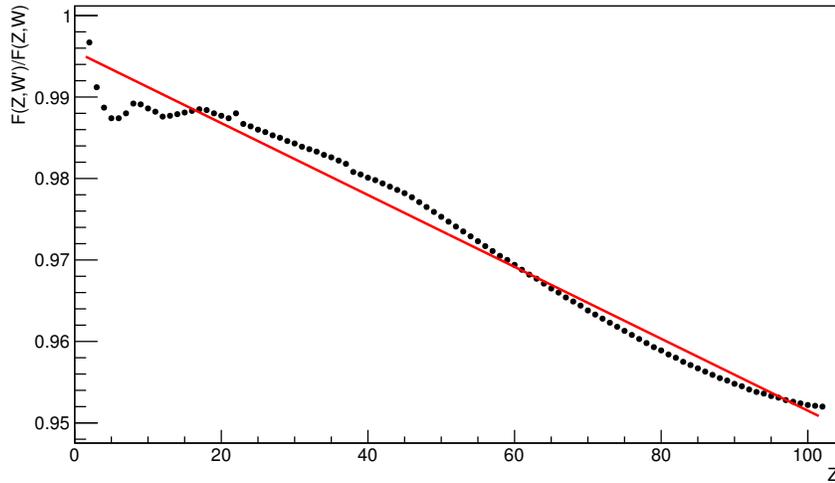


Figure 4.26: Fitting a straight line to the minimum value of the ratio of the screened to unscreened Fermi function as a function of Z provides a lower limit to the size of the screening correction, $Q(Z, W)$.

To use this screening correction, I altered the computation of the Fermi function to use the value of the electron wavefunction at its unscreened energy W , while the screened value W' is still used in the computation of the screening correction above Q_{\min} (details are in [114]).

Figure 4.27 shows the electron spectral shapes produced by `decaychain` and `decay0` for the first two branches table 4.5 lists. The `decaychain` spectra both before and after the corrections I implemented are shown. The figure also provides calculated beta spectra [117] from the Joint Institute for Nuclear Research (JINR)⁵¹ and, for the main ^{60}Co decay branch, a comparison to a calculation of the electron spectrum [118] supplied by I. Towner⁵², which we take to be the best spectrum as its source is well-known. As figure 4.27 shows, both

⁵¹These spectra are calculated semi-empirically, though the details of the calculation are not clear. The data appear to no longer be available, either.

⁵²I. Towner provided us several calculations of this spectrum upon request. The one shown uses a large Gamow-Teller matrix element, which Towner suggested was likely the most accurate. All calculations included a nuclear size and a screening correction factor.

the corrected `decaychain` simulation and the `decay0` simulation agree with the Towner calculation of the main branch electron energy spectrum. Nevertheless, the two simulations do not completely agree on the spectral shape for the second branch. Because `decay0` is generally regarded as the more accurate simulation⁵³, I take it to have produced the more accurate spectral shape. Also, note how the uncorrected `decaychain` simulation underestimated the emission of low-energy electrons, producing unphysical spectral features for energies below ~ 50 keV. The corrections I made to the `decaychain` simulation solved this error, causing the simulation to match with high precision the Towner calculation for the main ^{60}Co decay branch. Also, the spectra calculated by JINR agree with those calculated by `decay0` above ~ 50 keV, but appears to have difficulty producing an accurate spectrum below that energy.

Lastly, figure 4.28 shows the opening angle between the two gamma-rays emitted in the main ^{60}Co decay branch using `decay0`. The data are shown with the theoretical curve for a $4+\rightarrow 2+\rightarrow 0+$ cascade, given by [120]

$$N(\theta) = 1 + \frac{1}{8} \cos^2 \theta + \frac{1}{24} \cos^4 \theta , \quad (4.10)$$

overlaid. Clearly, `decay0` correctly performs this angular sampling. `decaychain` does not simulate this effect, so its corresponding distribution is flat in $\cos \theta$.

Because `decay0` correctly simulates the electron energy spectra for all ^{60}Co decay branches and includes all appropriate gamma-rays and the angular correlations between gamma-rays emitted in a cascade, I conclude it is the preferred simulation for ^{60}Co decays and use it for the RAT ^{60}Co source generator.

⁵³SNO+ examined the electron spectrum shapes for a host of isotopes [119], and in most cases `decay0` outperformed `decaychain`. The conclusion is that we use `decay0` when possible, and only use `decaychain` when the isotope of interest is not included in the `decay0` arsenal.

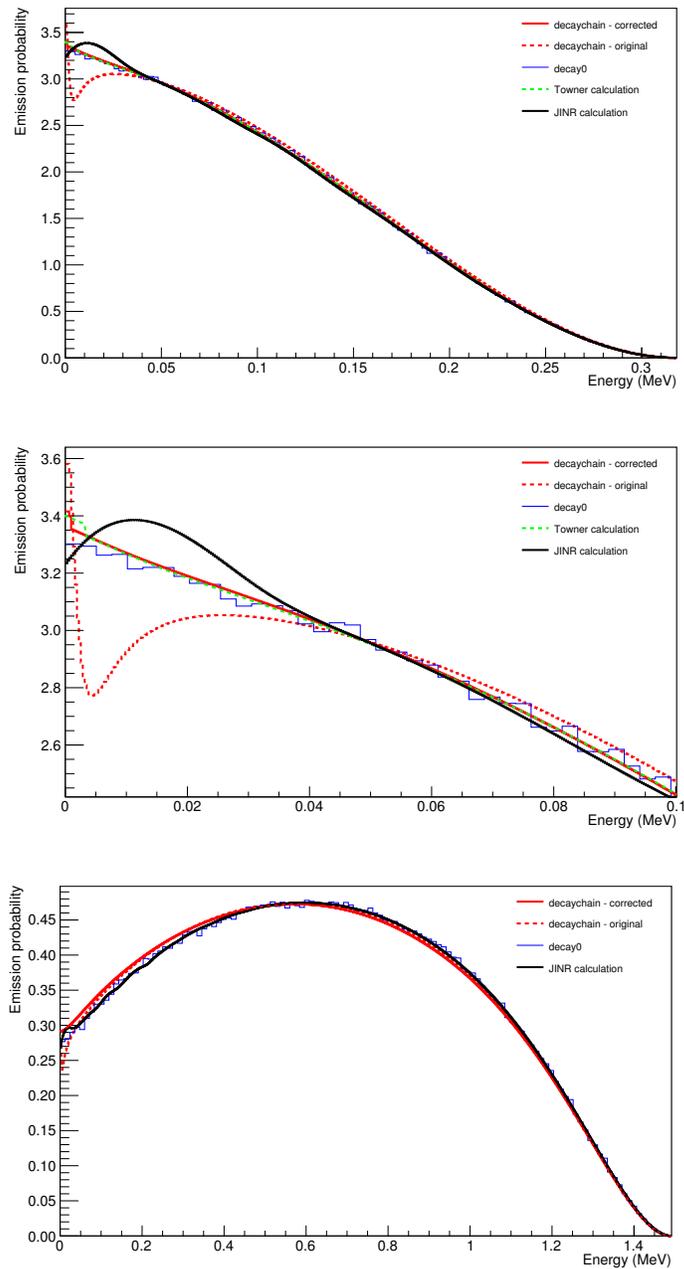


Figure 4.27: The main (top, middle) and second (bottom) ^{60}Co decay branches are shown for both the original and corrected `decaychain` simulations, the `decay0` simulation, JINR calculations and Towner calculation (main branch). The original `decaychain` simulation (red) poorly calculated the spectra at low energies. For the main branch, both the corrected `decaychain` (red dashed) and `decay0` (blue) simulations are well-matched to the Towner calculated spectrum (green dashed). For the second branch, neither `decaychain` simulation are a great match to `decay0`. The JINR calculations (black) mostly agree with `decay0`, deviating below ~ 50 keV.

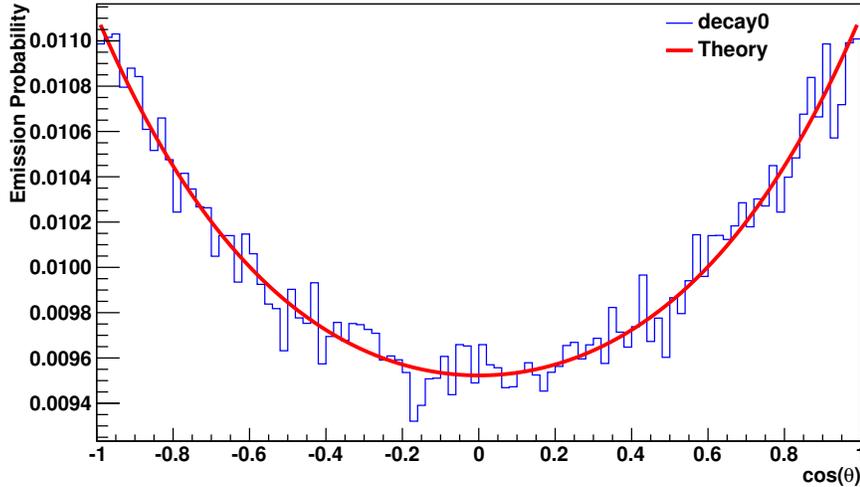


Figure 4.28: The opening angle between the gamma-rays emitted in the main ^{60}Co decay branch in simulations using `decay0` follow that predicted by equation 4.10. `decaychain`, on the other hand, does not take this correlation into account, which results in a flat distribution of $\cos\theta$.

4.2.5 Simulation results

To investigate the effect of using the ^{60}Co calibration source in the SNO+ detector, I simulated calibration events using the ^{60}Co source generator and compared the result to a simulation of two simultaneous gamma-rays of energies 1.3325 MeV and 1.1732 MeV with no source geometry present. The latter simulation I take as equivalent to having the ^{60}Co source in the detector, while removing any effect the source geometry has on the propagation of light through the LS, which we call shadowing.

Figure 4.29 shows the result of simulating ^{60}Co source events and the equivalent double gamma-ray events at the centre of the LS volume (and using the nominal SNO+ LS). The resultant full-deposition peak from the ^{60}Co source causes 2.64% fewer PMTs to observe at least one photoelectron (nHits is the

number of such PMTs) per event than when the source geometry is not present. This implies an energy scale of 1254.3 nHit at 2.5057 MeV. The large low-nHit tailing in the ^{60}Co source distribution is due to interactions of one of the gamma-rays with the source geometry (a Compton scatter) before entering the LS volume. This tends to broaden the full-deposition peak. To determine the mean and resolution of that peak, I fit it with a Gaussian function from $-0.25\sigma_0$ to $1.5\sigma_0$, where I take $\sigma_0 = \sqrt{\text{nHit}}$ as an initial estimate of the resolution. The nHit resolution I obtain from this fit is 3.24% at 1222 nHit. This is slightly broader than expected from a Poisson distribution. Figure 4.29 also shows the nHit distribution of ^{60}Co decays at the centre of the detector. The distribution is broadened toward higher nHit as the electron adds its energy to the event. The second ^{60}Co decay branch is visible as the few events scattered near 800 nHit.

Another study I performed was to look at the effect of particle type on the response of the LS. During the $2\nu\beta\beta$ phase, the energy contribution from two electrons will be producing scintillation light in the TeLS, whereas the ^{60}Co calibration source will use two gamma-rays to produce scintillation light. While both of these particle types eventually create scintillation light via electrons, the ways in which they interact with the LS are somewhat different (the gamma-rays typically Compton scatter to first produce electrons that then create scintillation light, for example). Because the production of scintillation light is more quenched (see section 5.1) for low-energy electrons, events in which multiple low-energy electrons produce the scintillation light will look different than events in which a single high-energy electron of equivalent energy produces the scintillation light.

Figure 4.30 shows this effect for single- and double-electron and gamma-

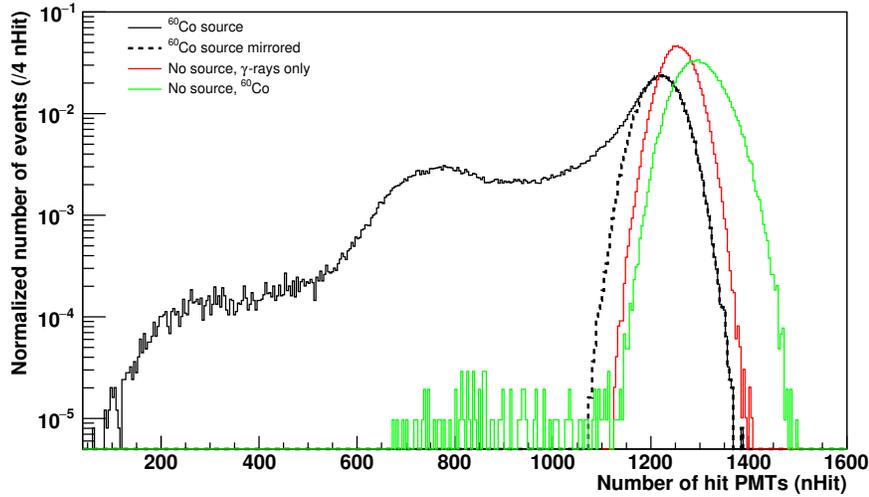


Figure 4.29: Simulating ^{60}Co source (black) and equivalent double gamma-ray (red) events from the centre of the LS volume shows the source geometry shadows the propagation of light in the detector, leading to a 2.64% downward shift in the full-deposition peak position. The geometry also causes a low-nHit tail on the ^{60}Co events as some gamma-rays interact in the source geometry before entering the LS volume. To highlight this low-nHit broadening, a mirrored ^{60}Co source distribution (dashed black) is drawn where the distribution above the mean is reflected about the mean. For comparison, ^{60}Co events at the centre of the detector (green) tend toward higher nHit as the electron contributes additional light to the event.

ray events, where the total energy deposited in the LS volume is always 2.5057 MeV. Gamma-rays produce less light than equivalent-energy electrons, as the gamma-rays tend to undergo multiple Compton scatters (producing several low-energy electrons in the process) in order to deposit their total energy in the LS. In the same way, the single-particle events tend to produce more light in the LS than the double-particle events. Relative to the double-gamma-ray events, which produce the least light in the detector of the events simulated, the double-electron events produce, on average, 3.38% more nHits, while the single-gamma-ray and electron events produce, on average, 2.23% and 4.26% more nHits, respectively.

It is imperative, then, that we correctly understand how quenching of the LS affects the data we collect from various types of events. This is particularly true as the calibration sources SNO+ plans to employ provide a calibration for gamma-rays in the TeLS, whereas the $0\nu\beta\beta$ signal is the summed energy deposition of two electrons in the TeLS. Understanding how to set the energy scale is critical for defining the SNO+ ROI for $0\nu\beta\beta$ events. Complicating this further is the fact that radioactive backgrounds in the SNO+ detector typically produce a combination of gamma-rays and electrons, or alphas (which are even more highly quenched), which makes placing all events on the same energy scale a non-trivial task.

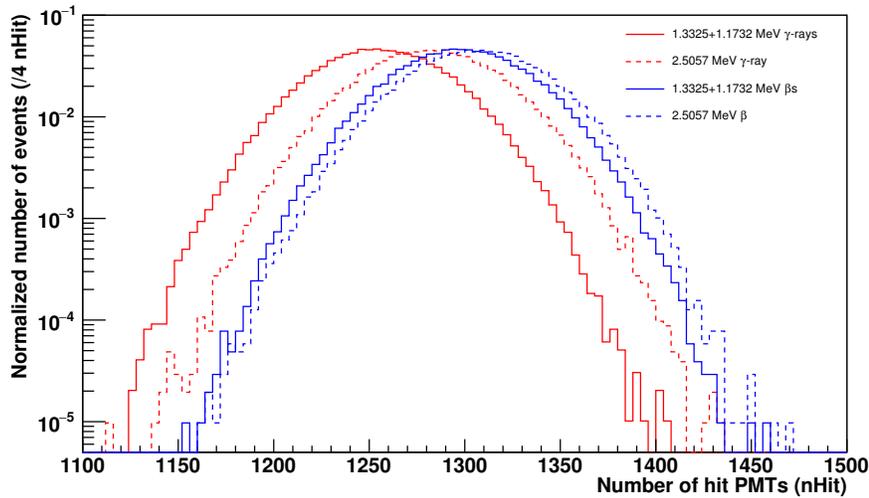


Figure 4.30: The response of the LS depends not only on the energy of the particles interacting with it, but also on the type of particle. The amount of light produced in the LS (parametrized here as nHits) is higher for electrons (blue) than gamma-rays (red) with equal energy. Also, the amount of light a single particle (dashed) produces is higher than the amount two particles (solid) with an equivalent sum energy produce, a result of quenching of the LS for multiple lower-energy particles. Gamma-rays appear more heavily quenched than electrons.

Chapter 5

Quenching of low energy electrons

I do not know what I may appear to the world, but to myself I seem to have been only like a boy playing on the sea shore, and diverting myself in now and then finding a smoother pebble or a prettier shell than ordinary, while the great ocean of truth lay all undiscovered before me.

- Isaac Newton, *From the Anecdotes of Joseph Spence*

The simulation results of the ^{60}Co calibration source in section 4.2 clearly show a decrease in the LS light output (the number of scintillation photons produced) for multiple lower-energy electrons together depositing energy E in the LS, versus a single electron depositing E . The key difference is that the stopping power (the change in E along a particle's track) for a low-energy electron is significantly higher than that for a high-energy electron, which means multiple low-energy electrons tend to produce more regions of high ionization in the LS. Ionization quenching—this reduction in light output—occurs in organic liquid scintillators when a particle with high stopping power interacts,

causing non-linear changes, characterized by Birks' constant, to the energy-dependent LS light output. Section 5.1 gives an overview of the scintillation process and quenching effect in organic liquid scintillators.

Quenching is strongly pronounced for particles with high stopping power, like alphas, protons and low-energy electrons. This is significant for SNO+, as many of the radioactive backgrounds the detector will observe produce alphas and the $0\nu\beta\beta$ signal is the sum energy of two electrons interacting in the TeLS, where one electron may have substantially lower energy than the other. Because we expect the light output of the LS and TeLS to drop in these cases, it is critical we understand *how* it drops in order to correctly set the energy scale for these events. This is particularly important for events with energy in the SNO+ ROI near the ^{130}Te $2\nu\beta\beta$ endpoint, where how well we determine a particle's energy may have a significant effect on the experiment's sensitivity.

This chapter presents the first measurement of low-energy electron quenching in TeLS. Section 5.2 provides details regarding the technique I used and apparatus we built to make this measurement. How I processed the raw data coming out of the apparatus is the subject of section 5.3, while I describe the analysis techniques I used to examine the quenching effect in section 5.4.

As a final note, similar measurements have been made for low-energy electron quenching in LS and neodymium-loaded LS [121], as well as for alphas and protons in TeLS [122, 67, 123], both as exercises for the SNO+ experiment. This measurement of low-energy electron quenching in TeLS complements these other measurements, while still being unique and significant on its own. Section 5.5 collates and compares the results of these various measurements, and compares them to the result of the low-energy electron quenching measurement in TeLS.

5.1 Scintillation and quenching

Aromatic and conjugated organic molecules, in particular polycyclic hydrocarbon systems of benzenoid rings (like LAB), have the particular property that a portion of the outer electron orbitals of the carbon atoms are shared in such a way that these electrons exist in free orbitals about the entire perimeter of the molecule [124]. As such, the entire molecule exhibits a series of excited states, comprised of electronic, vibrational and rotational levels, with spacings on the order of a few eV for the electronic levels, a fraction of an eV for the vibrational levels and near negligible spacing for the rotational levels [125].

This complex band structure is what enables scintillation by these molecules. Interaction with ionizing radiation may push the molecule into an excited state. From this excited state, the molecule may relax back to its ground state, S_0 , through either a radiative or non-radiative release of energy. Vibrational states typically interact with surrounding molecules until the vibrational energy is completely released in the form of heat. Electronic states may internally transfer energy to a vibrational state, which then dissipates as heat, or may relax via the emission of a photon. Because some energy is typically lost via vibrational dissipation, any emitted photon is typically lower in energy than that initially absorbed. This implies these molecules' absorption and emission spectra do not coincide, so that self-absorption of the emitted light by the medium does not readily occur.

A typical LS consists of a solvent and a small fraction of one or more solutes. In such a system, the absorption spectrum of the solute typically overlaps the emission spectrum of the solvent. Energy transfer in these systems proceeds as follows [124]. Ionizing radiation excites a molecule of the solvent, which relaxes

via non-radiative energy transfer to a vibrational energy level associated with its lowest electronic excited state, S_1 . The vibrational energy is dissipated non-radiatively before the energy associated with S_1 is efficiently transferred, typically non-radiatively, to a molecule of the solute. This process is more efficient than either the radiative or non-radiative de-excitation of the solvent molecule to S_0 . The solute then relaxes to its S_1 state through vibrational energy transfer to surrounding molecules, then relaxes radiatively from S_1 to its S_0 state. Because the energy levels of the solute are lower than those of the solvent, energy cannot be transferred back to the solvent molecules. The emission from the solute is then away from the absorption of the solvent, so the liquid becomes transparent to the scintillation light it emits. Complex systems of solutes may be designed to increase the efficiency of this process and to move the emission spectrum of the liquid away from its absorption spectrum (so to improve the liquid's transparency) and toward the absorption spectrum of the instrument (say, a PMT) being used to observe the scintillation photons. Typically, we refer to the solvent as the *scintillator*, the solute as a *fluor* and any secondary solute as a *wavelength shifter*.

We may further divide the excited electronic states of the molecule into spin-singlet and spin-triplet states, where the energy of the triplet states is necessarily somewhat below that of the corresponding singlet state. While transitions from S_0 into the first triplet excited state T_1 are forbidden by spin-selection criteria, it is possible for a molecule to undergo non-radiative de-excitation from an excited singlet state into T_1 in a process we call inter-system crossing. Once in T_1 , the molecule will typically remain there until it interacts with another molecule also in T_1 , which produces one molecule in S_0 and another in S_1 , which then de-excites radiatively (for the solute). This leads to

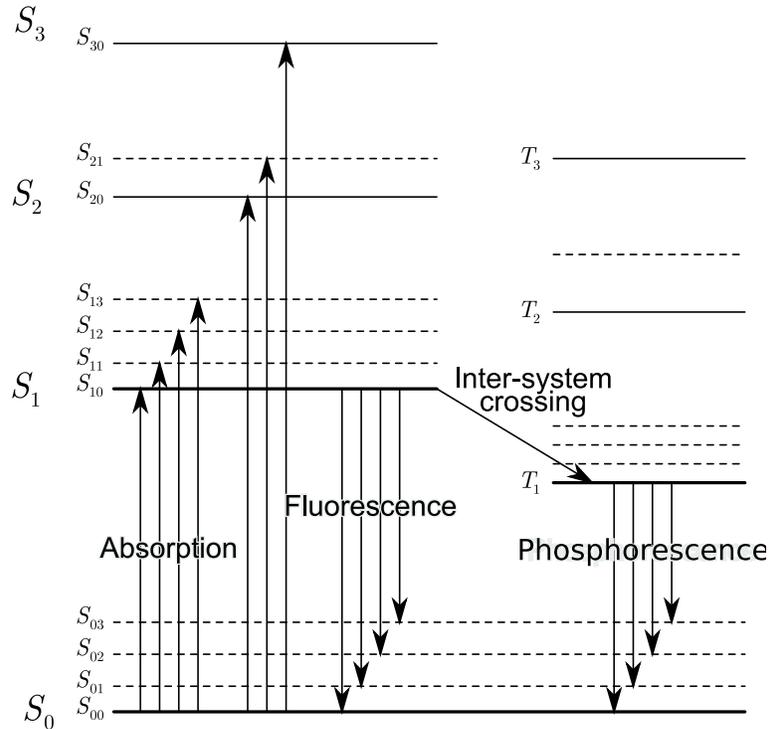


Figure 5.1: A molecule may be excited from its ground state S_0 into an excited singlet state S_{ij} , where i corresponds to the electronic level and j the vibrational level. Non-radiative de-excitation to the first excited singlet state S_1 or to the lowest excited triplet state T_1 then occurs before the molecule relaxes radiatively to S_0 via fluorescence (fast scintillation) or phosphorescence (slow scintillation) [126].

a delayed scintillation component and is the reason why the scintillation timing profile for incident alphas is longer than the profile for other incident radiation (see section 2.1.2). We typically refer to the fast scintillation component as *fluorescence* and this slow component as *phosphorescence*. Figure 5.1 provides a simplified schematic of the molecular energy levels.

Particles with low stopping power dE/dx (say electrons with energies above ~ 1 MeV) produce relatively sparse ionization (and subsequent molecular excitation) along their trajectory through a typical LS. Because the separation

between these excitations is on the order of a few molecular distances, interactions between them are rare, so the scintillation process I described above happens basically unhindered [124]. As such, the light output L of the LS is proportional to the total energy E a particle deposits in the LS, where

$$\begin{aligned} L &= L_0 E \\ \frac{dL}{dx} &= L_0 \frac{dE}{dx} . \end{aligned} \tag{5.1}$$

L_0 is the absolute scintillation efficiency, which we commonly call light yield.

Particles with large dE/dx , like alphas or low-energy electrons (<125 keV), have dense ionization along their tracks, where BdE/dx is the number of ionizations created by a particle depositing an energy dE in a track of length dx , with B a constant. In this case, the primary molecular excitation efficiency may be reduced, which is known as ionization quenching [127, 124, 128]. Assuming this quenching is limited to a single molecule (so interactions between molecules are not considered), equation 5.1 becomes

$$\frac{dL}{dx} = \frac{L_0 dE/dx}{1 + kBdE/dx} , \tag{5.2}$$

where k is the quenching parameter [127]. In principle, the parameter we measure that describes quenching is kB , as it is difficult to decouple the two, though they have different physical interpretations. Note that when dE/dx is small, this reduces to equation 5.1. For very large dE/dx , the differential light output saturates at $dL/dx = L_0/kB$, becoming constant. We typically refer to the parameter kB as *Birks' constant*, and treat it as a measure of the amount of quenching in a medium (equation 5.2 applies to non-liquids, as well).

There are other models in addition to that in equation 5.2 that describe ionization quenching. A bimolecular model that assumes interactions between nearby similar molecules [124] suffers from an inability to correctly predict the saturation effect for high dE/dx . We may also generalize equation 5.2 to include second order effects of the ionization density, where

$$\frac{dL}{dx} = \frac{L_0 dE/dx}{1 + B' dE/dx + C' (dE/dx)^2} , \quad (5.3)$$

with B' and C' constants. Typically, though, data are consistent with $C' = 0$. Birks' model in equation 5.2 is generally regarded as the standard description for quenching, even though the exact physical process is not completely understood [128].

Other types of quenching are also possible, with the common feature that all types decrease the light output of the LS. Intramolecular quenching occurs when excitation energy is dissipated vibrationally, or through other non-radiative mechanisms. If the concentration of solute in the LS is too high, it may lead to concentration quenching, where interactions between solute molecules tend to lead to non-radiative energy losses (and which is more prevalent in liquids where the molecular mobility is high). Chemical quenching is the product of impurities in the LS that can trap excitation energy and which are very inefficient at releasing energy radiatively. Quenching as a result of oxygen contamination in a LS sample is an example of chemical quenching, where oxygen is one of the strongest chemical quenching agents available [125]. Lastly, colour quenching may occur if impurities in the LS have a strong absorption spectrum overlapping the scintillation spectrum, so scintillation photons are absorbed, with that energy being dissipated non-radiatively.

5.2 Apparatus

Ideally, to measure the LS light output as a function of electron energy, we would like to have a source of mono-energetic electrons contained within our LS sample. In this way, we know the exact electron energy inciting the LS response, whereas any apparatus holding the LS sample would invariably affect the energy of electrons originating external to the apparatus. On the other hand, adding any electron source into the LS sample could affect the chemistry of the LS, again altering its response.

One way around this is to *produce* an electron inside the LS using a gamma-ray that originates outside the apparatus with energy E_γ and then Compton scatters inside the LS. The energy $E_{\gamma'}$ of the outgoing gamma-ray constrains the electron energy E_{e^-} via

$$\begin{aligned} E_{e^-} &= E_\gamma - E_{\gamma'} \\ \frac{1}{E_{\gamma'}} - \frac{1}{E_\gamma} &= \frac{1 - \cos \theta}{m_{e^-} c^2} \end{aligned} \quad (5.4)$$

where θ is the angle between the incoming and outgoing gamma-rays¹.

This Compton coincidence technique, whereby $E_{\gamma'}$ is measured coincidentally with the LS light output, is a well-tested technique for performing scintillator response and absolute efficiency measurements [129, 130, 131, 132, 133]. In my setup, a HPGe detects the gamma-ray while the LS is viewed by two PMTs that measure the LS light output. Figure 5.2 shows a schematic of

¹If the electron is too energetic or if the LS volume is too small, then the electron could escape the LS volume, in which case equation 5.4 overestimates the amount of energy in the LS. Also, the gamma-ray could undergo multiple Compton scatters, either in the LS or in its surrounding apparatus, before being detected, so could be measured with a lower energy than would balance the energy of the electron in equation 5.4.

the experiment. While this technique is not novel, the use of two PMTs in conjunction with an HPGe does not appear to be common (though [133] performs a similar experiment using two PMTs). This triple instrument (two detector) setup allows me to measure the production of very few photons in the LS [134] (down to the SPE level for each PMT), and the HPGe provides a high-precision measurement of $E_{\gamma'}$, so a precise measurement of E_{e^-} .

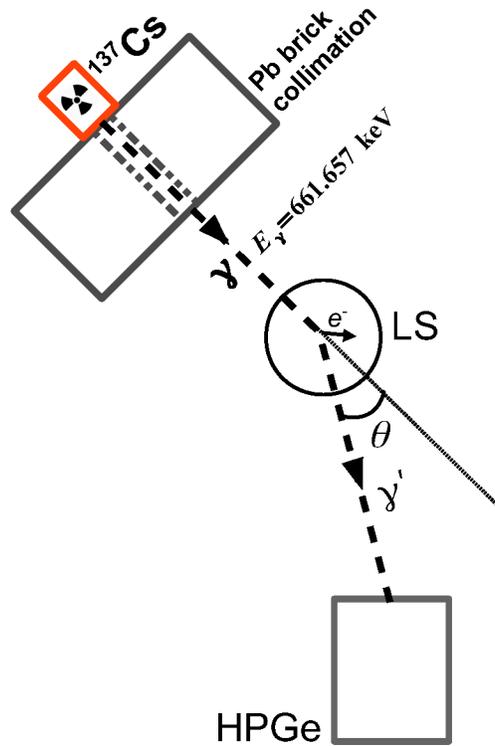


Figure 5.2: 661.657 keV gamma-rays from a ^{137}Cs source Compton scatter inside the LS volume. The scattered electron produces scintillation light in the LS, which two PMTs couples to the LS volume observe. The resultant gamma-ray may subsequently interact in the HPGe (depending on the scattering angle θ), which, given a full-deposition of the gamma-ray energy in the HPGe, completely constrains the electron energy.

Section 5.2.1 describes the apparatus we built to contain the LS samples,

with an overview of the LS processing system I used to degas the LS samples and fill the apparatus contained in section 5.2.2. Section 5.2.3 describes how I assembled the detectors, including both the PMT and HPGe components. Finally, a description of the electronics, trigger logic and data acquisition is in section 5.2.4.

5.2.1 Acrylic chamber

We designed and built an acrylic chamber (AC) to hold the LS sample. Acrylic is chemically compatible with the LS² and is optically clear at the wavelengths the LS emits, enabling readout of the light by a PMT. The AC is a cylinder $3\frac{1}{2}$ in long and $1\frac{1}{2}$ in in diameter. On each end, 1 in of solid acrylic acts as a light guide from the LS volume to a PMT, one optically coupled to each light guide (see section 5.2.3). The $1\frac{1}{2}$ in long central volume with wall thickness $\frac{1}{8}$ in holds approximately 30 mL of LS. One light guide is a plug that slides into an open end of the AC. The plug has a $1\frac{1}{4}$ in diameter for $\frac{3}{4}$ in of its length, with the last $\frac{1}{4}$ in of length having a $1\frac{1}{2}$ in diameter. A Viton[®] o-ring (size 028) sits circumferentially in a groove 0.55 in up from the plug's bottom. A $\frac{7}{64}$ in diameter elbow is bored into the plug $\frac{1}{4}$ in up from its bottom, extending $\frac{3}{16}$ in into the plug before turning 90° and exiting the plug's bottom. The AC body has two ports aligned with these elbow joints, each $\frac{1}{8}$ in in diameter. A 0.29 in o-ring groove is cut on the inside of the AC body about each port to accommodate an additional Viton[®] o-ring (size 006). Turning the plug so the ports and elbows are not aligned isolates the LS volume from atmosphere, where the three o-rings act as seals. To prevent shearing of the port o-rings on the elbow edges as we turn the plug to isolate the LS volume, we lightly

²This is critical for the SNO+ experiment!

ran one side of each o-ring over sand paper until one side was visibly flattened. To prevent shearing the port o-rings on the plug when pushing the plug into the AC, we cut small angled notches into the base of the plug that allow the o-rings to be eased into their compressed states. Figure 5.3 contains the technical drawings of the AC and figure 5.4 shows pictures of the AC as built.

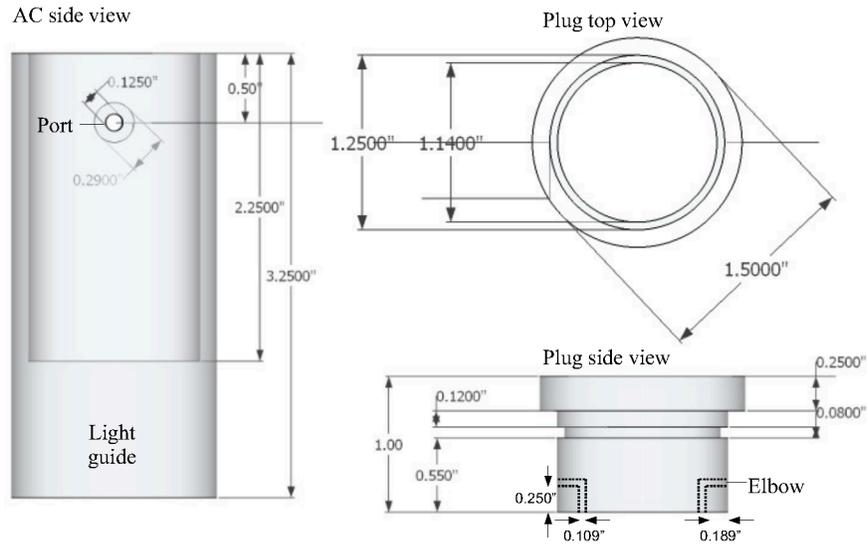


Figure 5.3: We designed an acrylic chamber (AC) to hold the LS sample. The central $1\frac{1}{2}$ in of the AC holds the sample, while 1 in light guides on either end allow for optical coupling to PMTs. The AC has an outer diameter of $1\frac{1}{2}$ in (which matches the PMT diameter) and a wall thickness of $\frac{1}{8}$ in, corresponding to approximately 30 mL of LS sample.

5.2.2 Liquid scintillator process system

To measure quenching in the LS sample, it is critical to first remove excess O_2 and H_2O , which act as quenching agents [124]. Also, since reoxygenation may

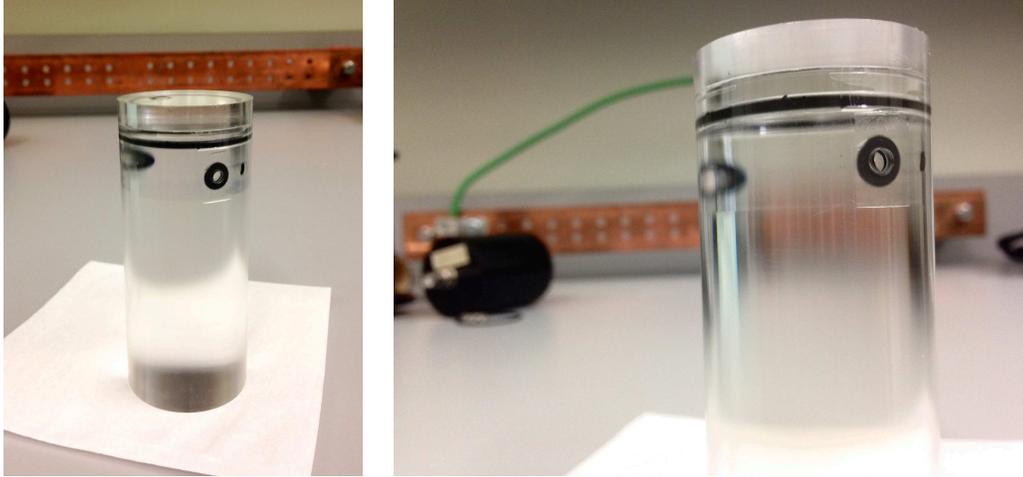


Figure 5.4: Photos of the AC show how the o-rings sit in the plug and around the ports. In these photos, the AC is filled with LS, so the ports and elbows are not aligned. At these viewing angles, the bottom of the plug is also visible.

occur quite rapidly when the LS is left in contact with the atmosphere [125], the deoxygenation process should either take place with the LS *in situ* or in such a way that the LS is transferable into the AC in isolation from the atmosphere.

I designed a LS process system (LSPS) that enables deoxygenation of the LS sample in isolation from the atmosphere via vacuum distillation³, thus satisfying the two criteria above. Vacuum distillation is also one of the techniques SNO+ employs to purify the LS (see section 2.1.2). Figure 5.5 shows a schematic of the LSPS. The degassing occurs in a LACO LVC0812-1121-VC acrylic vacuum chamber. An analog pressure gauge in the lid of the chamber allows me to monitor the pressure. A second port in the lid connects the chamber to an AdixenTM (Alcatel[®]) ACP 15 dry vacuum pump that drives the degassing. A digital pressure gauge at the pump allows me to monitor the

³The ability to bubble N₂ gas through the LS exists in the LSPS, but I never used this alternate deoxygenation method.

pressure in the system. I pump on the LS until the pressure reaches approximately 0.40 mbar (about 7 minutes of pumping time in this setup), which is sufficiently long that no bubbles are visible in the sample.

The AC is connected into the LSPS via two $\frac{1}{8}$ in tube fittings. These fittings pass through the ports in the AC body until they are flush with the inner AC surface. Pushing the plug into place in the AC compresses the internal o-rings onto the fittings, forming a seal against the atmosphere⁴. However, because it is difficult to form a robust seal against the curved surface of the plug when the ports and elbows are aligned, a second set of identical o-rings are compressed on the outside of the AC body using two aluminum bars that are brought together by tightening joining screws. I position the AC in the LSPS with the plug on the top, which causes the base of the plug where the elbows pass through to be the highest point in the internal AC volume.

After isolating the vacuum chamber from the pump (after the LS degassing is complete), I use the pump to evacuate the remaining lines, as well as the AC and a Mason jar that acts as a buffer between the LS in the LSPS and the pump (to prevent LS from reaching the pump and damaging it). The AC is situated in the LSPS between the vacuum chamber and the Mason jar. In terms of vertical positioning, the bottom of the Mason jar sits approximately 11 in below the bottom of the vacuum chamber, but completely above the AC, the top of which sits approximately 15 in below the bottom of the vacuum chamber. This geometry enables me to use a gravity drainage method to fill LS into the AC by opening a valve at the bottom of the vacuum chamber as I continue to pump through the Mason jar. Also, because the AC is positioned

⁴I put the connector tubes in place while assembling the AC. This later allows me to easily position the AC into LSPS.

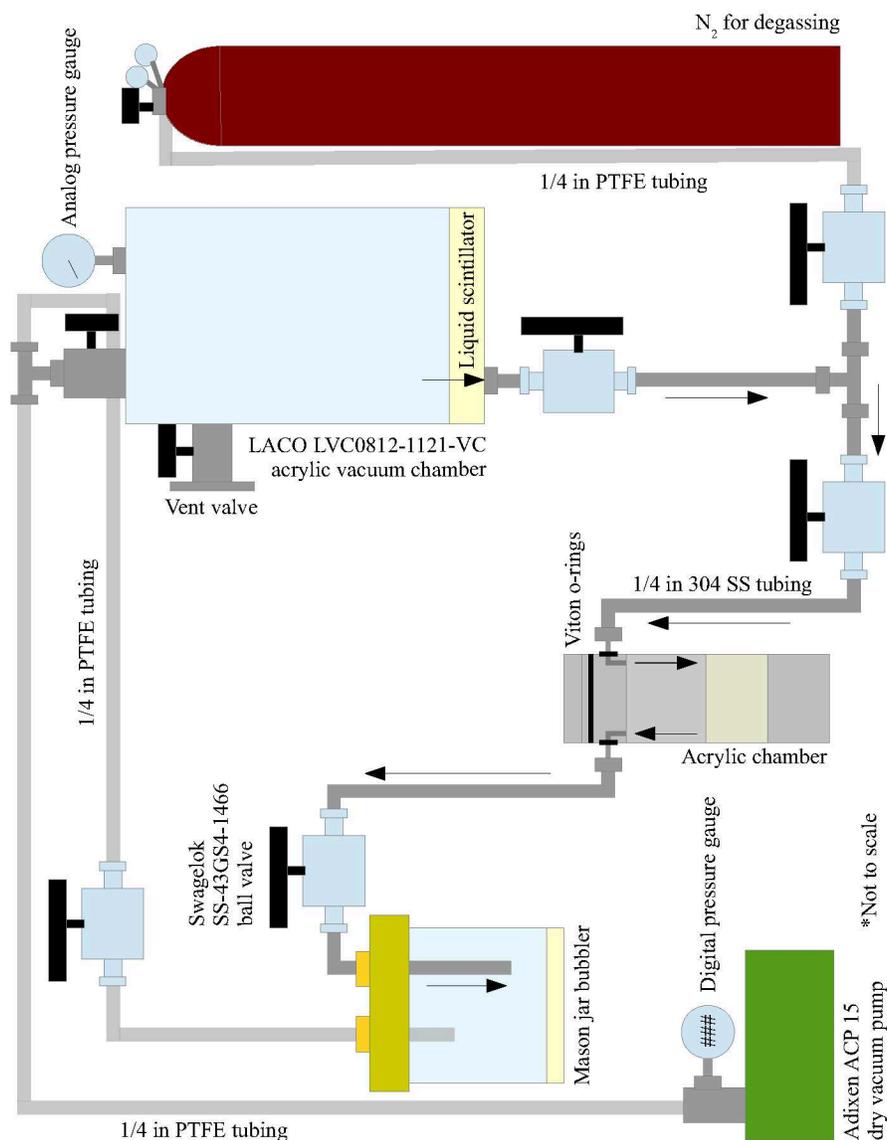


Figure 5.5: The LSPS uses an acrylic vacuum chamber for vacuum degassing of the LS, which is then gravity-drained into the AC in isolation from the atmosphere. A Mason jar protects the pump from overflow LS in the tubing. All tubing in which LS flows is $\frac{1}{4}$ in 304 stainless steel. Other lines are $\frac{1}{4}$ in PTFE. The 304 stainless steel ball valves are non-lubricated. An option for bubbling N₂ through the LSPS exists. Black arrows indicate the direction of LS flow.

with the elbows as the top-most point, any residual gas bubbles present in the LSPS easily pass from the AC through the elbows on to the Mason jar. Typically, some handling of the AC is necessary to ensure bubbles are dislodged from within the elbows and from around the o-rings. I visually monitor the AC as I slowly tip it to and fro to watch for all bubbles to clear the volume.

Once the AC is full of LS, I close valves on either side of the AC to isolate it from the rest of the LSPS. At this point, I carefully rotate the plug⁵ until the opening of the elbows are completely on the outside of the internal o-rings. This creates an effective and robust seal against the atmosphere, even after I have removed the external o-rings, preventing the ingress of O₂ or other quenching contaminants. I then remove the AC from the LSPS and clean any residual LS or other contaminants from the outside using a Kimwipe wetted with RO water containing Alconox.

I thoroughly clean each component of the LSPS before the addition of the LS into the vacuum chamber. I use a 15 min ultrasonic bath of RO water and Alconox at 55 °C to clean the acrylic vacuum chamber and AC, followed by a 15 min ultrasonic bath at room temperature in RO water. To ensure the acrylic does not contain residual water (which could leach into the LS), I dry all acrylic parts in a forced air oven⁶. All stainless steel tubing, connectors and valves, as well as the o-rings undergo an identical cleaning procedure. After rinsing, I leave these parts to air dry (under a clean Kimwipe) for approximately 24 h. I also give all stainless steel surfaces that will come into contact with the LS

⁵The plug fits *very* snugly in the AC body because of the circumferential o-ring seal. This means rotating the plug requires me to secure a hose clamp to the $\frac{1}{4}$ in protruding end of the plug and use pliers to do the rotation—it is nearly impossible to accomplish by hand.

⁶I raise the temperature of the oven to 85 °C over 75 min, hold it at 85 °C for 4 hours, then cool back to room temperature over 6 hours, roughly following the appropriate acrylic annealing cycle (though not quite as warm) set out in [97].

a rinse with methanol, then allow them to air dry, once again.

5.2.3 Detector assembly

To turn the LS-filled AC into a particle detector, I optically couple two Hamamatsu R580 PMTs to the AC, one on each light guide. The optical coupling compound is BC-630. I apply a thin layer to the PMT face, then press it onto the AC, sliding it over the acrylic until there are no visible air bubbles trapped between the PMT face and light guide, which could otherwise adversely affect the detector's optical efficiency. The PMT tends to seal well to the acrylic, so that, holding the AC vertically, I am able to couple both PMTs effectively.

I next wrap the two PMTs and AC in a sheet of aluminum foil, which acts as a reflector to increase the optical efficiency of the setup [124]. To provide an opaque seal against ambient light in the laboratory, I wrap the foil in a layer of black vinyl tape. No light leaks are evident in this configuration.

Custom bases for the PMTs were built in the Department of Physics, according to the standard voltage distribution ratios supplied by Hamamatsu [135]. E678-12A sockets connect the PMT to the electronics board, both of which are held within an aluminum cap. This cap connects to an aluminum tube that runs the length of the PMT body and supports the PMT. I connect the two tubes together using threaded rods, which help maintain the structural integrity of the detector (so the PMTs do not decouple from the AC), and which ensure the AC is centred between the aluminum tubes (providing detector symmetry). Figure 5.6 shows the detector at various steps along the assembly procedure.

A HPGe detector⁷ is the other component comprising the detector system.

⁷It is a GC12023 HPGe detector from Canberra Industries.



Figure 5.6: (Top left) I coat the face of an R580 PMT with BC-630 coupling compound before (top right) pressing the AC onto the face until no air bubbles remain between the face and the light guide. Next, I wrap the AC in aluminum foil (middle right) followed by an opaque layer of black vinyl tape (bottom right). Aluminum caps house the PMT electronics, with the two sides held in place using threaded rods held in aluminum sleeves (bottom left).

It records the energy of the gamma-ray that Compton scatters within the LS, thus providing an inferred measurement of the electron energy inside the LS volume. The HPGe detector consists of a germanium crystal 87.1 mm in diameter by 90.1 mm in length housed in a 101.6 mm long aluminum end cap. The HPGe sits as the centre of a multi-component shield, consisting of a 1 in thick copper box with 12 in side lengths (including the thickness) surrounded by 10–20 in of lead. For this experiment, I removed the shielding above the

HPGe detector to expose the HPGe aluminum end cap.

I constrain the energy of the electrons being produced in the LS volume by restricting the range of possible Compton scatter angles that could lead to the subsequent gamma-ray reaching the HPGe detector. A jig I constructed from extruded aluminum framing holds the gamma-ray source and AC detector along one axis while the centres of the AC detector and HPGe end cap are fixed on a second axis that forms an angle θ , the Compton scattering angle, with the first. The finite size of the HPGe end cap⁸ allows a range of scattering angles about θ to reach the HPGe detector, implying a range of possible electron energies in the LS for any geometry. The jig also holds a 2 in thick lead brick with a 7.5 mm bore hole in its centre between the gamma-ray source and the AC detector. This loosely collimates the gamma-rays from the source toward the centre of the LS volume. The 2 in of lead attenuates more than 99.8% of the 661.7 MeV gamma-rays from a ¹³⁷Cs source that do not pass through the bore hole [136], helping to more tightly constrain the possible electron energies in the LS for any geometry. Figure 5.2 shows a schematic of this setup.

Because of the geometry of the HPGe detector and shielding, changing the angle θ requires a vertical rotation of the jig arm holding the gamma-ray source. The jig also enables the centre of the AC detector to be shifted horizontally relative to the centre of the HPGe end cap in such a way that a direct path connecting the gamma-ray production point to the HPGe detector via a Compton scatter in the LS is always available. As a result of the distance between the gamma-ray source and AC detector (25 cm) and the AC and HPGe detectors (typically 30 cm), θ typically covers a 10–20° range. Figure 5.7 shows

⁸The gamma-ray source, AC detector and HPGe detector are not far enough separated to treat each as a point with respect to the others.

the jig holding a ^{137}Cs source and the AC detector in its position on the HPGe shielding.

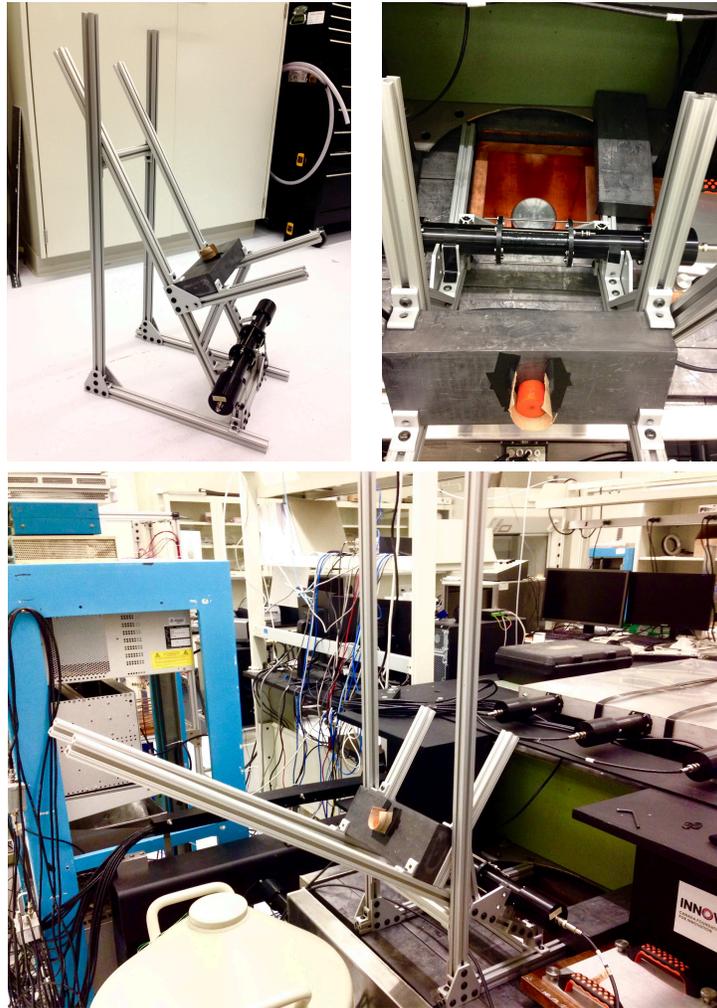


Figure 5.7: A jig constructed from extruded aluminum framing holds the AC detector in place (top left). The jig constrains the geometry of the AC detector relative to the HPGe detector, with a relatively high-angle scattering geometry shown (bottom). The bored lead brick located between the ^{137}Cs source (orange cylinder in top right picture) and AC detector limits the solid angle from the gamma-ray production point to the HPGe detector after a single Compton scatter in the LS. The HPGe end cap is visible just above the AC detector, though it sits below the AC detector inside the copper shielding.

5.2.4 Electronics, trigger logic and data acquisition

A series of Nuclear Instrumentation Modules (NIMs) provides timing, shaping and discrimination of the raw signals arriving from each of the three sensitive instruments (one HPGe and two PMTs) this experiment utilizes. In addition, the NIMs perform logic summing of these signals to provide a trigger that indicates a coincident pulse between either of the two PMTs and the HPGe. I describe the electronics, trigger and data acquisition in more detail below. Figure 5.8 provides a schematic of the electronics set up.

A Canberra model 2004 preamplifier converts the raw charge in the HPGe to a voltage pulse, the height of which is proportional to the charge (which is itself proportional to the energy deposited by a gamma-ray in the HPGe). This voltage pulse then passes to an Ortec 572 Amplifier⁹, which shapes the voltage pulse (shortening it to a $2.4 \mu\text{s}$ width) while maintaining the linearity between the pulse height and the input energy. The unipolar output of the amplifier connects to one channel on a Tektronix DPO 5204 Digital Oscilloscope, which records the amplified pulses. The bipolar output feeds an Ortec 455 Timing Single Channel Analyzer (SCA) that discriminates pulses between approximately 59 keV and 837 keV (corresponding to unipolar pulses ranging from 106 mV to 1373 mV, where the full deposition energy of a ^{137}Cs gamma-ray is 1087 mV). Pulses accepted by the SCA feed a LeCroy 222 Dual Gate Generator, which produces a 780 ns wide gate for the HPGe pulse. This gate is input into a LeCroy 365AL 4-Fold Logic Unit.

⁹While I use a Canberra Lynx digital signal analyzer to supply the HPGe with a bias voltage of +1500 V, I do not use it to acquire and analyze the HPGe pulse. The Lynx is not capable of performing coincidence spectroscopy with other detectors, which is necessary for this experiment. It also does not supply the times at which the HPGe events occur, which could otherwise be used for offline comparison of HPGe events and PMT events.

A WIENER MPOD EHS 8 020p high voltage module supplies a high voltage of +2000 V to a custom-designed splitter box that strips the PMT signal from the coaxial cable carrying the high voltage. The box design reduces the high voltage on the PMT by a factor of $\frac{4}{5}$ to +1600 V, above the manufacturer-suggested operating voltage, but below the maximum-specified voltage of +1750 V. The higher operating voltage increases the PMT gain, which is beneficial in observing very low light levels. By keeping the PMT operating voltage at +1600 V, the PMT remains in its prescribed linear response region. The oscilloscope collects the raw PMT pulses before they continue on to an Ortec 9302 Amplifier Discriminator. I use this module only to amplify the PMT pulses prior to their passage on to a LeCroy 623 Octal Discriminator¹⁰. Pulses that survive the discriminator are delayed approximately 3.5 μ s by an Ortec 427A Delay Amplifier in order to better match the timing of the HPGe pulse, which is much slower than those of the PMTs. Having the PMT and HPGe pulses closely matched in time allows for a narrower trigger window when looking for temporal coincidences between the two (or three) instruments, thus reducing the number of random or uninteresting pulses that satisfy the coincidence condition. The output from the delay amplifier feeds a LeCroy 222 Dual Gate Generator, which produces an approximately 210 ns wide gate pulse for each PMT. These gates are input into the same logic unit as the HPGe pulse.

The logic unit first evaluates the gates arriving from the two PMTs, producing a pulse if a gate from either or both PMTs arrives. As a result of the

¹⁰The discriminating ability of the Ortec 9302 is limited to pulses of 50 mV, which are larger than those in which I am interested. The LeCroy discriminator provides a lower discriminator threshold at 30 mV. It was also not clear whether the discriminator circuit in the Ortec module was functioning in a robust way, whereas the LeCroy discriminator provided a predictable result.

timing along the PMT electronics chain, the gate from PMT B is always delayed compared to that from PMT A by approximately 40 ns. The logical sum of the the two PMT channels passes to a second logic input that compares it to the input HPGe gate. If both gates exist and are coincident, then the logic unit outputs a pulse that acts as an external trigger input to the oscilloscope. Because the HPGe gate is much wider than the PMT gates (and the subsequent PMT logic pulse), it always fully contains the PMT pulse, if one exists. The gates generated for each PMT pulse are added to this trigger gate using a LeCroy 428F Linear Fan-in/Fan-out. In this way, the trigger pulse indicates which PMTs had a pulse coincident with the HPGe signal. Small time offsets for each gate make this determination unambiguous.

As I mentioned above, a Tektronix DPO 5204 Digital Oscilloscope records the amplified HPGe pulse and raw pulses from the two PMTs, each on a different channel. The fourth channel reads in the logic pulse, on which the oscilloscope triggers. Upon triggering, the full waveform from each channel is recorded to an external drive using a library that allows a connection to the oscilloscope via the VXI-11 RPC (ethernet) protocol [137]. I modified this library to enable the simultaneous readout of multiple oscilloscope channels and to provide real-time graphical output of the data using Root. Two additional data files are also recorded. One file contains the oscilloscope settings during the data collection period (to ensure the oscilloscope is set up the same way for each set of data collected), while the other contains parameters related to the oscilloscope voltage and time gains for each recorded channel, which are required for later data processing.

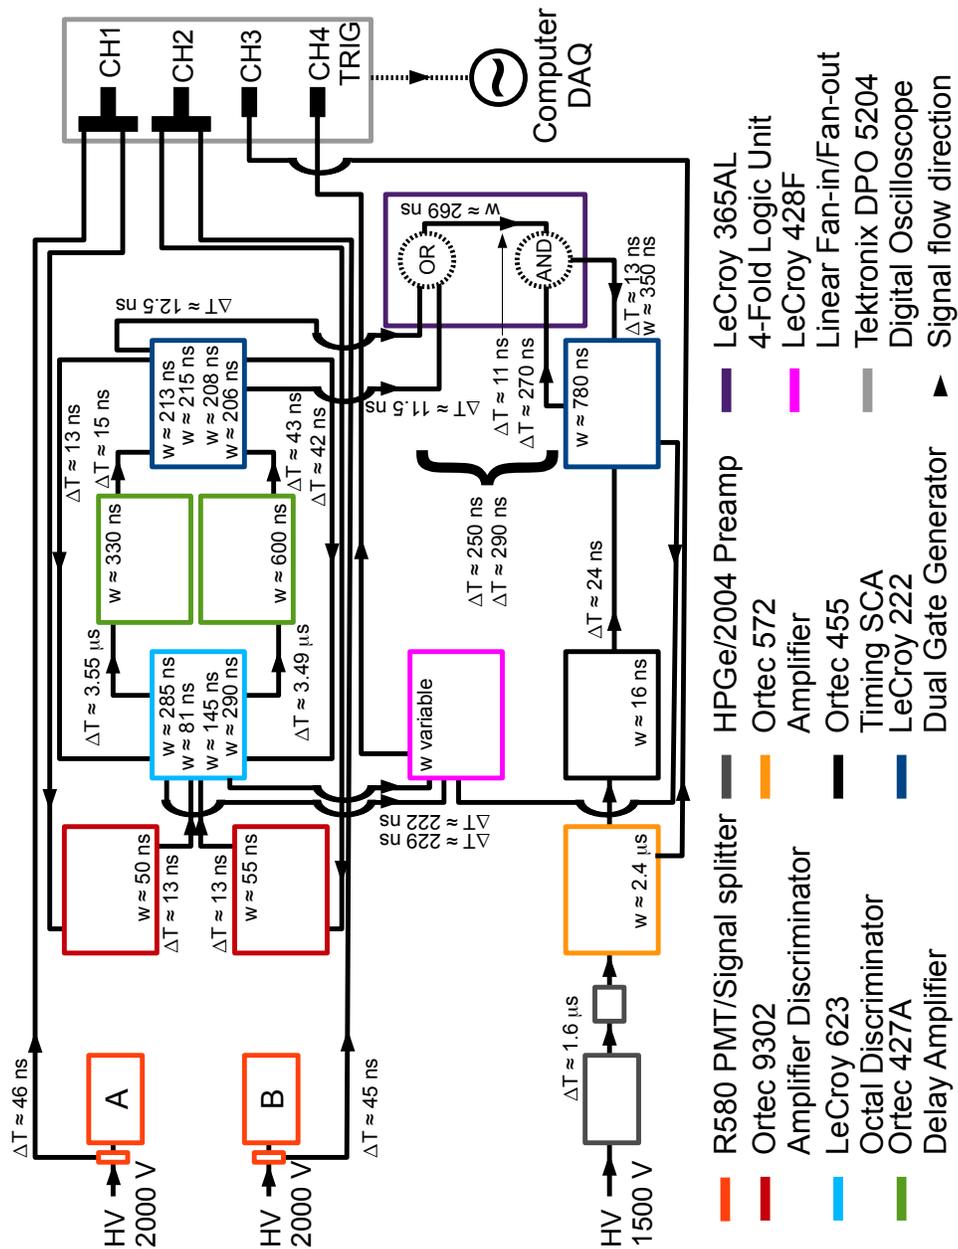


Figure 5.8: This diagram shows the set up of quenching measurement electronics. w is the width of the pulse output by that particular module and ΔT is the time difference between the leading edge of pulses output by that module and the next. The ΔT s shown next to the left parenthesis refer to the time difference between the pulse of each PMT output by the Ortec 9302 Amplifier Discriminator and the HPGe pulse output by the Ortec 572 Amplifier.

5.3 Data

The data I collect in this experiment are HPGe and PMT pulses resulting from particle interactions in the HPGe and LS. A Tektronix DPO 5204 Digital Oscilloscope samples and records the voltage of each waveform 1000 times in a $5 \mu\text{s}$ snapshot. The voltage-axis of the oscilloscope consists of 256 analog-to-digital converter (ADC) channels, 250 of which are segmented into 10 divisions (with three above and three below those 10 divisions). The voltage gain is variable and set independently for each oscilloscope channel. For each data set, the HPGe channel is set at a gain of 120 mV/div, whereas the PMT channels vary from 1 mV/div (the minimum possible) up to 24 mV/div, depending on the setup of the experiment (and, consequently, the energy being deposited in the LS).

A ^{137}Cs 661.657 MeV gamma-ray creates light in the LS via Compton scattering and the scattered gamma-ray may then produce charge in the HPGe. When pulses from either of the PMTs and the HPGe are coincident, the oscilloscope triggers and writes the waveforms of each instrument to file. The ^{137}Cs source I use has an activity of approximately 370 kBq, which, given the geometry of the experiment (see section 5.2.3), leads to a trigger rate of up to a few Hz. I collect several independent sets of data for each LS sample in order to scan the LS response over a range of energies. Again, because of the experimental geometry, each data set covers a 50–100 keV energy range, whereas the collated data cover the approximate energy range 0–300 keV.

Section 5.3.1 describes how I calibrate the voltage axis of each channel to correspond to the relevant quantity each channel is measuring (an energy for the HPGe and the average number of observed photoelectrons for the PMTs).

Each waveform is written to file for later offline processing. Section 5.3.2 outlines how this processing works, and describes how the raw data are converted into useful information for later analysis.

5.3.1 Calibration

I calibrate the energy response of the HPGe using a set of gamma-ray calibration standards, including ^{57}Co , ^{109}Cd , ^{137}Cs and ^{133}Ba . Pulse height distributions from each source are collected using the same electronics chain I describe in section 5.2.4 where the only gate that contributes to the logic pulse is that from the HPGe. Peaks from each source indicate the ADC to energy conversion and also provide a measure of the energy-dependent resolution of the detector. Figure 5.9 shows the calibration data, including the calibration and resolution functions.

I determine the peak positions for each gamma-ray source using Root's `TSpectrum` class¹¹. I assume a linear model for the energy scale calibration, where the best fit to the data results in

$$E = (0.614 \pm 0.003) \text{ keV/ADC} \cdot \text{ADC} + (-5.72 \pm 1.62) \text{ keV}. \quad (5.5)$$

To determine the resolution, I fit a Gaussian function to each peak (the three bins to either side of the peak position found by `TSpectrum`) and take the fitted standard deviation as the energy-dependent resolution. Fitting the energy and resolution data, I find the energy resolution σ of the HPGe in this configuration

¹¹`TSpectrum` uses the second derivative of the waveform to identify peaks on top of a noisy background. It identifies peaks when the change in the derivative is large compared to the average change.

to be, with σ and E in units of keV,

$$\sigma(E) = (0.86 \pm 1.85) \times 10^{-1} \sqrt{E} + (-1.34 \pm 5.52) \times 10^{-3} E + (1.90 \pm 1.44) . \quad (5.6)$$

I calibrate the PMTs in terms of the mean number of photoelectrons they observe by dividing the PMT pulse height by the average height of the PMT's single photoelectron (SPE) response. To measure the SPE response, I placed each PMT in a dark box and supplied it with its nominal +1600 V. I used a Tektronix AFG 3021B Single Channel Arbitrary/Function Generator to drive an LED (blue) with 30 ns wide and 3 mV tall square pulses at a rate of 1 kHz. Light from the LED travelled into the dark box via an optical fibre, where the open end of the fibre was pointed toward the centre of the PMT face. I collected the PMT SPE data in the same way as all other PMT data via the oscilloscope, except the TTL output from the function generator provided the trigger signal. To ensure the injected light was producing an exclusively SPE response from the PMT (or as near to it as possible), I ensured that about only 10% of triggers contained a true PMT pulse¹².

Using the model in [138] (with additional notes on this model coming from [139]), I fit the pulse height distribution of each PMT to determine the mean pulse height of each PMT's SPE response. This model assumes the number of photons in each pulse of light from the LED is Poisson distributed and that the production of a photoelectron (PE) given a photon incident on the PMT photocathode is binomially distributed. Thus, the probability of

¹²If the function generator produced pulses that were much wider than 30 ns or any taller than 3 mV, then the PMT response clearly contained multi-photoelectron pulses.

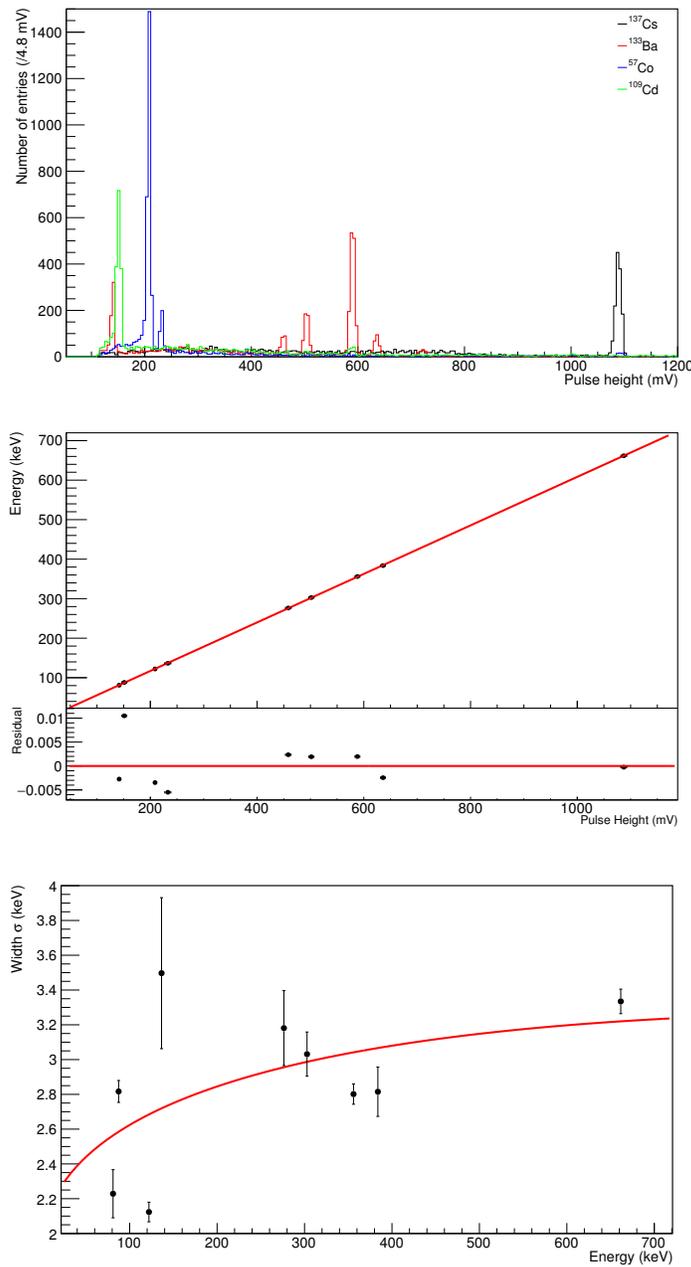


Figure 5.9: I use ^{57}Co , ^{109}Cd , ^{137}Cs and ^{133}Ba gamma-ray sources to calibrate the energy scale and resolution of the HPGe detector. The position of the main gamma-ray peaks from each source (top) provide a linear energy scaling (middle), while the width of each determines the energy-dependent resolution of the HPGe (bottom).

observing n PE when the mean number of PE the first dynode collects is μ is

$$P(n; \mu) = \frac{\mu^n e^{-\mu}}{n!} , \quad (5.7)$$

where μ depends both on the PMT quantum efficiency and on the number of photons striking the photocathode. Assuming a large gain at the first dynode, the charge output at the end of the dynode chain is well-represented using a Gaussian distribution

$$G_1(x) = \frac{1}{\sqrt{2\pi\sigma_1^2}} e^{-\frac{(x-Q_1)^2}{2\sigma_1^2}} , \quad (5.8)$$

with x the variable charge, and Q_1 and σ_1 the average charge and the standard deviation of the charge distribution for SPE. For n -photoelectron (n PE) pulses, the charge distribution becomes the convolution of n SPE distributions

$$G_n(x) = \frac{1}{\sqrt{2\pi n\sigma_1^2}} e^{-\frac{(x-nQ_1)^2}{2n\sigma_1^2}} . \quad (5.9)$$

The response for an ideal, noiseless PMT is then the convolution of equation 5.7 and equation 5.9, giving

$$\begin{aligned} S_{\text{ideal}}(x) &= P(n; \mu) \otimes G_n(x) \\ &= \sum_{n=0}^{\infty} \frac{\mu^n e^{-\mu}}{n!} \frac{1}{\sqrt{2\pi n\sigma_1^2}} e^{-\frac{(x-nQ_1)^2}{2n\sigma_1^2}} . \end{aligned} \quad (5.10)$$

The model also accounts for two types of background that add to the n PE response distribution. The first is due to low-charge processes that cause a spread in the signal when the photocathode emits zero PE. This so-called pedestal is Gaussian distributed. The second type of background uses an

exponential function to model discrete noise processes, like thermal emission, internal and external radioactivity, &c. If, given a pulse resulting from one of these background processes, we set the probability it is of the second type to w , then the background parametrization is

$$B(x) = \frac{1-w}{\sqrt{2\pi\sigma_0^2}} e^{-\frac{x^2}{2\sigma_0^2}} + w\theta(x)\alpha e^{-\alpha x} . \quad (5.11)$$

Here, σ_0 describes the width of the pedestal, which the model assumes is always present with mean charge Q_0 , α describes the falloff of the second background type with charge x , and $\theta(x)$ demands that backgrounds of the second type have a charge $x \geq 0$.

The real PMT response, then, is the convolution of the ideal PMT response in equation 5.10 with the background in equation 5.11

$$S_{\text{real}}(x) = \sum_{n=1}^{\infty} \frac{\mu^n e^{-\mu}}{n!} [(1-w)G_n(x-Q_0) + wI_n(x-Q_0)] , \quad (5.12)$$

where

$$\begin{aligned} I_n(x-Q_0) &= \int_{Q_0}^x G_n(x'-Q_0)\alpha e^{-\alpha(x-x')} dx' \\ &= \frac{\alpha}{2} e^{-\alpha(x-Q_n-\alpha\sigma_n^2/2)} \\ &\left[\operatorname{erf}\left(\frac{|Q_0-Q_n-\alpha\sigma_n^2|}{\sqrt{2\sigma_n^2}}\right) + \operatorname{sign}(x-Q_n-\alpha\sigma_n^2)\operatorname{erf}\left(\frac{|x-Q_n-\alpha\sigma_n^2|}{\sqrt{2\sigma_n^2}}\right) \right] \end{aligned} \quad (5.13)$$

and

$$Q_n = Q_0 + nQ_1 \quad (5.14)$$

$$\sigma_n = \sqrt{\sigma_0^2 + n\sigma_1^2} . \quad (5.15)$$

Q_1 and σ_1 correspond to the average SPE charge (with the pedestal charge Q_0 subtracted) and the equivalent width of the SPE distribution, respectively.

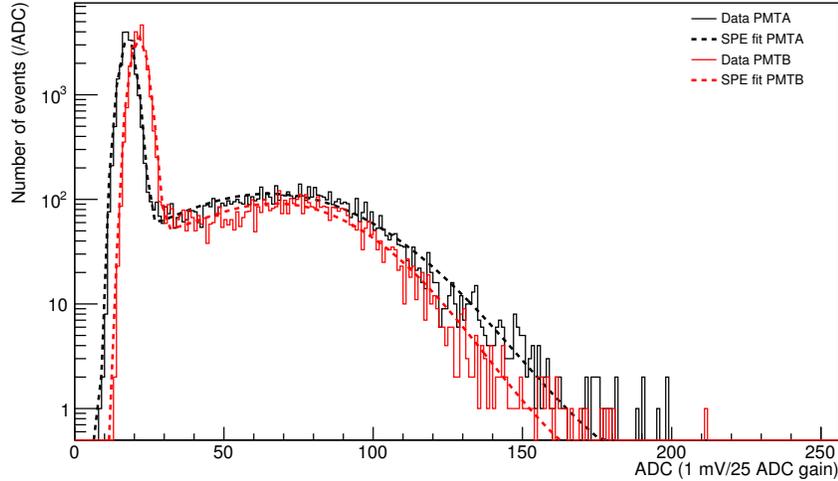


Figure 5.10: Fits (dashed lines) to the pulse height distributions for PMTA (black) and PMTB (red) indicate average SPE pulse heights of 47.81 ADC and 46.69 ADC, respectively. The pedestal is the large peak at the left of each distribution. The different position for each pedestal simply implies that the baseline for each oscilloscope channel was set differently.

Figure 5.10 shows the SPE data for both PMTs along with the best fit of equation 5.12 to each data set. The best fit for each PMT, based on the reduced χ^2 of the fit, was for $n = 1$ in equation 5.12 (I attempted fits up to $n = 4$). Table 5.1 lists the best-fit parameters for each PMT (aside from an overall normalization parameter). Note that the SPE data correspond to the distribution of PMT pulse heights, whereas equation 5.12 requires the total *charge* output in the pulse. Because the total charge and pulse height for these PMT pulses are linearly related, equation 5.12 applies equally well to the PMT pulse height distribution.

Parameter	PMTA		PMTB	
Q_0 (ADC)	17.78	± 0.02	21.52	± 0.02
σ_0 (ADC)	2.19	± 0.01	2.32	± 0.01
Q_1 (ADC)	47.81	± 0.46	46.69	± 0.60
σ_1 (ADC)	29.56	± 0.42	25.28	± 0.55
μ	0.353	± 0.006	0.222	± 0.006
w	0.039	± 0.004	0.054	± 0.005
α	0.026	± 0.003	0.046	± 0.003
χ^2 / NDF	710.4 / 162		829.1 / 149	

Table 5.1: The fit SPE parameters for PMTA and PMTB indicate average SPE pulse heights of 47.81 ADC and 46.69 ADC, respectively, for the two PMTs. These correspond to 1.91 mV and 1.87 mV (given the PMT gain of 1 mV / 25 ADC). The gain for each PMT is slightly different, with that for PMTA being slightly higher.

5.3.2 Data processing

I search each raw waveform (both PMTs, the HPGe and the trigger) for pulses using a derivative-based pulse-finding algorithm. Figure 5.11 provides an example waveform for PMTA and the HPGe. The data I extract from each waveform are the pulse height and the pulse time. The pulse heights encode information about the energy each instrument observes, while the pulse times enable the search for events coincident between the detectors. Though the electronics supply a trigger only upon coincidence between at least one PMT and the HPGe, the software assists in refining that signal. The software identifies which combination of PMTs provided a signal, enabling the extraction of events where the two PMTs together observe only a single PE. It also determines the relative timing of the pulses from each instrument, which is useful in discriminating true coincident pulses from, say, PMT pulses that randomly fall within the relatively wide HPGe gate and cause a hardware trigger. Lastly, it screens events where PMT noise caused a hardware trigger, which typically results in a PMT waveform that contains no discernible pulse.

I begin by determining the average times at which the largest pulse occurs in the first 1000 collected PMT and HPGe waveforms. Fitting distributions of these times for each instrument with a Gaussian function gives me an estimate of where to expect to find pulses in each waveform (which is relatively constant in time, given the electronics setup I describe in section 5.2.4). For the PMTs, I search for pulses in a $\pm 20\sigma$ time window about the mean minimum time, whereas for the HPGe I search for pulses in a $\pm 10\sigma$ time window about the mean maximum time, where σ refers to the standard deviation of the respective fit Gaussian function. I also calculate the mean derivative of the the waveform baseline by calculating the derivative δ_i at each of the first 32 time samples. I calculate the derivative of each waveform at time sample i as

$$\delta_i = y_{i+1} - y_{i-1} , \quad (5.16)$$

where y_i is the waveform ADC value at sample i . The derivative threshold for a waveform is the standard deviation of a Gaussian function fit to the distribution of δ_i for the first 32 time samples of the first 1000 collected waveforms.

I scan each PMT waveform for negative pulses in its respective time window. The waveform baseline is the average of the waveform over the first 30 time samples, with a width equal to the standard deviation of the waveform over that window. A pulse begins when the derivative of the waveform drops below the derivative threshold and y_i is more than five times the baseline width below the baseline. The pulse ends when the waveform derivative is positive and less than three times the derivative threshold away from the mean baseline derivative, and y_i is less than three times the baseline width below the baseline. If more than one pulse is found in the time window, then I discard

all but the pulse with the largest pulse height (effectively removing any after pulses from the data set). I define the pulse height to be the absolute value of the difference between the pulse minimum y_i and the waveform baseline. This transforms negative PMT pulses into positive pulse heights. I define the pulse time as that at which the waveform derivative (within a pulse) is minimal.

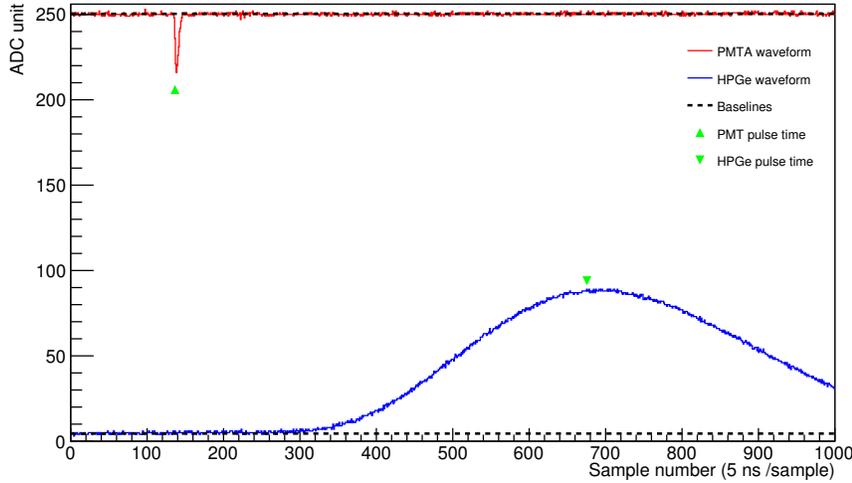


Figure 5.11: These sample waveforms for PMTA (red) and the HPGe (blue) are shown in their relative positions as they appeared on the oscilloscope. Dashed black lines indicate the baseline for each waveform. Green markers indicate the samples that the pulse-finding algorithm selected as the pulse times.

Searching for HPGe pulses proceeds in a similar way. I first determine the waveform baseline and baseline width, then search for the HPGe pulse in its respective time window. Because the HPGe pulse is very long compared to the PMT pulses, I calculate the derivative as

$$\delta_i = y_{i+4} - y_{i-4} , \quad (5.17)$$

which helps avoid finding false pulses as a result of jitter between adjacent

waveform samples. Pulses begin when the derivative is positive and more than three times larger than the derivative threshold, and y_i is more than three times the baseline width above the baseline. The pulse ends when the derivative is negative and less than three times the derivative threshold, and y_i is less than three times the baseline width above the baseline (or upon reaching the end of the waveform). The pulse height is the difference between the pulse maximum y_i and the baseline, and the pulse time is that at which the waveform (within a pulse) is maximal (different than the PMT timing definition). Again, if more than one pulse is found in the time window, then I discard all but the pulse with the largest pulse height.

I calibrate the HPGe pulse height using equation 5.5 to convert the ADC pulse height to an energy, $E_{\gamma'}$. I infer the apparent electron energy in the LS from the gamma-ray energy via

$$E_{e^-} = E_{\gamma} - E_{\gamma'} = 661.657 \text{ keV} - E_{\gamma'} . \quad (5.18)$$

I calculate the average number of PE each PMT observed by dividing the pulse height in each PMT by the average pulse height per SPE given by the respective values of Q_1 in table 5.1. Converting the PMT responses to the average number of PE observed then allows me to add the PMT responses, regardless of any difference in gain between the two PMTs.

To select events that correspond to single Compton scatters in the LS volume with a coincident gamma-ray in the HPGe, I make a series of cuts on the PMT and HPGe pulse timing and the HPGe energy distribution. First, I fit the distribution of time difference between pulses in PMTA and PMTB, $\Delta t = t_A - t_B$, with a Gaussian function, with mean $\overline{\Delta t}$ and width $\sigma_{\Delta t}$. I then

place a cut, $c_{\Delta t}$, on events for which Δt is large and a pulse in both PMTA (P_A) and PMTB (P_B) exists, where

$$c_{\Delta t} == |\Delta t - \overline{\Delta t}| < 2\sigma_{\Delta t} \ \&\& \ P_A \ \&\& \ P_B . \quad (5.19)$$

Figure 5.12 shows the Δt distribution and corresponding cut range.

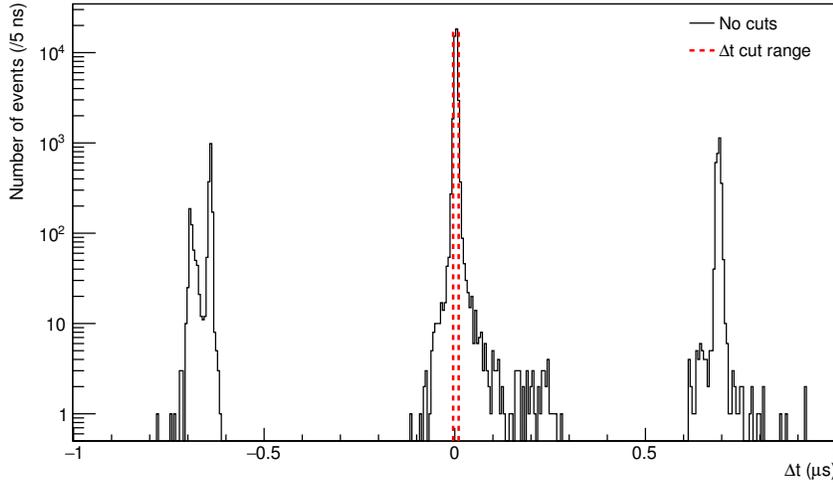


Figure 5.12: The distribution of time differences Δt between pulses in the two PMTs exhibits a strong central peak near $0 \mu\text{s}$, with two peripheral peaks that correspond to events where only one of the PMTs records a pulse. The dashed red lines correspond to the tight cut region of equation 5.19.

For events surviving $c_{\Delta t}$, I then form distributions of t_A and t_B . Each of these is fit with a Gaussian function to form timing cuts on the times of individual pulses, where I expect all single scatter event PMT pulses to occur near the mean pulse time for each PMT. I do this for events for which only one PMT registers a pulse, resulting in cuts c_{t_A} and c_{t_B} , and for events where both PMTs register pulses¹³, resulting in cut $c_{t_{AB}}$. These cuts are

¹³These pulses may be far away from the expected pulse time even though the time difference between them satisfies $c_{\Delta t}$.

$$\begin{aligned}
c_{t_A} &== |t_A - \overline{t_A}| < 3\sigma_{t_A} \ \&\& \ P_A \ \&\& \ !P_B \\
c_{t_B} &== |t_B - \overline{t_B}| < 3\sigma_{t_B} \ \&\& \ !P_A \ \&\& \ P_B \\
c_{t_{AB}} &== |t_A - \overline{t_A}| < 3\sigma_{t_A} \ \&\& \ |t_B - \overline{t_B}| < 3\sigma_{t_B} \ \&\& \ P_A \ \&\& \ P_B . \quad (5.20)
\end{aligned}$$

For events satisfying the overall PMT timing cut

$$(c_{\Delta t} \ \&\& \ c_{t_{AB}}) \ || \ c_{t_A} \ || \ c_{t_B} , \quad (5.21)$$

I further cut on the distribution of the HPGe pulse times t_H , again to ensure pulse times are near the average pulse time. The HPGe timing cut c_H is

$$c_H == |t_H - \overline{t_H}| < 3\sigma_{t_H} \ \&\& \ P_H , \quad (5.22)$$

where P_H refers to the HPGe having registered a pulse in the event.

For HPGe pulses surviving the cuts in equations 5.21 and 5.22, I perform a cut on E_{e^-} to select events that correspond to single Compton scatters in the LS volume that lead to full energy depositions of the gamma-rays in the HPGe. This removes events where the gamma-ray does not fully deposit its energy in the HPGe, which artificially inflates E_{e^-} . It also removes multi-scatter events, which have a different energy profile than those of interest. I fit the peak in the E_{e^-} distribution with a Gaussian function with mean $\overline{E_{e^-}}$ and width $\sigma_{E_{e^-}}$, and place a cut c_E on the remaining events, where

$$c_E == \overline{E_{e^-}} - 1.8\sigma_{E_{e^-}} < E_{e^-} < \overline{E_{e^-}} + 1.0\sigma_{E_{e^-}} . \quad (5.23)$$

The asymmetry in the cut energy region is a result of the overlap between the single scatter energy peak and the Compton edge of that peak. This overlap is due to there being a range of possible electron energies in any data sample resulting from the geometry of the experiment, as I discussed in section 5.2.3.

I consider events that pass the combined cut

$$c == (c_H \&\& c_E) \&\& ((c_{\Delta t} \&\& c_{AB}) || c_A || c_B) \quad (5.24)$$

to be candidate single Compton scatter, full deposition events. Figures 5.13 and 5.14 show the distributions of PMT and HPGe pulse times, and the distributions of the average total number of PE and inferred electron energies, respectively, as the cuts in equations 5.20, 5.22 and 5.23 are applied. These use a single set of TeLS data with a central scattering angle near 27° .

While c effectively removes most noise from the PMT and HPGe distributions, a small number of background events remain. These consist of a small number of PMT noise events that pass the PMT timing cut in equation 5.21 that sit in the 0–1 PE range. Additionally, some pulses in the HPGe that correspond to single or multiple Compton scatters of the gamma-ray rather than a full deposition survive the cut on E_{e^-} . This is due to the nature of the geometry admitting a range of electron energies in each sample of data and that the cut on E_{e^-} does not account for the correlation between the PMT and HPGe responses. However, the spatial separation of the ^{137}Cs source, AC and HPGe in the geometry, and the collimation of the gamma-rays minimizes this background contribution. Figure 5.15 shows two-dimensional distributions of the average total number of PE versus E_{e^-} as the timing and E_{e^-} cuts are applied. These correspond to the same data as in figures 5.13 and 5.14.

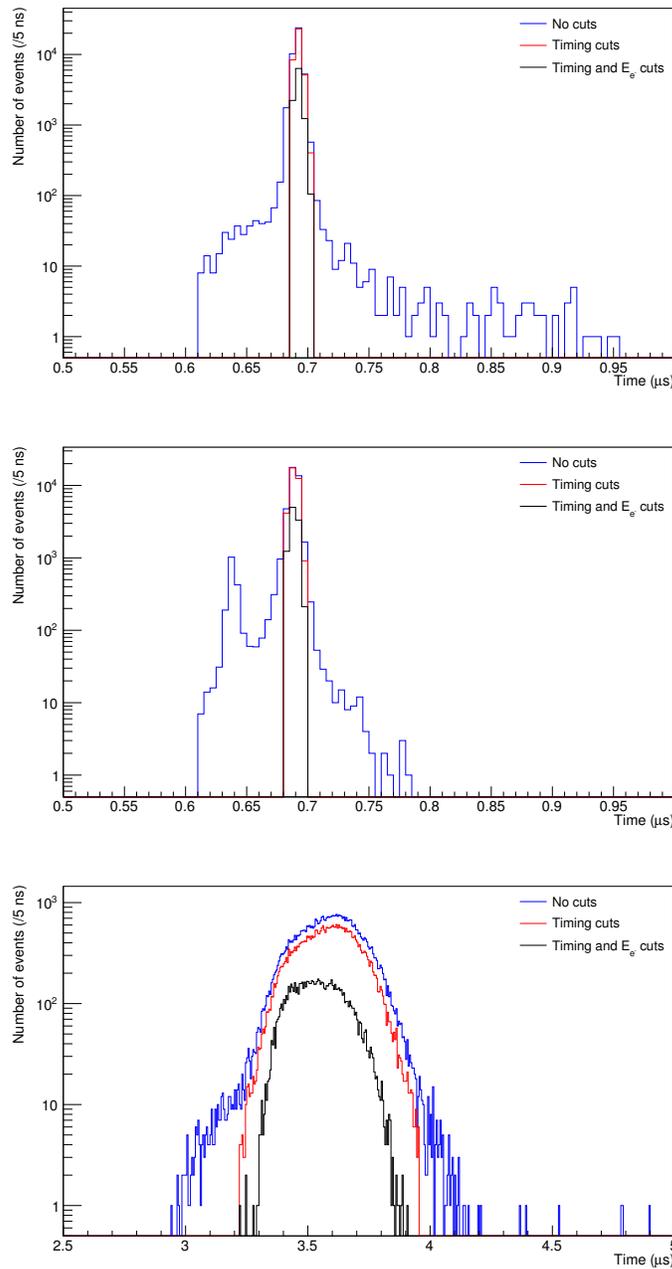


Figure 5.13: The spread in the uncut (blue lines) distributions of pulse times for PMTA (top) and PMTB (middle) is effectively reduced by the timing cuts of equations 5.20 and 5.22 (red lines). Cutting on E_{e^-} restricts the number of accepted events (black lines) without significantly changing the distributions' widths. The distribution of pulse times for the HPGe (bottom) is much broader than those for the PMTs. Applying the same timing cuts removes the distribution's tails and cutting on E_{e^-} again reduces the number of accepted events, while further narrowing the distribution.

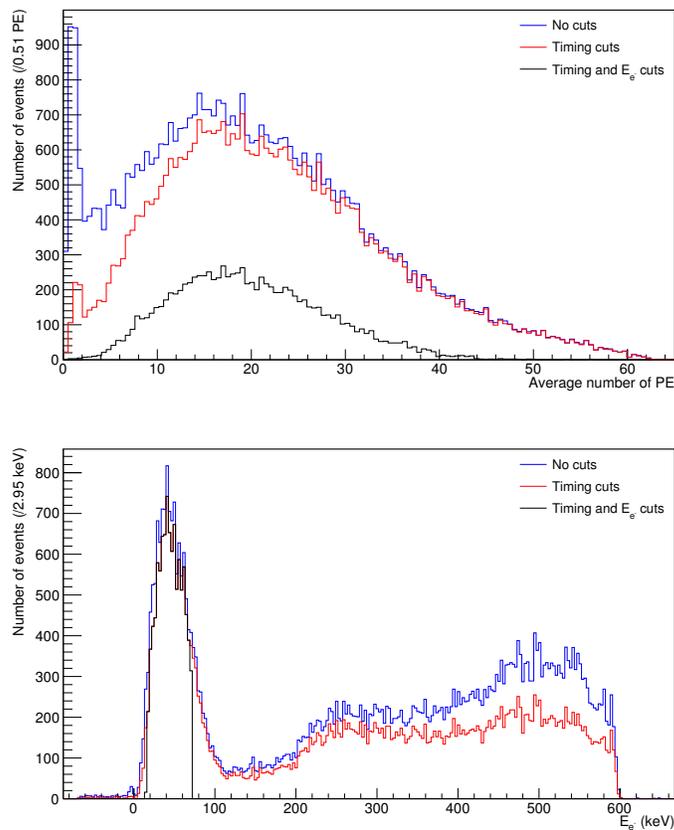


Figure 5.14: The distribution of the average total number of PE observed (top) exhibits a strong peak at low PE resulting from PMT noise (blue lines). The timing cuts of equations 5.20 and 5.22 (red lines) significantly reduce the low PE noise, while the cut on E_{e^-} further reduces low PE noise (black lines). The distribution of E_{e^-} (bottom) has a strong peak near the energy range $\sim 20\text{--}80$ keV. While the timing cuts have a small effect on that peak's shape, they do reduce the higher E_{e^-} region, which corresponds to the Compton continuum region of observed gamma-rays. The cut on E_{e^-} selects a specific region of that peak and is tuned to select mostly single Compton scatter, full deposition candidate events.

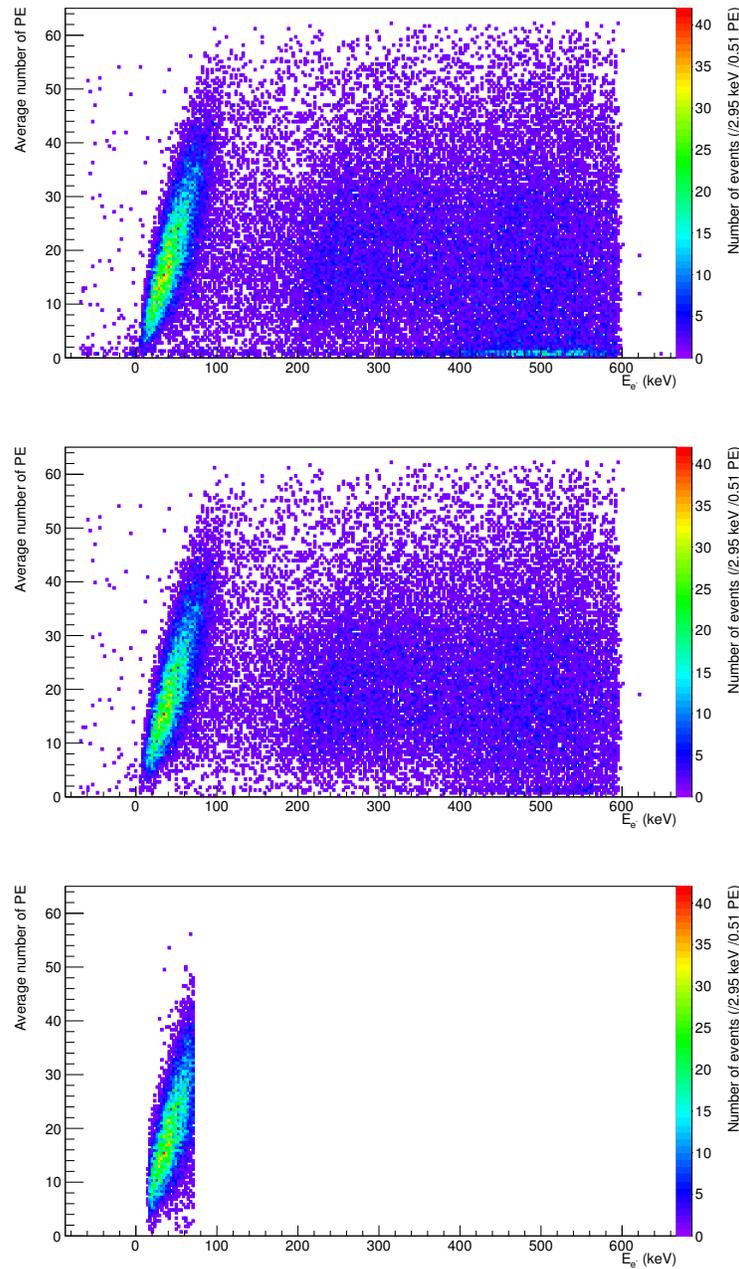


Figure 5.15: With no cuts applied (top), the distribution of average PE versus E_{e^-} shows several distinct regions. Between ~ 20 – 80 keV the data show a clear indication of single Compton scatter, full deposition events, with a Compton continuum in the region above ~ 200 keV. Low PE noise is visible near 0 PE across the entire energy range, with a significant number of events near the gamma-ray backscatter region of 500 keV. The timing cuts (middle) clear most of the low PE noise. Adding the cut on E_{e^-} (bottom) selects mainly single Compton scatter, full deposition events, cutting out noise from multiple Compton scatter gamma-rays.

5.4 Analysis

I measured the response to low-energy electrons of two different LS samples in this experiment. The first is the nominal SNO+ LS, which consists of LAB + 2 g/L PPO. The second sample is a TeLS cocktail, which is the nominal LS with the addition of 15.3 mg/L of bisMSB, 5% by mass of PRS and 0.3% by mass of telluric acid ($\text{Te}(\text{OH})_6$)¹⁴. The PRS is a surfactant that enables the $\text{Te}(\text{OH})_6$ to be dissolved in the LAB, while the bisMSB is a secondary fluor that further aids in preventing self-absorption of the scintillation photons. I processed each sample in the LSPS according to the method in section 5.2.2.

To investigate the LS response over a range of electron energies required I collect several sets of data for each sample, as the geometry restricted the available Compton scattering angles for each set. I collected nine sets of data for each sample, each containing 50000 triggered events. Each set of data covered a different range of scattering angles (with overlaps between the data sets to ensure sufficient coverage of the electron energy). Table 5.4 lists the approximate angular range and expected electron energy range for each data set. I calculated the expected energy based on the geometric setup of the experiment, where the physical extent of the surface of the HPGe end cap enables the detection of a range of Compton scattering angles.

To examine how the detector assembly (see section 5.2.3) affected the data, I collected two data sets for a LS sample covering the same angular range of 35–56 °. Between trials, I completely dismantled the PMT assembly, including decoupling the PMTs from the AC, then rebuilt it before placing it back in

¹⁴The LAB was distilled at SNOLAB. The PPO is from FLUKA (Sigma-Aldrich), the bisMSB was supplied by Brookhaven National Laboratory and the PRS is Stepan BioSoft N-411. The $\text{Te}(\text{OH})_6$ is from SEASTAR CHEMICALS, Inc. Both samples were prepared by O. Chkvorets at SNOLAB.

θ range ($^\circ$)	E_{e^-} (keV)	PMT gain (mV/div)		Event Rate (Hz)	
		LS	TeLS	LS	TeLS
0–16	0–33	1	1	3.00 ± 0.11	0.93 ± 0.04
4–20	2–46	1	3	1.92 ± 0.09	0.91 ± 0.03
14–29	25–90	3	6	1.72 ± 0.11	0.92 ± 0.04
21–35	51–123	3	6	1.34 ± 0.02	0.79 ± 0.02
24–37	64–138	6	12	1.23 ± 0.08	1.43 ± 0.07
27–48	81–198	6	24	0.88 ± 0.05	1.25 ± 0.07
35–56	128–243	12	24	0.78 ± 0.03	1.46 ± 0.07
48–69	198–302	12	24	0.76 ± 0.04	1.20 ± 0.08
54–77	230–331	24	24	0.88 ± 0.05	0.52 ± 0.02

Table 5.2: I collected nine sets of data for each of the LS and TeLS samples, where each set spanned a different range of Compton scattering angles, θ . I set the PMT gain for each data set in order to utilize the full range of the oscilloscope so that true pulses did not saturate the oscilloscope. To calculate the average event rate, I used the time required to obtain 100 triggers for five separate trials for each data set.

the detector jig. I neither replaced the LS in the AC nor moved the jig. Figure 5.16 shows the resulting average total number of PE and E_{e^-} for each trial. The distributions for E_{e^-} confirm that the geometry between trials did not significantly change, with a mean of 177.1 keV (on the distribution of cut E_{e^-} values, which is not shown) for both trials. The average total number of PE detected changed from 44.6 PE for trial 1 to 46.9 PE for trial 2, a discrepancy of 5.5%, which I take to be the systematic uncertainty on the absolute light output I measure.

I attribute this to a change in the optical efficiency of the PMT assembly resulting from the assembly process. The optics of the PMT assembly rely on an air gap between the reflective aluminum wrapping and the AC, as well as efficient coupling of the PMTs to the AC. The coupling compound may couple a small area of the aluminum to the AC, particularly near the PMT faces.

This would reduce total internal reflections in the AC and lower the optical efficiency. Small air bubbles between the PMT faces and the AC (inefficient coupling) would also lower the optical efficiency¹⁵. These two effects could account for the 5.5% discrepancy I measure. Though I only dismantled the PMT assembly between collecting the LS and TeLS data (not between each individual measurement) to refill the AC, I take the 5.5% discrepancy to be an estimate of the systematic uncertainty on the light output for each data set.

I also investigated the dependence of the measured gamma-ray energy on the measurement geometry. Depending on the geometry, gamma-rays may interact with the HPGe closer or farther from the crystal centre. Because the crystal is large (5 in), interactions near the edge may suffer from incomplete charge collection, appearing lower in energy than equivalent depositions near the crystal centre. Using 661.657 keV gamma-rays from the ¹³⁷Cs source, I measured the full energy deposition for interactions near the crystal centre and for interactions near the crystal edge by masking these regions of the HPGe end cap with 2 in thick lead bricks. Between these two independent measurements, I found a 0.3% shift in the mean full energy deposition, which I take to be a systematic uncertainty on E_{e^-} , though it is insignificant compared to the 5.5% systematic error on the measured light output.

Using the data processing techniques in section 5.3.2, I create a distribution of average total number of PE versus E_{e^-} for candidate single Compton scatter full energy deposition events for each data set. I then merge the data sets together to create one distribution spanning the entire electron energy range for each of the LS and TeLS samples. Figure 5.17 shows these distributions.

¹⁵Each time I dismantled the PMT assembly, I noted small areas where the PMT was not coupled to the AC, and roughly estimated the area to be about 5% of the PMT face.

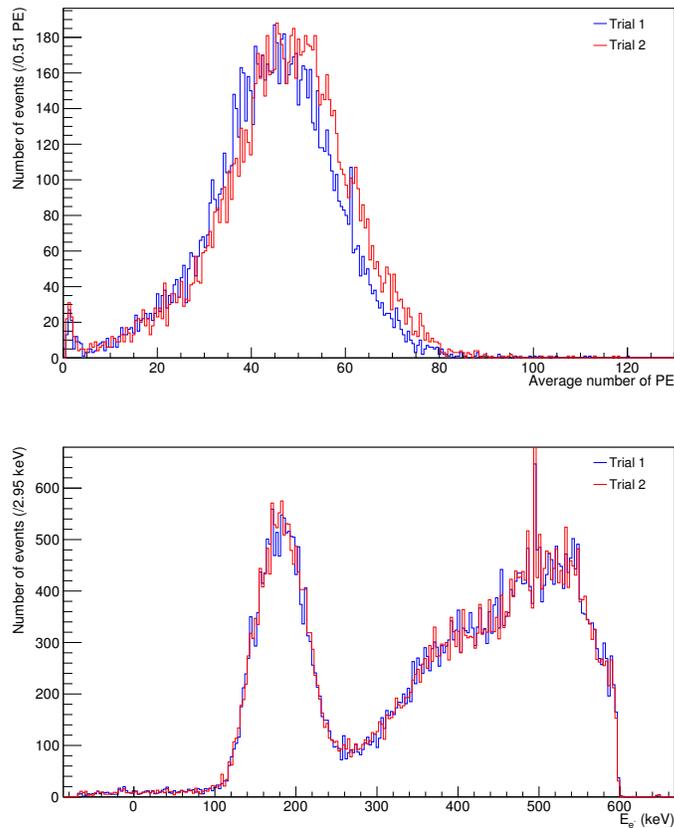


Figure 5.16: The distributions of the average total number of PE observed (top) for two trial assemblies in the same geometric configuration show that the optical efficiency of the PMT assembly may change by approximately 5% as a result of the assembly process. The distribution of E_{e^-} (bottom) is essentially identical for each trial, implying the geometric configurations were approximately the same and the change in the average total number of PE was, in fact, due to the assembly process.

To determine how the LS response varies with E_{e^-} , I determine the mean LS response (in terms of the average total number of PE) in each energy bin. Fitting the PE distribution with a Poisson function, I define the region $\pm 3\sigma$ about the Poisson mean to contain true events, where a majority of the background (in particular, the background near 1 PE that I attribute to PMT noise pulses) falls outside this region. I then take the mean of the data in that

region and use the standard deviation of the mean as the statistical error on the mean LS response. The error in E_{e^-} is equal to the energy resolution at the energy $E_{\gamma'} = E_{\gamma} - E_{e^-}$. I add the 5.5% and 0.3% systematic errors to these in quadrature. The mean LS response as a function of energy overlays the processed data in figure 5.17.

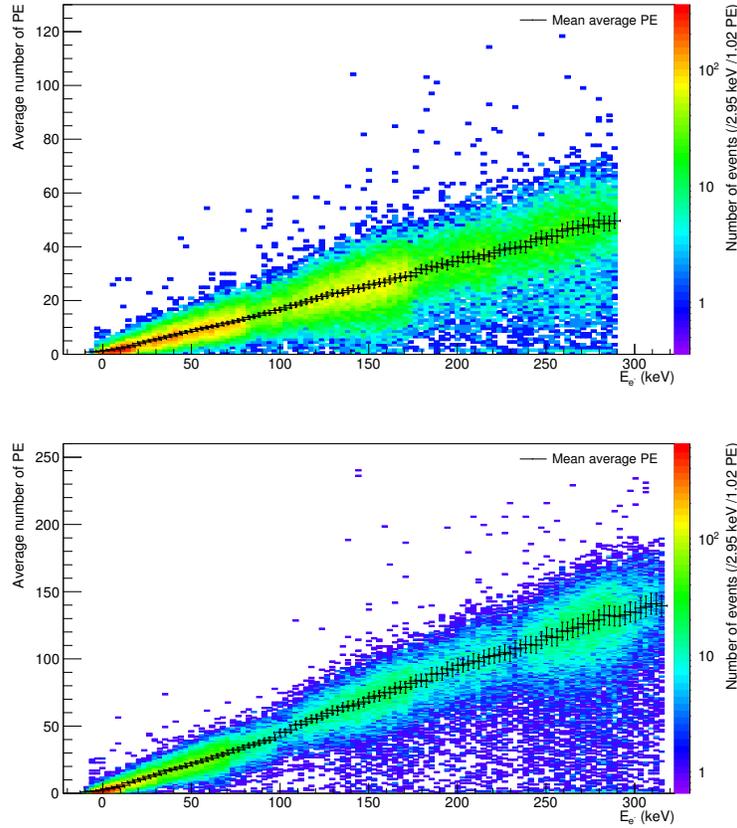


Figure 5.17: The processed complete data set for the LS sample (top) and TeLS sample (bottom) are overlaid with the mean response (in terms of the average total number of PE). Error bars show statistical and systematic errors added in quadrature.

5.4.1 Monte Carlo simulation

There are two different optical efficiencies that affect the measurement of the LS response [134]. The first is the efficiency with which a scintillation photon created in the LS volume reaches the photocathode of the PMT, which I label ξ . The second is the efficiency of creating a SPE at the photocathode given an incident photon, ϵ . Each PMT has a characteristic ξ and ϵ , so there are four efficiencies in total (ξ_A , ξ_B , ϵ_A and ϵ_B). While quenching is only coupled to ξ , in practice, ξ and ϵ are difficult to decouple and measure independently. Also, because the processes associated with each efficiency are binomially distributed, the statistics governing photon counting in a two-PMT system quickly become algebraically complex when more than a handful of photons are produced in the LS. Also, the two efficiencies cannot be coupled into a single efficiency unless the number of photons produced is large so we may approximate distributions as Poisson [134]. Because I am concerned with counting very few photons to assess the level of low-energy electron quenching, I performed a MC simulation of the photon counting statistics, which naturally incorporates this complexity.

I perform 1000 MC experiments at each energy E_{e^-} for which data has been collected (the centre of each energy bin in figure 5.17). To simulate the response of the HPGe in determining E_{e^-} , I sample an energy E from a Gaussian distribution with mean E_{e^-} and width equal to the HPGe energy resolution in equation 5.6 at $E_\gamma - E_{e^-}$. Next, I calculate the number of photons I expect to be generated at E by integrating equation 5.2 as

$$L(E) = \int_0^{E_s} \frac{L_0}{1 + kBdE/dx} dE . \quad (5.25)$$

I estimate the light output L_0 by fitting the data in figure 5.17 with a straight line from 80–180 keV and using the resulting slope as a measure of L_0 . This energy region is where we expect the dE/dx to be small enough for equation 5.1 to be valid. Because L_0 is intimately coupled with the values of ξ and ϵ , I make a simplified estimate of L_0 to be the fitted slope divided by the factor $(\xi_a\xi_B + \epsilon_A\epsilon_B)$. I correct for this simplification later (see below). I randomly sample the number of photons to generate at E for each MC event from a Poisson distribution with mean $L(E)$.

Integrating over dE/dx is non-trivial, as dE/dx is itself a complicated function of energy. I obtain values for dE/dx from the NIST ESTAR database [140], calculated for the specific chemical compositions of the LS and TeLS. Because these values are only available down to 1 keV, I perform a linear extrapolation down to 0 keV using the values of dE/dx at 1 keV and 1.5 keV¹⁶. Figure 5.18 shows the calculated $L(E)$ function.

For each sample photon created in each MC event, I first determine which, if either, PMT face it reaches. Upon reaching a PMT face, I then sample the efficiency for creating a SPE. If a SPE is created in a PMT, I then sample the SPE pulse height distribution for that PMT, governed by equation 5.12 and shown in figure 5.10, but where I have removed the component that describes the PMT pedestal. I do this for each photon created in the MC event, and sum the sampled pulse heights for each to produce a simulated PMT response. I assume efficiencies of $\xi_A = 0.45$, $\xi_B = 0.45$, $\epsilon_A = 0.25$ and $\epsilon_B = 0.25$. This implies an overall optical efficiency of 90% for the PMT assembly (respecting the systematic error of 5.5% on that value). The ϵ values are approximately

¹⁶Alternatives to this could include holding dE/dx constant between 1 keV and 0 keV, or to include an integration constant in equation 5.25.

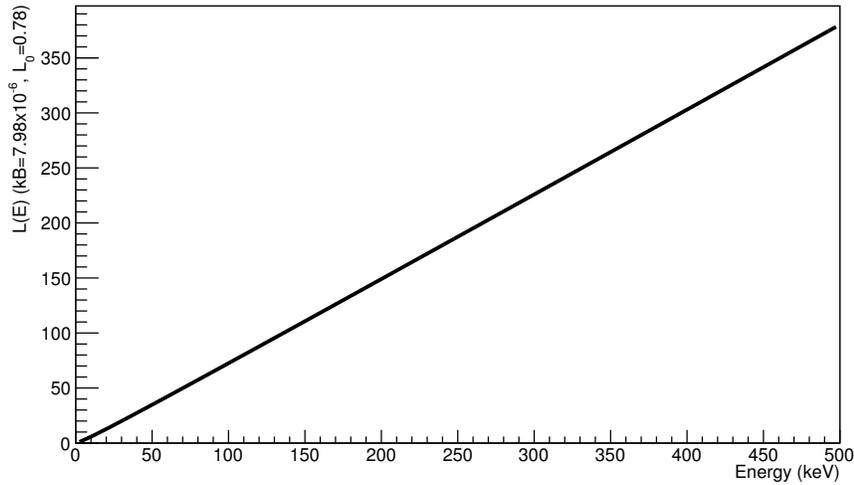


Figure 5.18: Using data for dE/dx from [140], I calculate the light output $L(E)$ using equation 5.25. I estimate L_0 from the data in figure 5.17, giving 0.78. Here, kB is set to the nominal value of 7.98×10^{-6} cm/keV.

the quantum efficiency of the PMTs near the peak of the LS scintillation spectrum. Any error in these assumptions is absorbed into the value of L_0 .

Figure 5.19 shows the resulting MC data for the LS, assuming a nominal value of $kB = 7.98 \times 10^{-6}$ cm/keV and $L_0 = 0.78$. I also compute the mean MC LS response using the same method as I use for the data, which figure 5.19 shows overlaid. Taking a ratio of the data to the MC data reveals that the mean LS response in the data is systematically higher than the MC in the region 200–275 keV. To correct for this, I fit a flat line to the ratio in that energy region, which, in figure 5.19, suggests L_0 should be scaled by a factor of 1.23. I then repeat the MC including this scaling factor.

I created MC data sets for several values of kB ¹⁷ between 0 cm/keV and the nominal value of 7.98×10^{-6} cm/keV, which is the value previously measured

¹⁷Changes in L_0 change the slope of the linear portion of the $L(E)$ curve, whereas changes in kB affect the non-linearity of the curve at low E . It is this non-linearity I am interested in measuring, so I scan over kB , whereas I determine L_0 beforehand.

kB (7.98×10^{-6} cm/keV)	L_0 scaling factor	
	LS	TeLS
0	1.17830	1.08182
0.001	1.17896	1.08256
0.002	1.18212	1.08303
0.005	1.17966	1.08260
0.008	1.18031	1.08193
0.010	1.18050	1.08197
0.020	1.18154	1.08457
0.050	1.18082	1.08535
0.100	1.18805	1.08748
0.150	1.19243	1.09112
0.200	1.19173	1.09352
0.250	1.19339	1.09604
0.300	1.19794	1.09850
0.400	1.20175	1.10382
0.500	1.21127	1.10804
0.800	1.22442	1.12235
1.000	1.23411	1.13228

Table 5.3: I determined these scaling factors by fitting a constant to the ratio of the mean LS response in data to MC in the energy region 200–275 keV. I then repeat the MC simulation with L_0 multiplied by these factors.

for the SNO+ LS [121]. Table 5.4.1 contains the L_0 scaling factors for the LS and TeLS samples for each value of kB . To assess the level of quenching the LS and TeLS data exhibit, I determine for which value of kB the difference in the response between data and MC is the most consistent with 0. Figures 5.20 and 5.21 show the difference between the data and MC mean LS and TeLS response, respectively, as functions of E_{e^-} . I calculate a χ^2 value for the difference in the energy region 5.8–251.7 keV for LS (83 degrees of freedom (DOF)) and 2.3–248.2 keV for TeLS (84 DOF), relative to a value of 0. The value of kB for which the data is most consistent with the MC is that at which the χ^2 is minimal. Using the above-mentioned ranges to calculate the χ^2 leads to $kB=0$ cm/keV for both samples.

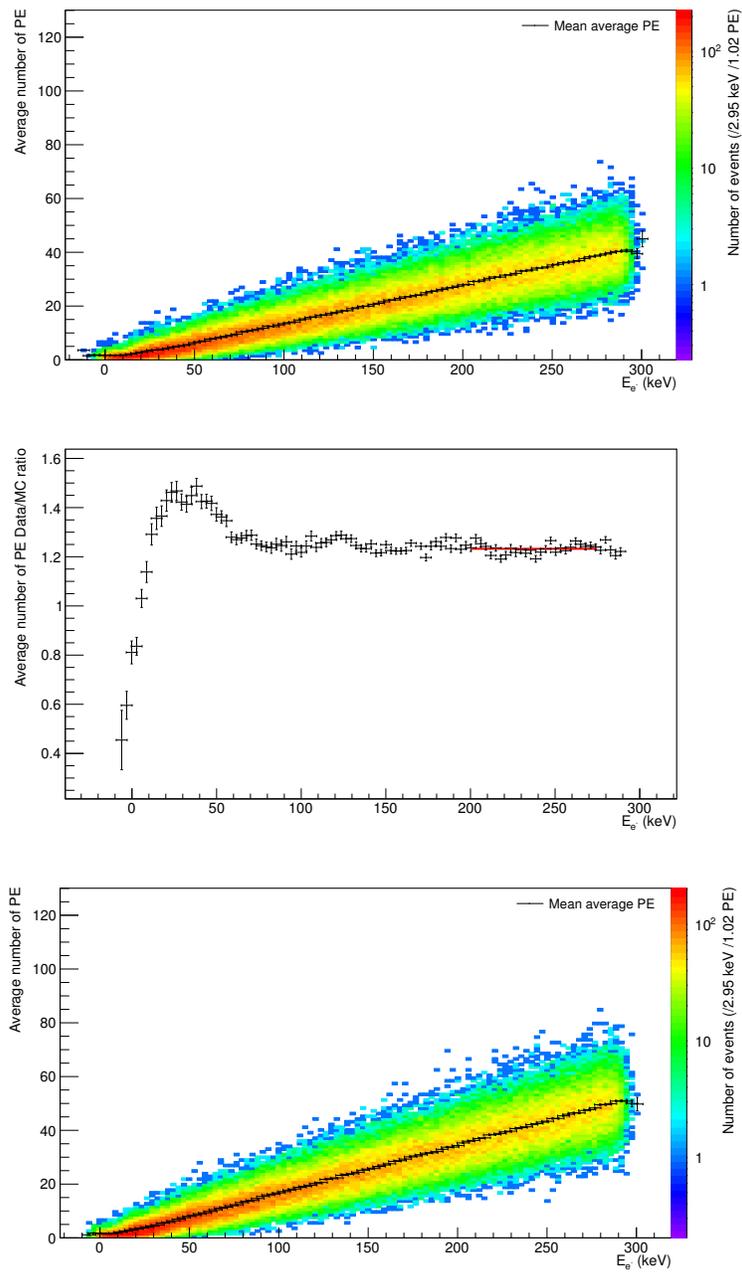


Figure 5.19: The MC data for LS (top) assumes a nominal value of $kB = 7.98 \times 10^{-6}$ cm/keV and $L_0 = 0.78$. Taking the ratio of the data to this MC data (middle) reveals that a scaling factor of 1.23 (red line) for L_0 is necessary to have the MC agree with the data in the energy range 200–275 keV. Incorporating that value into the MC results in scaled MC data (bottom) that better matches the data in figure 5.17.

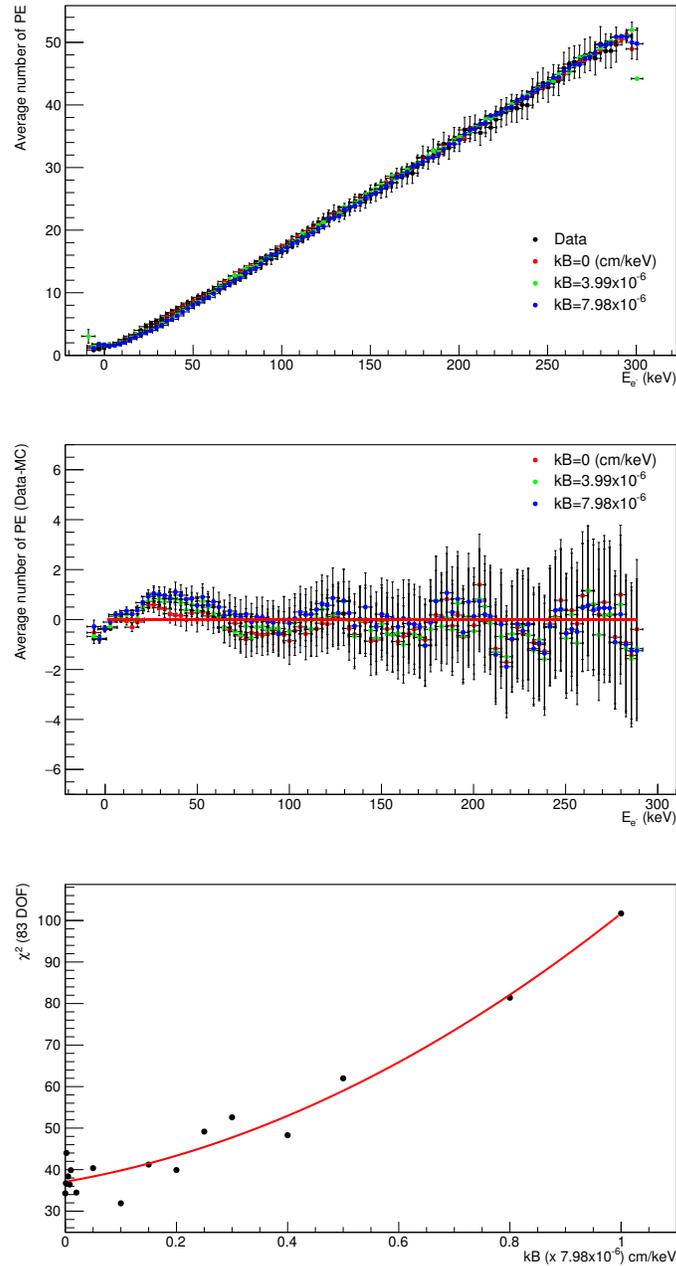


Figure 5.20: The LS data are compared to MC data that assume values of kB of 0 (red), 0.5 (green) or 1 (blue) times the nominal value of $kB = 7.98 \times 10^{-6}$ cm/keV. The top plot shows the mean response of each, while the centre plot shows the difference between data and MC of the mean LS response. The bottom plot shows the calculated χ^2 of the difference between data and MC relative to 0, as a function of kB , in the energy range 5.8–251.7 keV. The red line is a quadratic fit to the χ^2 data.

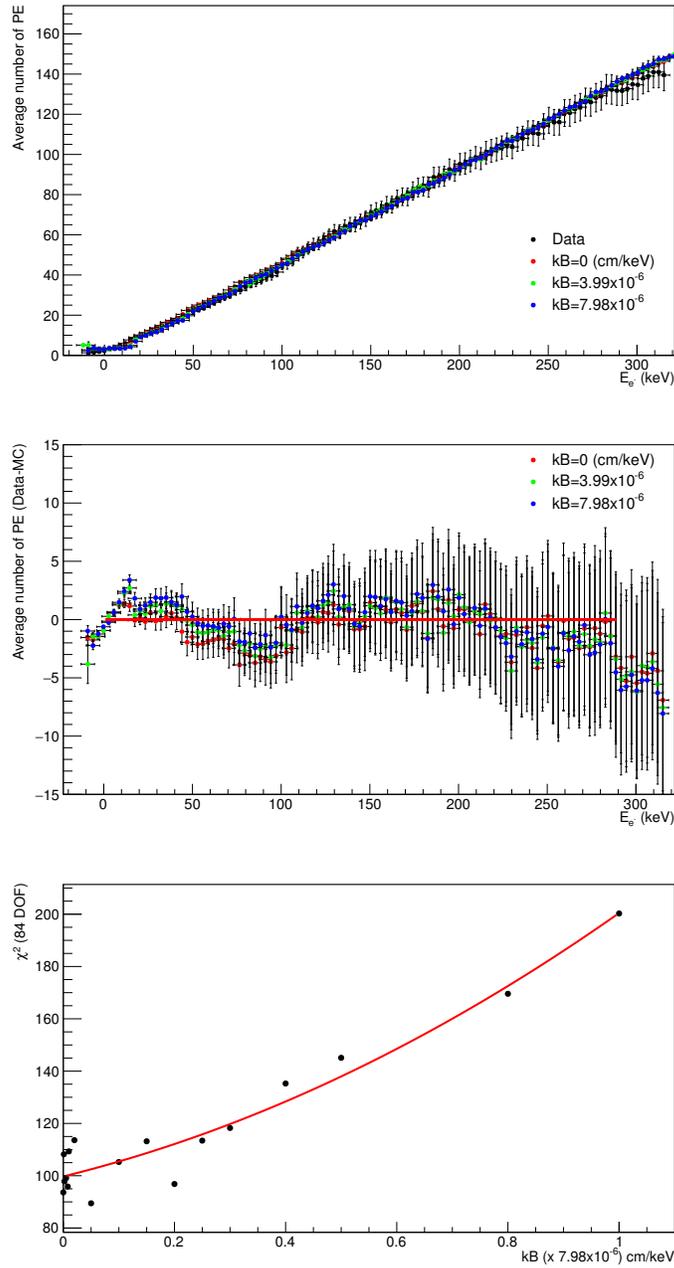


Figure 5.21: The TeLS data are compared to MC data that assume values of kB of 0 (red), 0.5 (green) or 1 (blue) times the nominal value of $kB = 7.98 \times 10^{-6}$ cm/keV. The top plot shows the mean response of each, while the centre plot shows the difference between data and MC of the mean TeLS response. The bottom plot shows the calculated χ^2 of the difference between data and MC relative to 0, as a function of kB , in the energy range 2.3–248.2 keV. The red line is a quadratic fit to the χ^2 data.

5.5 Results

Referring to the bottom plots in figures 5.20 and 5.21, the shape of a quadratic function fitted to the computed χ^2 as a function of kB indicates that the MC is most consistent with the data when $kB = 0$ cm/keV for both the LS and TeLS samples. Nevertheless, figure 5.22 shows that the fitted minimum of the χ^2 actually depends strongly on the energy range used for its computation.

The centre plots in figures 5.20 and 5.21, which show the difference between data and MC of the average PE response, indicate there are strong features at low energies in both sets of data. Below approximately 10 keV, the average number of PE in the data drops below 1 (top plots in figures 5.20 and 5.21), whereas the MC sets the average PE here above 1. This is due to contamination in the data of small noise pulses that the data processing algorithm did not remove. This noise, which typically sits below one PE in height, would tend to pull the average PE distribution to lower values. I explicitly leave out the noise peak in the MC generation. Also, in the LS sample, the data set that covers the energy region 20–40 keV is systematically above the MC, while in the TeLS sample, the data set that covers the energy region 50–100 keV is systematically below the MC. The 5.5% systematic uncertainty on the LS response (see section 5.4) conservatively accounts for these shifts.

To assess the systematic error in the value of kB resulting from this dependence on the χ^2 computation range, I determine the value of kB using 400 randomly selected energy ranges. Using a flat distribution, I randomly sample the lower bound of the energy range between 0–100 keV and the upper bound between 150–290 keV. This forces the χ^2 computation range to be at least 50 keV wide, which corresponds to approximately five data points.

Figure 5.22 shows the distributions of the values of kB I obtain using the method of section 5.4.1 for both the LS and TeLS samples. I take the mean of the distribution to be the best measured value of kB and the root mean squared (RMS) of the distribution to be the systematic error in the value of kB . For the LS sample, these values are 0.41 and 0.35 times the nominal value $kB = 7.98 \times 10^{-6}$ cm/keV, respectively. Similarly for the TeLS sample, these values are 0.51 and 0.33 times the nominal value of kB , respectively.

To obtain the statistical error in the value of kB , I determine, for each χ^2 computation range, the value of kB for which $\chi_{\min}^2 \rightarrow \chi_{\min}^2 + 1$. I take the average of the distributions of these statistical errors, which figure 5.22 also shows, to be the statistical error on the best measured value of kB for each sample. For the LS sample, this results in a statistical error of 0.17 times the nominal value $kB = 7.98 \times 10^{-6}$ cm/keV. For the TeLS sample, the statistical error is 0.13 times the nominal value of kB .

I determine the total error by taking the quadrature sum of the statistical and systematic errors. This results in values of $kB = (3.3 \pm 3.1) \times 10^{-6}$ cm/keV for the LS sample and $kB = (4.1 \pm 2.9) \times 10^{-6}$ cm/keV for the TeLS sample. Table 5.5 summarizes the measured values of Birks' constant for protons, alphas and electrons in the SNO+ LS and TeLS, and includes this measurement of kB for low-energy electrons in LS and TeLS.

Particle	kB ($\times 10^{-6}$ cm/keV)	
	LS	TeLS
p	9.7 ± 0.3	
α	7.6 ± 0.3	7.0 ± 0.4
β	7.98	
β	3.3 ± 3.1	4.1 ± 2.9

Table 5.4: Birks' constant (kB) for p [122], α [123, 67] and β [121] particles for the SNO+ LS and TeLS. The last row reports the results of this work.

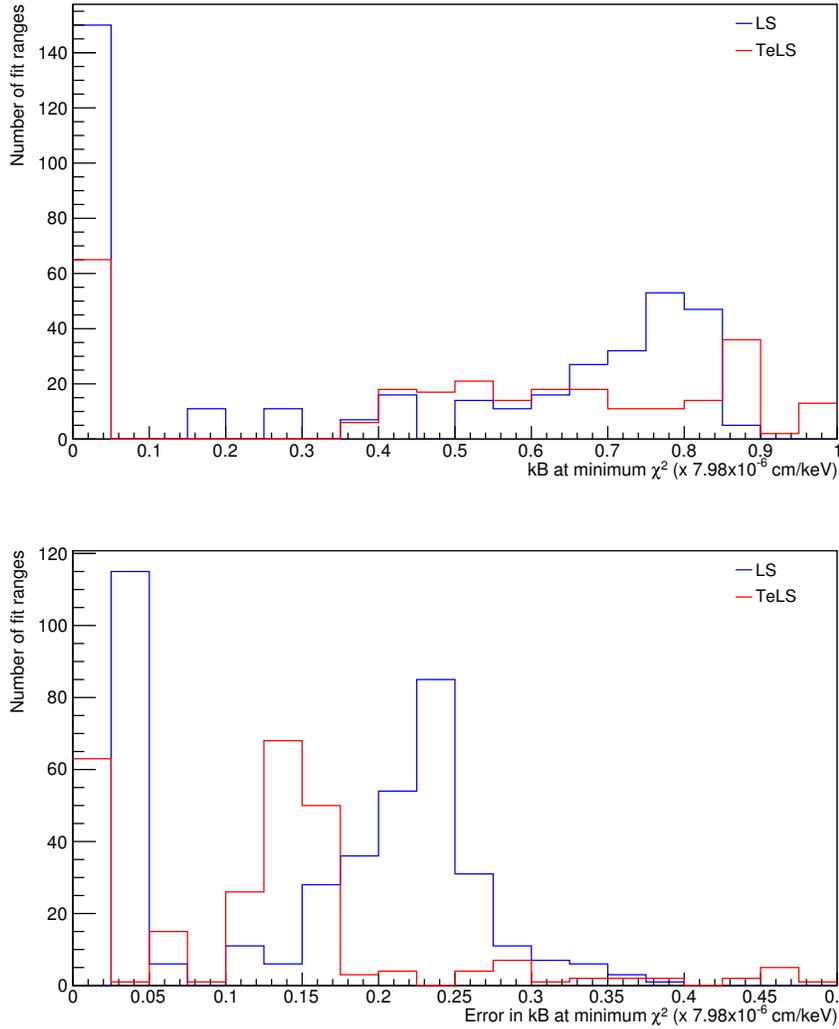


Figure 5.22: Determining kB using 400 randomly selected energy ranges beginning between 0–100 keV and ending between 150–290 keV leads to a mean kB for the LS (blue) and TeLS samples (red) of 0.41 and 0.51 times the nominal value $kB = 7.98 \times 10^{-6}$ cm/keV, respectively (top). Using the RMS of these distributions as the systematic uncertainty in kB leads to errors of 0.35 and 0.33 times the nominal kB value for the LS and TeLS samples. The statistical error on the mean value of kB is the mean value of the statistical error from each of the 400 computations (bottom). The statistical errors for the LS (blue) and TeLS (red) samples are 0.17 and 0.13 times the nominal kB value, respectively.

Conclusion

Jump at da sun!

- Lucy Hurston, *Words of encouragement to her daughter,*
Zora Neale Hurston

For SNO+ to make a meaningful measurement of $0\nu\beta\beta$ using a tellurium-loaded liquid scintillator, it is critical to understand the liquid scintillator response to various radiations. This work presents the first measurement of low-energy electron ionization quenching in the SNO+ tellurium-loaded, LAB-based liquid scintillator. Comparing data taken using a Compton coincidence technique with ^{137}Cs 661.657 keV gamma-rays to a Monte Carlo simulation of the light output statistics, I find the level of quenching of the scintillation light produced by low-energy electrons is small. I measured Birks' constant to be $kB = (3.3 \pm 3.1) \times 10^{-6}$ cm/keV for a nominal LAB-based liquid scintillator and $kB = (4.1 \pm 2.9) \times 10^{-6}$ cm/keV for the tellurium-loaded liquid scintillator.

Using Birks' model to calculate the light response, the range of kB this result suggests corresponds to a variability in the possible light response of the SNO+ tellurium-loaded liquid scintillator of 1.2% at an electron energy of 2.5 MeV. I also performed a Monte Carlo simulation of the SNO+ detector response, which indicates a systematic error of this size in the measurement of

the SNO+ energy scale could lead to a systematic error on the $0\nu\beta\beta$ half-life limit SNO+ measures of approximately 1%. Systematic errors in measuring the SNO+ energy resolution could affect the sensitivity in a similar way.

To minimize any misstatement of the energy scale and resolution, I created a ^{60}Co calibration source that will provide a practically background-free set of calibration data and enable SNO+ to measure the energy scale and resolution near the ^{130}Te $2\nu\beta\beta$ endpoint with high precision. I also created a Monte Carlo simulation of the calibration source within the SNO+ simulation framework. Analysis of the simulated SNO+ response to the ^{60}Co calibration source predicts the source will measure a 3.24% energy resolution (1σ half-width) at 2.51 MeV in the nominal SNO+ liquid scintillator.

The precision of this kB measurement may introduce a systematic error on par with other systematic errors that will affect a SNO+ $0\nu\beta\beta$ measurement, though the discovery of ^{130}Te $0\nu\beta\beta$ decay remains limited by statistical fluctuations in small numbers of background and signal events. This quenching measurement also compliments similar measurements of alpha particles and protons in the SNO+ liquid scintillator and tellurium-loaded liquid scintillator. Taken together, these measurements indicate there appears to be small dependence on the detector response to particle type. Knowing this dependence on particle type will enable SNO+ to place the two-electron $0\nu\beta\beta$ signal on the same energy scale as backgrounds that largely consist of alpha particles and gamma-rays. The consequence is improved accuracy in the estimate of the number of events SNO+ expects to observe near the ^{130}Te $2\nu\beta\beta$ endpoint, which will give the experiment confidence in its ability to extract a potential $0\nu\beta\beta$ signal.

Bibliography

- [1] W. Pauli. Letter, addressed to participants of the Tübingen conference on radioactivity. (Available from the CERN document server: <http://cdsweb.cern.ch/record/83282P>), December 1930.
- [2] C.L. Cowan Jr., 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.
- [3] J. Beringer, *et al.* (Particle Data Group). Review of Particle Physics. *Phys. Rev.*, **D86**(010001), 2012.
- [4] Y. Fukuda, *et al.* Measurement of a small atmospheric ν/μ / ν/e ratio. *Phys. Lett.*, **B433**:9–18, 1998.
- [5] Y. Fukuda, *et al.* Study of the atmospheric neutrino flux in the multi-GeV energy range. *Phys. Lett.*, **B436**:33–41, 1998.
- [6] Q. Ahmad, *et al.* Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, **89**(011301), 2002.
- [7] W.N. Cottingham and D.A. Greenwood. *An Introduction to the Standard Model of Particle Physics*. Cambridge University Press, 2nd edition, 2007.
- [8] Z. Xing and S. Zhou. *Neutrinos in Particle Physics, Astronomy and Cosmology*. Advanced Topics in Science and Technology in China. Zhejiang University Press, 2011.
- [9] P. Di Bari. Some notes on neutrino physics and particle cosmology. In *University of Southampton School of Physics and Astronomy*, 2008.
- [10] W. Rodejohann. Neutrino mass models. In *International Neutrino Summer School*, Virginia Tech, Blacksburg, VA, 2012.
- [11] S.M. Davidson and H.E. Logan. Dirac neutrinos from a second Higgs doublet. *Phys. Rev.*, **D80**:095008, 2009.

- [12] E. Majorana. Teoria simmetrica dell'elettrone e del positrone. *Il Nuovo Cim.*, 14(4):171–184, 1937.
- [13] S.M. Bilenky. Neutrino masses, mixing and oscillations. *Proc. R. Soc. Lond.*, **A460**:403–443, 2004.
- [14] K. Zuber. *Neutrino Physics*. CRC Press, 2nd edition, 2012.
- [15] B. Kayser. On the quantum mechanics of neutrino oscillation. *Phys. Rev.*, **D24**(1):110–116, 1981.
- [16] I.J.R. Aitchison and A.J.G. Hey. *Gauge Theories in Particle Physics Volume II: QCD and the Electroweak Theory*. Graduate Student Series in Physics. Institute of Physics Publishing, Philadelphia, 2004.
- [17] G.C. Branco, R. González Felipe, and F.R. Joaquim. Leptonic CP violation. *Rev. Mod. Phys.*, **84**(2):515–565, 2012.
- [18] E.K. Akhmedov. Do charged leptons oscillate? arXiv:0706.1216v3, 2007.
- [19] S. Boyd. Neutrino oscillations. In *PX435 Neutrino Physics lecture notes*, http://www2.warwick.ac.uk/fac/sci/physics/current/teach/module_home/px435/lec_oscillations.pdf, Coventry, UK, September 2012. The University of Warwick.
- [20] T. Schwetz, M. Tortola, and J. Valle. *New J. Phys.*, **13**(063004), 2011.
- [21] F. Dufour and D. Wark. Experimental status of neutrino physics. arXiv:1207.3983v1, 2012.
- [22] T.K. Kuo and J. Pantaleone. Neutrino oscillations in matter. *Rev. Mod. Phys.*, **61**(4):937–979, 1989.
- [23] B. Aharmim, *et al.* Combined analysis of all three phases of solar neutrino data from the Sudbury Neutrino Observatory. *Phys. Rev.*, **C88**:025501, 2013.
- [24] F. Capozzi, *et al.* Neutrino masses and mixings: Status of known and unknown 3ν parameters. *Nucl. Phys.*, **B908**:218–234, 2016.
- [25] J.L. Hewett, *et al.* Fundamental Physics at the Intensity Frontier. arXiv:1205.2671 [hep-ex], 2012.
- [26] S. Dell’Oro, S. Marcocci, M. Viel, and F. Vissani. Neutrinoless double beta decay: 2015 review. *Adv. High En. Phys.*, **2016**(Article ID 2162659), 2016.

- [27] P.A.R. Ade, *et al.* Planck 2015 results. XIII. Cosmological parameters. arXiv:1502.01589v3 [astro-ph.CO], 2015.
- [28] E.W. Otten and C. Weinheimer. Neutrino mass limit from tritium β decay. *Rep. Prog. Phys.*, **71**(086201), 2008.
- [29] B. Monreal. Neutrino Mass. In *International Neutrino Summer School*, Virginia Tech, Blacksburg, VA, 2012.
- [30] Max Planck Institute for Nuclear Physics. THE-TRAP Project. <http://www.mpi-hd.mpg.de/blaum/high-precision-ms/ptms/neutrino-mass.en.html>, Heidelberg, 2012. University of Heidelberg.
- [31] J. Wolf for the KATRIN collaboration. The KATRIN neutrino mass experiment. *Nucl. Instrum. Meth.*, **A623**:442–444, 2010.
- [32] M. Goeppert-Mayer. Double beta-disintegration. *Phys. Rev.*, **48**(6):512–516, 1935.
- [33] S.M. Bilenky and C. Giunti. Neutrinoless double-beta decay. A brief review. arXiv:1203.5250v3 [hep-ph], 2012.
- [34] J. Kotila and F. Iachello. Phase-space factors for double- β decay. *Phys. Rev.*, **C85**(3):034316, 2012.
- [35] S.M. Bilenky. Neutrinoless double beta-decay. arXiv:1001.1946v1 [hep-ph], 2010.
- [36] A. Dueck, W. Rodejohann, and K. Zuber. Neutrinoless Double Beta Decay, the Inverted Hierarchy and Precision Determination of θ_{12} . arXiv:1103.4152v2 [hep-ph], 2011.
- [37] J. Barea, J. Kotila, and F. Iachello. Nuclear matrix elements for double- β decay. *Phys. Rev.*, **C87**(1):014315, 2013.
- [38] H.V. Klapdor-Kleingrothaus, A. Dietz, H.L. Harney, and I.V. Krivosheina. Evidence for neutrinoless double beta decay. *Mod. Phys. Lett.*, **A16**(37):2409–2420, 2001.
- [39] H.V. Klapdor-Kleingrothaus and I.V. Krivosheina. The evidence for the observation of $0\nu\beta\beta$ decay: the identification of $0\nu\beta\beta$ evnets from the full spectra. *Mod. Phys. Lett.*, **A21**(20):1547–1566, 2006.
- [40] S. Biller. A New Technique to Load ^{130}Te in Liquid Scintillator for Neutrinoless Double Beta Decay Experiments. In *Neutrino 2016*, London, UK, July 2016.

- [41] A. Gando, *et al.* Search for Majorana Neutrinos near the Inverted Mass Hierarchy region with KamLAND-Zen. arXiv:1605.02889v1 [hep-ex], 2016.
- [42] M. Agostini, *et al.* Limit on neutrinoless double beta decay of ^{76}Ge by GERDA. *Physics Procedia: 13th International Conference on Topics in Astroparticle and Underground Physics, TAUP 2013*, **61**:828–837, 2015.
- [43] M. Agostini. First results from GERDA Phase II. In *Neutrino 2016*, London, UK, July 2016.
- [44] K. Alfonso, *et al.* Search for Neutrinoless Double-Beta Decay of ^{130}Te with CUORE-0. *Phys. Rev. Lett.*, **115**:102502, 2015.
- [45] The SNO Collaboration. The Sudbury Neutrino Observatory. *Nucl. Instrum. Meth.*, **A449**:177–207, 2000.
- [46] M. Chen. The SNO Liquid Scintillator Project. *Nucl. Phys. Proc. Suppl.*, **B145**:65–68, 2005.
- [47] C. Kraus for the SNO+ collaboration. SNO with liquid scintillator: SNO+. *Prog. Part. Nucl. Phys.*, **57**:150–152, 2006.
- [48] C. Kraus and S.J.M. Peeters. The rich neutrino programme of the SNO+ experiment. *Prog. Part. Nucl. Phys.*, **64**:273–277, 2010.
- [49] S. Andringa, *et al.* Current Status and Future Prospects of the SNO+ Experiment. *Adv. High En. Phys.*, 2016(Article ID 6194250), 2016.
- [50] Nobelprize.org. The Nobel Prize in Physics 2015. http://www.nobelprize.org/nobel_prizes/physics/laureates/2015/. Nobel Media AB 2014, 2015.
- [51] A. Bialek, *et al.* A rope-net support system for the liquid scintillator detector for the SNO+ experiment. *Nucl. Instrum. Meth.*, **A827**:152–160, 2016.
- [52] R. Knapik, R. Bonventre, and T. Shokair. SNO+ Electronic Upgrades. *Nucl. Phys. Proc. Suppl.*, **B229–232**:523, 2012.
- [53] S. Grullon. Light yield and scintillation decay time constants of Te-loaded liquid scintillator for the SNO+ experiment. In *Proceedings of the 26th International Conference on Neutrino Physics and Astrophysics (Neutrino '14)*, Boston, MA, June 2014. Boston University.

- [54] H. O’Keeffe, E. O’Sullivan, and M. Chen. Scintillation decay time and pulse shape discrimination in oxygenated and deoxygenated solutions of linear alkylbenzene for the SNO+ experiment. *Nucl. Instrum. Meth.*, **A640**(1):119–122, 2011.
- [55] M. Yeh, *et al.* A new water-based liquid scintillator and potential applications. *Nucl. Instrum. Meth.*, **A660**(1):51–56, 2011.
- [56] T. Kaptanoglu. TeDiol (0.5%) Optics in LAB+PPO+bisMSB. *Internal SNO+ document DocDB 3516-v3*, 2016.
- [57] R. Ford. A Scintillator Purification Plant and Fluid Handling System for SNO+. *AIP Conf. Pro. (Low Radioactive Techniques 2015 (LRT 2015))*, **1672**:080003, 2015.
- [58] S. Hans, *et al.* Purification of Telluric Acid for SNO+ Neutrinoless Double-beta Decay Search. *Nucl. Instrum. Meth.*, **A795**:132–139, 2015.
- [59] R. Alves, *et al.* The calibration system for the photomultiplier array of the SNO+ experiment. *J. Inst.*, **10**:P03002, 2015.
- [60] Z. Petriw. An Underwater Six-Camera Array for Monitoring and Position Measurements in SNO+. Master’s thesis, University of Alberta, 2012.
- [61] S. Agostinelli, *et al.* Geant4—a simulation toolkit. *Nucl. Instrum. Meth.*, **A506**(3):250–303, 2003.
- [62] J. Allison, *et al.* Geant4 developments and applications. *IEEE Trans. Nucl. Sci.*, **53**(1):270–278, 2006.
- [63] R. Brun and F. Rademakers. ROOT - An Object Oriented Data Analysis Framework. *Nucl. Instrum. Meth.*, **A389**:81–86, 1997.
- [64] National Nuclear Data Center. Information extracted from the nudat 2 database. [http://www.nndc.bnl.gov/nudat2/getdecayscheme.jsp?nucl eus=60NI&dsid=60co bM decay \(1925.28 d\)&unc=nds](http://www.nndc.bnl.gov/nudat2/getdecayscheme.jsp?nucl eus=60NI&dsid=60co bM decay (1925.28 d)&unc=nds).
- [65] S. Chu, L. Ekström, and R. Firestone. The Lund/LBNL Nuclear Data Search V2.0. <http://nucleardata.nuclear.lu.se/toi/>, 1999.
- [66] I. Coulter. *Modelling and reconstruction of events in SNO+ related to future searches for lepton and baryon number violation*. PhD thesis, University of Oxford, 2014.

- [67] B. von Krosigk. *Measurement of proton and α -particle quenching in LAB based scintillators and determination of spectral sensitivities to supernova neutrinos in the SNO+ detector*. PhD thesis, Technischen Universität Dresden, 2015.
- [68] P. Antonioli, R. Fienberg, and F. Fleurot, *et al.* SNEWS: the SuperNova Early Warning System. *New J. Phys.*, **6**(114), 2004.
- [69] A. Friedland, C. Lunardini, and C. Peña Garay. Solar neutrinos as probes of neutrino-matter interactions. *Phys. Lett.*, **B594**:347–354, 2004.
- [70] G. Bellini, J. Benziger, and D. Bick, *et al.* First evidence of pep solar neutrinos by direct detection in Borexino. *Phys. Rev. Lett.*, **108**(5):051302, 2012.
- [71] G. Bellini, *et al.* Final results of Borexino Phase-I on low-energy solar neutrino spectroscopy. *Phys. Rev.*, **D89**:112007, 2014.
- [72] M. Asplund, N. Grevesse, A.J. Sauval, and P. Scott. The Chemical Composition of the Sun. *Annu. Rev. Astro. Astrophys.*, **47**:481–522, 2009.
- [73] W.C. Haxton, R.G. Hamish Robertson, and A.M. Serenelli. Solar Neutrinos: Status and Prospects. *Annu. Rev. Astro. Astrophys.*, **51**:21–61, 2013.
- [74] Borexino Collaboration. Neutrinos from the primary proton-proton fusion process in the Sun. *Nature*, **512**:383–386, 2014.
- [75] A. Gando, *et al.* Constraints on θ_{13} from a three-flavor oscillation analysis of reactor antineutrinos at KamLAND. *Phys. Rev.*, **D83**(5):052002, 2011.
- [76] O. Šrámek, W.F. McDonough, E.S. Kite, V. Lekić, S.T. Dye, and S. Zhong. Geophysical and geochemical constraints on geoneutrino fluxes from Earth’s mantle. *Earth Planet. Sci. Lett.*, **361**:356, 2013.
- [77] M. Agostini, *et al.* Spectroscopy of geoneutrinos from 2056 days of Borexino data. *Phys. Rev.*, **D92**:031101(R), 2015.
- [78] A. Gando, *et al.* Partial radiogenic heat model for Earth revealed by geoneutrino measurements. *Nature Geoscience*, **4**:647–651, 2011.
- [79] C. Alduino, *et al.* Measurement of the Two-Neutrino Double Beta Decay Half-life of ^{130}Te with the CUORE-0 Experiment. arXiv:1609.01666v1 [nucl-ex], 2016.

- [80] S. Biller and N. Fatemi-Ghomi. Optimisation of Detector Parameters for Enriched Nd Loading Scenarios. *Internal SNO+ document DocDB 1153-v1*, 2011.
- [81] M. Yeh. Nd-LS Synthesis and Purification. *Internal SNO+ document DocDB 395-v1*, 2011.
- [82] E. Vazquez-Jauregui. Pile up studies with RAT for 0.1% and 0.3% Nd concentration. *Internal SNO+ document DocDB 1169-v1*, 2011.
- [83] N. Fatemi-Ghomi. Double beta decay plots. *Internal SNO+ document DocDB 1137-v4*, 2011.
- [84] F. James. *Statistical Methods in Experimental Physics*. World Scientific Publishing Co., Singapore, 2nd edition, 2006.
- [85] A.L. Read. Modified frequentist analysis of search results (the CL_s method). *CERN Note*, open-2000-205, 2000.
- [86] A.L. Read. Presentation of search results—the CL_s technique. In L. Lyons and J. Stirling, editors, *Advanced Statistical Techniques in Particle Physics*, www.ippp.dur.ac.uk/Workshops/02/statistics/proceedings/read.ps, Grey College, Durham, 18-22 March, 2002.
- [87] A. Wright. *Robust Signal Extraction Methods and Monte Carlo Sensitivity Studies for the Sudbury Neutrino Observatory and SNO+ Experiments*. PhD thesis, Queen’s University, Kingston, Ontario, 2009.
- [88] H. O’Keeffe and M. Chen. Expected radioactive backgrounds in the SNO+ experiment. *Internal SNO+ document DocDB 507-v2*, 2010.
- [89] J. Argyriades, *et al.* Measurement of the double- β decay half-life of ^{150}Nd and search for neutrinoless decay modes with the NEMO-3 detector. *Phys. Rev.*, **C80**(034501(R)), 2009.
- [90] A. Barabash, P. Huber, A. Nachab, and V. Umatov. Investigation of $\beta\beta$ decay in ^{150}Nd and ^{148}Nd to the excited states of daughter nuclei. *Phys. Rev.*, **C79**(045501), 2009.
- [91] F. Simkovic, G. Pantis, J. Vergados, and A. Faessler. Additional nucleon current contributions to neutrinoless double β decay. *Phys. Rev.*, **C60**(055502), 1999.
- [92] G.D. Orebi Gann, S.J.M. Peeters, and J. Wilson. Calibration & Analysis Plan. *Internal SNO+ document DocDB 1698-v6*, 2013.

- [93] E.L. Chupp, P.J. Lavakare, and A.A. Sarkady. Onboard calibration system for γ -ray spectrometers in space satellites. *IEEE Trans. Nucl. Sci.*, **NS-16**(1):309–313, 1968.
- [94] Saint-Gobain Crystals. BC-400, BC-404, BC-408, BC-412, BC-416 Premium Plastic Scintillators. <http://www.crystals.saint-gobain.com/uploadedFiles/SG-Crystals/Documents/SGC%20BC400-404-408-412-416%20Data%20Sheet.pdf>, June 2014.
- [95] W. Brouwer, W.J. Burris, B. Caron, J. Hewlett, L. Holm, A. Hamilton, W.J. McDonald, J.L. Pinfold, P. Price, J.R. Schaapman, L. Sibley, R.A. Soluk, and L.J. Wampler. The ALTA cosmic ray experiment electronics system. *Nucl. Instrum. Meth.*, **A539**(3):595–605, 2005.
- [96] D. Pocanic. Tips on milling and handling plastic scintillator surfaces. http://pibeta.web.psi.ch/docs/notes/fly_cutter.html, 1999.
- [97] J.D. Stachiw. *Handbook of Acrylics for Submersibles Hyperbaric Chambers and Aquaria*. Best Publishing Company, USA, 2003.
- [98] Saint-Gobain Crystals. Care Instructions Plastic & Scintillating Fiber. http://www.crystals.saint-gobain.com/uploadedFiles/SG-Crystals/Documents/Operation_Handling/Handling%20and%20Cleaning%20Plastic%20Scintillators.pdf, May 2014.
- [99] David B. Sims. Private communication. Customer Service Representative, Saint-Gobain Crystals, October 2012.
- [100] H. Grimes and J.A. Santos. The Structure and Colour of Anhydrous Cobalt Chloride, $CoCl_2$, at Room and very Low Temperatures. *Z. Krist.*, **88**:136, 1934.
- [101] University of Alberta Office of Environmental Health & Safety. *Radiation Safety Course manual*, (9-5), 2003.
- [102] International Organization for Standardization. *ISO 9978:1992 Radiation protection—Sealed radioactive sources—Leakage test methods*. Geneva, Switzerland, 1992.
- [103] American National Standards Institute. *ANSI/HPS N43.6-1997 Sealed Radioactive Sources Classification*. New York, 1997.
- [104] V.V. Krongauz and D. Reddy. Radio-scintillating probe for monitoring molecular transport in polymers. *Pol. Comm.*, **32**(1):7–9, 1991.
- [105] J. Crank and G.S. Park, editors. *Diffusion in Polymers*. Academic Press Inc. (London) Ltd., London, 1968.

- [106] J. Crank. *The Mathematics of Diffusion*. Oxford University Press, London, 2nd edition, 1975.
- [107] Hamamatsu Photonics K.K. Electron Tube Division. Photomultiplier Tube Modules. www.hamamatsu.com/resources/pdf/etd/PMTmodules_TPMO0010E02.pdf, September 2012.
- [108] Saint-Gobain Crystals. Detector Assembly Materials. <http://www.crystals.saint-gobain.com/uploadedFiles/SG-Crystals/Documents/Organic%20Product%20Accessories%20Data%20Sheet.pdf>, June 2014.
- [109] Hamamatsu Photonics K.K. Electron Tube Division. Photomultiplier Tubes and Related Products. www.hamamatsu.com/resources/pdf/etd/PMT_TMPZ0001E01.pdf, November 2010.
- [110] A. Bousselham, H. H. Barrett, V. Bora, and K. Shah. Photoelectron anticorrelations and sub-Poisson statistics in scintillation detectors. *Nucl. Instrum. Meth.*, **A620**(2–3):359–362, 2010.
- [111] V.I. Tretyak and Y.G. Zdesenko. Tables of double beta decay data - an update. *Atom. Dat. Nucl. Dat. Tab.*, **80**(1):83–116, 2002.
- [112] O.A. Ponkratenko, V.I. Tretyak, and Y.G. Zdesenko. Event generator DECAY4 for simulating double-beta processes and decays of radioactive nuclei. *Phys. Atom. Nucl.*, **63**(7):1282–1287, 2000.
- [113] A. Bialek. VertexGen_Decay0 - a new RAT generator. *Internal SNO+ document DocDB 592-v4*, 2010.
- [114] N.B. Gove and M.J. Martin. *Nucl. Data Tables*, pages 205–317, 1971.
- [115] J.J. Matese and W.R. Johnson. Screening Corrections to the Fermi Function for Allowed β Decay. *Phys. Rev.*, **150**(3):846–851, 1966.
- [116] H. Behrens and J. Janecke. Numerical Tables for Beta-Decay and Electron Capture. In H. Schopper, editor, *Numerical Data and Functional Relationships in Science and Technology*, volume 4 of *Landolt-Bornstein New Series*. Springer-Verlag, 1969.
- [117] Joint Institute for Nuclear Research. <http://159.93.76.15/betas/public/>, 2013.
- [118] I. Towner. Private communication, 2010.
- [119] J.R. Wilson. Validation of Radioactive Decay Simulations. *Internal SNO+ document DocDB 2188-v3*, 2015.

- [120] M.A. Preston. *Physics of the Nucleus*. Addison-Wesley Publishing Company, Inc., Massachusetts, 1962.
- [121] H. Wan Chan Tseung, J. Kaspar, and N. Tolich. Measurement of the dependence of the light yields of linear alkylbenzene-based and EJ-301 scintillators on electron energy. *Nucl. Instrum. Meth.*, **A654**:318–323, 2011.
- [122] B. von Krosigk, L. Neumann, R. Nolte, S. Röttger, and K. Zuber. Measurement of the proton light response of various LAB based scintillators and its implication for supernova neutrino detection via neutrino–proton scattering. *Eur. Jour. Phys.*, **C73**:2390, 2013.
- [123] B. von Krosigk, *et al.* Measurement of α -particle quenching in LAB based scintillator in independent small-scale experiments. *Eur. Jour. Phys.*, **C76**(3):109–123, 2015.
- [124] J.B. Birks. *The Theory and Practice of Scintillation Counting*, volume 27 of *International Series of Monographs in Electronics and Instrumentation*. Pergamon Press Ltd., 1964.
- [125] E. Schram. *Organic Scintillation Detectors*. Elsevier Monographs Physics Section. American Elsevier Publishing Company, Inc., New York, 1963.
- [126] Napy1kenobi. Pistates. <https://en.wiki2.org/wiki/File:Pistates.svg#/media/File:Pistates.svg>. Licensed under CC BY-SA 3.0 via Commons.
- [127] J.B. Birks. Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations. *Proc. Phys. Soc.*, **A64**(10), 1951.
- [128] S. Beddar and L. Beaulieu, editors. *Scintillation Dosimetry*. CRC Press, 2016.
- [129] R. Hofstadter and J. McIntyre. Measurement of Gamma-Ray Energies with Two Crystals in Coincidence. *Phys. Rev.*, **78**:619, 1950.
- [130] M.N. Peron and P. Cassette. COCO, a Compton coincidence experiment to study liquid scintillator response in the 1–20 keV energy range. *Nucl. Instrum. Meth.*, **A353**:41–45, 1994.
- [131] M.N. Peron and P. Cassette. Absolute efficiency of LS cocktails using a Compton coincidence method. *Liquid Scintillation Spectrometry 1994: Proceedings of the International Conference on Advances in Liquid Scintillation Spectrometry*, August 1994.

- [132] M.N. Peron and P. Cassette. A Compton coincidence study of liquid scintillator response to low-energy electrons. *Nucl. Instrum. Meth.*, **A369**:344–347, 1996.
- [133] S.V. Bakhlanov, *et al.* A method for measuring the detector response function for monochromatic electrons based on Compton scattering. *Inst. Exp. Tech.*, **59**(3):333–336, 2016.
- [134] R. Broda, K. Pochwalski, and T. Radoszewsky. Calculation of Liquid-scintillation Detector Efficiency. *Appl. Radiat. Isot.*, **39**(2):159–164, 1988.
- [135] Hamamatsu Photonics K.K. Electron Tube Division. Photomultiplier Tube R580. http://www.hamamatsu.com/resources/pdf/etd/R580_TPMH1100E.pdf, June 2014.
- [136] J.H. Hubbell and S.M. Seltzer. Tables of X-Ray Mass Attenuation Coefficients and Mass Energy-Absorption Coefficients from 1 keV to 20 MeV for Elements $Z=1$ to 92 and 48 Additional Substances of Dosimetric Interest. www.nist.gov/pml/data/xraycoef/index.cfm. Radiation Physics Division, PML, NIST, January 13, 2016.
- [137] S.D. Sharples. Linux Drivers for Tektronix Oscilloscopes and Arbitrary Function Generators: Using the VXI-11 RPC Protocol over Ethernet. <http://optics.eee.nottingham.ac.uk/tek/> and <http://optics.eee.nottingham.ac.uk/vxi11/>, May 2007.
- [138] E.H. Bellamy, *et al.* Absolute calibration and monitoring of a spectrometric channel using a photomultiplier. *NIM*, **A339**:468–476, 1994.
- [139] Q. He. Fitting Single Photo-electron peak. <http://www.hep.princeton.edu/~mcdonald/dayabay/He/GainCalib.pdf>, 2010.
- [140] M.J. Berger, J.S. Coursey, M.A. Zucker, and J. Chang. Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions. <http://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html>. Physical Measurement Laboratory, NIST, December 30, 2015.
- [141] N. Fatemi-Ghomi. Double beta decay sensitivity code. Packaged with RAT, 2011.
- [142] L. Moneta. MINUIT2 better than MINUIT? <http://root.cern.ch/phpBB3/viewtopic.php?f=15&t=10211>, 2010.
- [143] Z. Cui. DBD Sensitivity Code Testing and DBD Sensitivity Code Version 0.0 Modification. *Internal SNO+ document DocDB 936-v1*, 2011.

Appendix A

^{150}Nd sensitivity method

A.1 Robustness of the sensitivity method

The code available in RAT to perform the $0\nu\beta\beta$ sensitivity study (the “original” code) [141] did not perform in a robust manner. Approximately 17% of the fits it performed in an ensemble of simulated experiments would fail. This high rate of failure prompted me to explore the original code more deeply, which ultimately led to me rewriting the code in full. In this section, I outline tests I performed to ensure this new version of the code (the “new” code) functioned appropriately. Ultimately, switching the minimizer from Minuit to Minuit2 (with the maximum number of function calls set to 7000 and the tolerance to 1×10^{-8}) alleviated this problem completely¹ [63].

¹The algorithm Minuit2 uses is effectively the same as that of Minuit, so the two minimizers should not technically be behaving differently, particularly as they both always force the calculation of the Hessian matrix in order to calculate the errors on the fit parameters (see [142] for a discussion on the differences between Minuit and Minuit2 in this regard). However, the two algorithms are clearly behaving differently, where Minuit2 is able to handle the error calculations that cause Minuit to fail. The cause for Minuit’s failure seems to lead back to how it is calculating the derivative of the likelihood function, but I was never able to discern the exact source.

The original code used Minuit as the function minimizer, and it often failed when performing the fit. To understand the Minuit results, I investigated the effect of changing the fit tolerance and fit range on the number of failing fits. Table A.1 shows the number of fits that fail with varying tolerance. The number of failing fits decreases with increasing tolerance, only going to zero for very large values of the tolerance (where the default tolerance in Minuit is 1×10^{-6}). It seems that too strict a tolerance was causing the fits to fail. I also varied the fit energy range to see the effect this had on the number of failing fits. The results of this test are in table A.2. As the lower edge of the fit range increases, the number of failing fits decreases. An explanation for this behaviour, put forth in [143], is that the huge disparity in the number of $2\nu\beta\beta+2\nu\beta\beta^*$ events compared to the other backgrounds (a factor of approximately 10000) is the culprit for the failing fits. Increasing the lower edge of the fit range substantially drops the number of $2\nu\beta\beta+2\nu\beta\beta^*$ events in the fit, setting each background on a more equal footing²; however, changing the fit range also caused the upper error in the number of $0\nu\beta\beta$ events to increase, as figure A.1 shows, where the average upper error increased from 16.6 for the fit range 2.5–5.0 MeV, to 18.6 for the fit range 3.0–5.0 MeV. As such, I kept the fit range to be 2.5–5.0 MeV, which [87] also found to be optimal.

To attempt to discriminate between good and failed fits, I manually calculated an equivalent χ^2 for the fit, expecting it to be higher for failed fits than that for good fits. Figure A.2 shows that the χ^2 is, on average, the same for both good and failed fits. This suggests that the value of the fit parameters

²I also noticed that the number of failed fits would drop significantly when scaling the number of events of each background type to be of the same order (and at most one order of magnitude lower) than the number of $2\nu\beta\beta+2\nu\beta\beta^*$ events, as in table 3.1.

Precision	Number of failed fits per year					Average Number
	5	4	3	2	1	
1E-4	166	177	181	147	170	168.2
1E-3	145	172	157	151	128	150.6
1E-2	121	111	95	85	77	97.8
1E-1	65	65	35	27	17	41.8
1E+0	41	44	38	31	4	31.6
1E+1	2	2	0	3	0	1.4
1E+2	0	0	0	0	0	0

Table A.1: The number of failed fits using Minuit, out of 1000 total fits, decreases as the tolerance is increased (a less strict requirement). The default tolerance for Minuit is 1×10^{-6} .

Fit range (MeV)	Number of failed fits
0.0–15.0	943
3.0–15.0	1
2.5–5.0	168
2.8–5.0	49
2.9–5.0	7
3.0–5.0	0

Table A.2: The number of failed fits in Minuit, out of 1000 total fits, varies as the fit energy range varies.

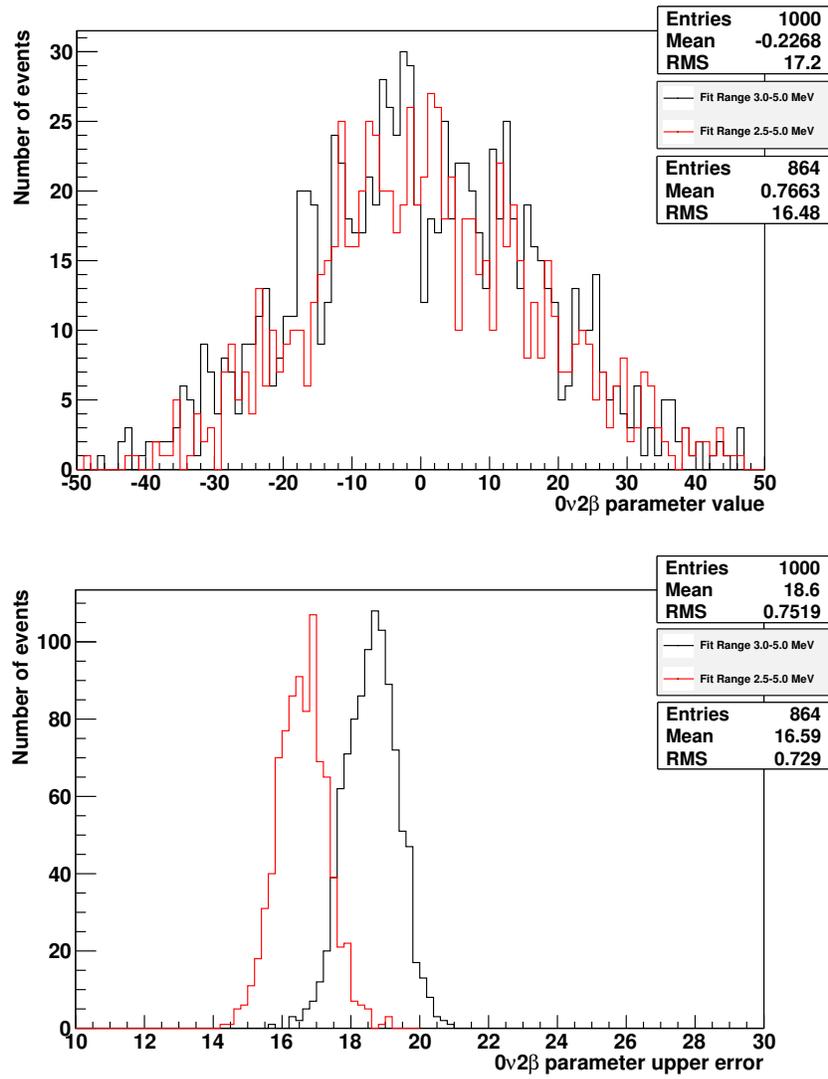


Figure A.1: The upper error (bottom) in the fit number of $0\nu\beta\beta$ events (top) increased significantly when using a narrower fit energy range.

Minuit was returning were not incorrect when the fit failed (as the reduced χ^2 remained close to one), but that it was a problem with the error calculation only. A scan over the χ^2 for each variable, (holding the value of each other variable at the returned fit value), also in figure A.2, shows that the χ^2 functions are all well behaved near their minima. The χ^2 curves vary substantially in width, being very narrow for the $2\nu\beta\beta$ background and very broad for the $0\nu\beta\beta$ signal (because the number of fit events of each type are so different). This lends support to the idea that Minuit may be struggling to find a good step size to use when calculating the derivative of the log likelihood function³, which is required for the error calculation. The conclusion is that failing fits are a result of a problem with calculating the errors on the fit parameters, not because the value of the fit parameters are incorrect (i.e. Minuit was able to find the correct minimum).

Switching to Minuit2, if it fails in its fitting attempt (which is rare), then it returns a fit status that is not 0 (0 indicates a good fit). In this case, the fit is redone with the background parameters fixed to those the failed fit returned and the signal parameter floated (since I expect the fit values to be approximately correct, as per my above argument). This method seems to always lead to a good fit; however, if this fit also fails, then I abandon the fit and redo the experiment (this very rarely happens). The original code uses the value of the upper error itself (later modified to be the value of the global correlation coefficient [143]) to determine if the fit failed by assuming a very low value of the error implied a failed fit. However, when the fit fails as a result of being unable to calculate the upper error correctly, Minuit simply does not return a value for the upper error (or the global correlation coefficient), and

³This plot shows the equivalent reduced χ^2 , not the log likelihood.

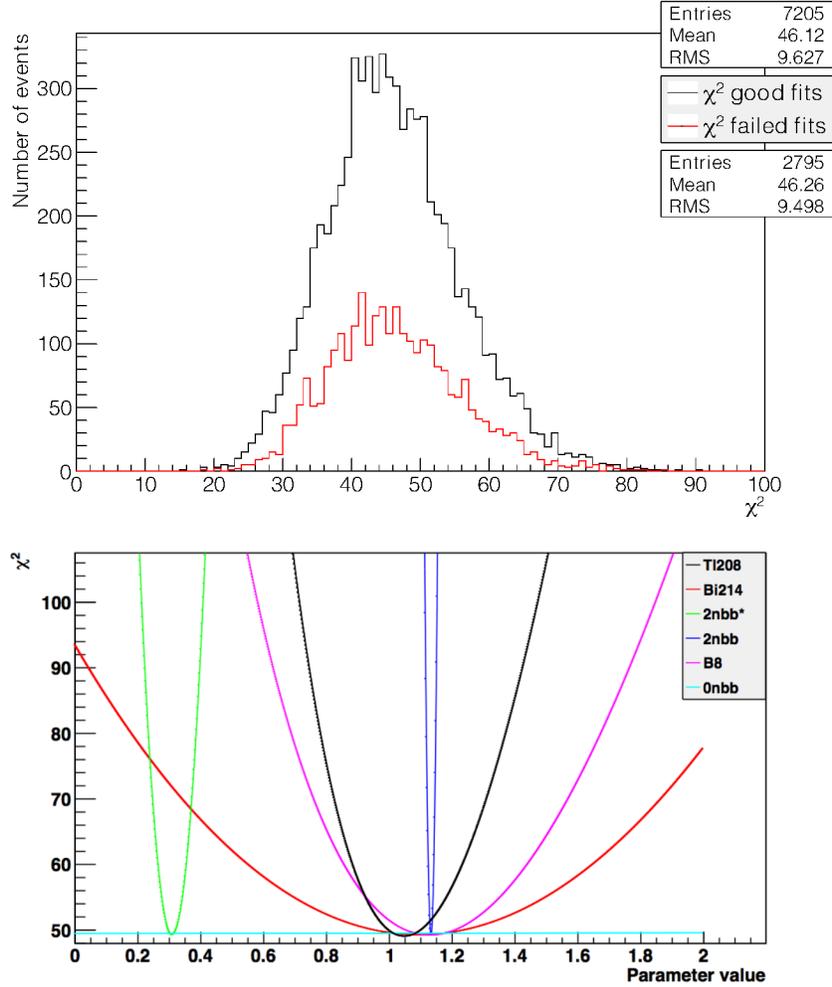


Figure A.2: The equivalent χ^2 (neglecting correlations between parameters) is, on average, the same for good (with a mean of 46.1) and failed (with a mean of 46.3) fits (top). There are 44 degrees of freedom (50 energy bins minus 6 free parameters, where here I treated the $2\nu\beta\beta$ and $2\nu\beta\beta^*$ backgrounds using two distinct parameters). The χ^2 scan (bottom) of a good fit shows that each parameter is nicely behaved near its minimum value.

the original code was discriminating against a value for a variable that did not exist, and that was instead being pulled as a random value. As I already stated, the new code uses the fit status that Minuit2 returns to discern whether a fit has failed.

If the fit returns a fit status equal to 0, indicating a good fit, then I retrieve the upper error on the fit parameters. In the original code, if Minuit returned a negative fit parameter, then the fit was thrown away. This actually tends to bias the fit parameter to a higher value, which figure A.3 shows. In the new code, I keep these negative parameter values, because, while the result may be unphysical, on average it is not, and the sensitivity still lies in a physical region. Also, the original code calculates the 90% confidence limit on the number of observed signal events as the mean number of observed events plus the mean 90% error on the number of observed events (i.e. $\langle \alpha_{0\nu} \rangle + \langle \Delta_{\alpha_{0\nu}} \rangle$). Figure 3.6 shows that the number of observed events and the 90% error on that number are correlated, so the new code calculates the 90% confidence limit on the number of observed events on an event-by-event basis.

A.2 Obtaining 90% confidence limits on ^{150}Nd

$$\tau_{1/2}^{0\nu} \text{ and } m_{eff}^{0\nu}$$

Armed with a Gaussian distribution for the 90% confidence limits on the number of observed $0\nu\beta\beta$ events in an ensemble of experiments, I can predict the shape of the distribution of the 90% confidence limits on the half-life $\tau_{1/2}^{0\nu}$. I obtain this by transforming the random variable $\alpha_{0\nu} + \Delta_{\alpha_{0\nu}}$ (call it x) to the random variable $\tau_{1/2}^{0\nu}$ (call it y) via the continuous transformation given by

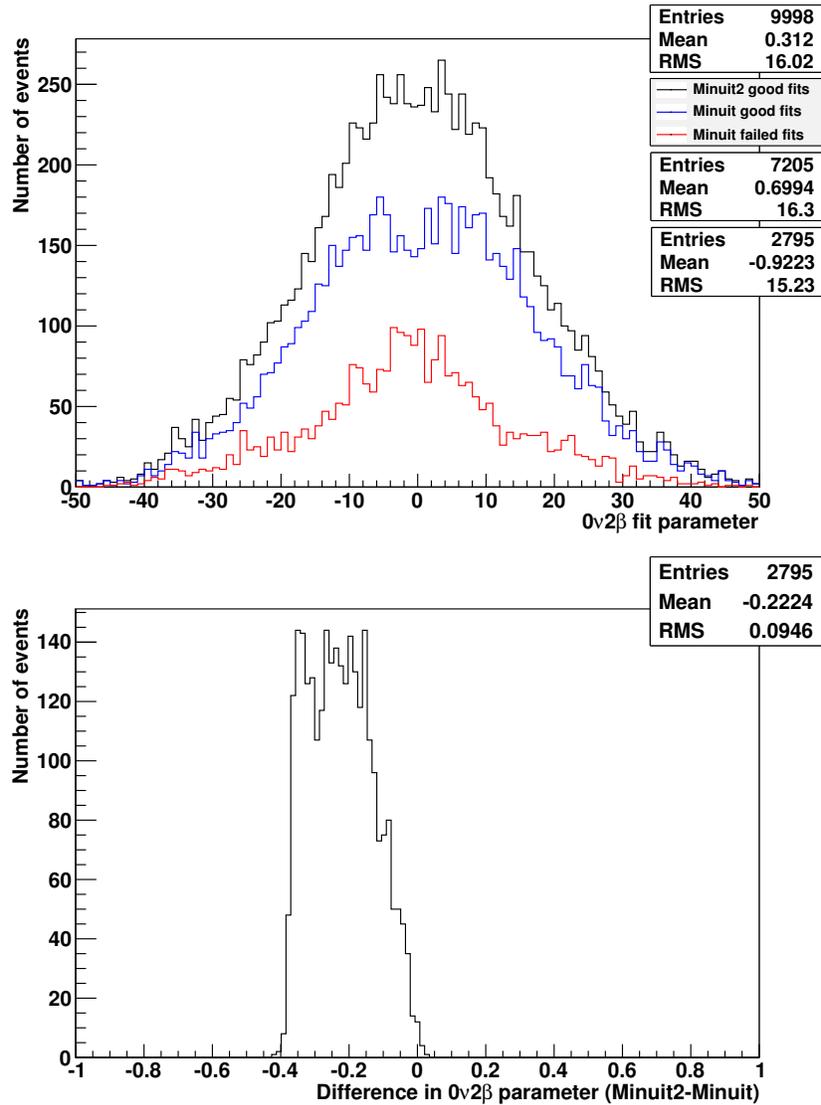


Figure A.3: Throwing away fits that return a negative number of events for any event type (as the original code did) upward-biases the average value of the parameter (top). Using Minuit, the mean $0\nu\beta\beta$ fit parameter for failed fits is -0.92, whereas for good fits it is 0.70. Minuit2 gives a mean value of 0.31. A histogram of the difference in the parameter value between Minuit2 and Minuit in the cases where Minuit failed (bottom) shows that when Minuit fails, Minuit2 actually finds the best value of the parameter to be less negative, on average (although the difference is small compared to the width of the distributions (top)). This difference is essentially 0 using good Minuit fits.

equation 3.4. Thus, given

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-x_o)^2/2\sigma^2}, \quad (\text{A.1})$$

where x_o is the distribution's mean/median (so the mean/median 90% confidence limit of the number of $0\nu\beta\beta$ events) and σ^2 is the variance, if I perform a variable transform

$$y = c/x, \quad (\text{A.2})$$

where c is a constant (about 121.6×10^{24} , as read from equation 3.4), then I may write the distribution of my random variable y as

$$\begin{aligned} g(y) &= |\partial x / \partial y| f(y(x)) \\ &= \frac{c}{y^2} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(c/y-x_o)^2/2\sigma^2}. \end{aligned} \quad (\text{A.3})$$

The extrema of this distribution occur where $dg/dy = 0$, with solutions:

$$\begin{aligned} y_{min} &= 0, \pm \infty \\ y_{max} &= \frac{c}{4\sigma^2} \left(-x_o \pm \sqrt{x_o^2 + 8\sigma^2} \right). \end{aligned} \quad (\text{A.4})$$

If I simplify this and assume that $x_o = 0$ (only as a demonstration), then the non-trivial solution becomes

$$y_{max} = \frac{\pm c}{\sqrt{2\sigma^2}} \quad (\text{A.5})$$

and I obtain

$$\begin{aligned}
g(y = 0) &= 0 \\
g(y = \pm c/\sqrt{2\sigma^2}) &= \sqrt{\frac{2}{\pi}} \frac{\sigma}{ce}, \tag{A.6}
\end{aligned}$$

where the non-trivial solutions are maxima. If I consider now the value of my original random variable at the maxima of $g(y)$, $x = c/y = \pm\sqrt{2\sigma^2}$, I get

$$f(x = \pm\sqrt{2\sigma^2}) = \frac{1}{\sqrt{2\pi\sigma^2}} \frac{1}{e}. \tag{A.7}$$

Thus, the maxima of the transformed function occur at the point where the original function falls to $1/e$ its maximum (at $x_{max} = x_o = 0$). If I do not set $x_o = 0$, so my non-trivial solution remains as in equation A.4, then I find

$$\begin{aligned}
x_{max} &= c/y_{max} \\
&= \frac{-4\sigma^2}{x_o \mp \sqrt{x_o^2 + 8\sigma^2}} \tag{A.8}
\end{aligned}$$

and the value of my original function at x_{max} becomes

$$f(x_{max}) = \frac{1}{\sqrt{2\pi\sigma^2}} \frac{1}{e} e^{-x_o^2/4\sigma^2(1 \mp \sqrt{1+8\sigma^2/x_o^2})}. \tag{A.9}$$

So, it does not correspond to the point where the original function goes to $1/e$ its maximum value and the positive and negative solutions are not symmetric about x_o . In terms of the function $g(y)$, we note that, as $x_o \rightarrow \infty$, the y position of the positive maximum approaches 0 and the height of the peak approaches ∞ , whereas the y position of the negative maximum goes to $-\infty$ with a height that goes to 0. In other words, the function $g(y)$ quickly becomes

asymmetric.

I may also use the cumulative distribution for $g(y)$ to determine which value of y contains a particular percent of the total area of $g(y)$ (i.e. to find the statistical confidence limits). Beginning with the region $y \leq 0$, I have

$$\begin{aligned}
 G(y \leq 0) &= \int_{-\infty}^0 g(y') dy' \\
 &= \int_{-\infty}^0 \frac{c}{y'^2} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(c/y' - x_o)^2/2\sigma^2} dy' \\
 &= \frac{1}{2} - \frac{1}{2} \operatorname{erf} \left(\frac{x_o}{\sqrt{2\sigma^2}} \right). \tag{A.10}
 \end{aligned}$$

So, if $x_o = 0$, then $G(y \leq 0) = 1/2$, which implies that if the median of $f(x)$ is 0, then the median of the transformed distribution $g(y)$ is also 0. If, however, $x_o > 0$, then the median of $g(y)$, call it y_m , occurs where

$$\begin{aligned}
 \left(\int_{-\infty}^0 + \int_0^{y_m} \right) g(y') dy' &= \frac{1}{2} \\
 \Rightarrow \int_0^{y_m} g(y') dy' &= \frac{1}{2} \operatorname{erf} \left(\frac{x_o}{\sqrt{2\sigma^2}} \right). \tag{A.11}
 \end{aligned}$$

Computing this integral, I arrive at

$$\int_0^{y_m} g(y') dy' = \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{x_o - c/y_m}{\sqrt{2\sigma^2}} \right), \tag{A.12}$$

where, if $y_m = \infty$, then

$$\begin{aligned}
 \int_0^{\infty} g(y') dy' &= \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{x_o}{\sqrt{2\sigma^2}} \right) \\
 \Rightarrow \left(\int_{-\infty}^0 + \int_0^{\infty} \right) g(y') dy' &= \frac{1}{2} - \frac{1}{2} \operatorname{erf} \left(\frac{x_o}{\sqrt{2\sigma^2}} \right) + \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{x_o}{\sqrt{2\sigma^2}} \right) \\
 &= 1, \tag{A.13}
 \end{aligned}$$

as expected. Solving equations A.11 and A.12 for y_m leads to

$$1 = \operatorname{erf}\left(\frac{x_o}{\sqrt{2\sigma^2}}\right) - \operatorname{erf}\left(\frac{x_o - c/y_m}{\sqrt{2\sigma^2}}\right). \quad (\text{A.14})$$

Letting $y_m = c/x_o$, the transform of the median x_o under $g(y)$, results in equation A.14 becoming

$$1 = \operatorname{erf}\left(\frac{x_o}{\sqrt{2\sigma^2}}\right), \quad (\text{A.15})$$

which is only true if $x_o = \infty$.

Thus, the median of the distribution $f(x)$ does not, seemingly, transform to the median of the distribution $g(y)$. However, noting that

$$\begin{aligned} \int_0^{c/x_o} g(y') dy' &= \frac{1}{2} + \frac{1}{2} \operatorname{erf}\left(\frac{x_o - \frac{c}{x_o}}{\sqrt{2\sigma^2}}\right) \\ &= \frac{1}{2}, \end{aligned} \quad (\text{A.16})$$

if I consider the median of $g(y)$ to be the point where 1/2 of its area lies in $[0, y_m]$, with the other half of the area in the region $(-\infty, 0) + (y_m, \infty)$, then the median x_o does indeed transform to the median of $g(y)$. This makes sense if I note that the region $x \geq x_o$ of $f(x)$ transforms to region $0 \leq y \leq c/x_o$ of $g(y)$, $0 \leq x < x_o$ transforms to $c/x_o < y < \infty$ and $x < 0$ transforms to $y < 0$. Thus, defining the median of $g(y)$ to be y_m such that

$$\int_0^{y_m} g(y') dy' = \frac{1}{2}, \quad (\text{A.17})$$

implies that the median of $f(x)$ does transform to the median of $g(y)$. To

compute the value of y_l that gives the one-sided limit of the distribution σ_l , I must, therefore, compute

$$\begin{aligned}\sigma_l &= \left(\int_{-\infty}^0 + \int_{y_l}^{\infty} \right) g(y') dy' \\ &= \frac{1}{2} - \frac{1}{2} \operatorname{erf} \left(\frac{x_o - c/y_l}{\sqrt{2}\sigma^2} \right),\end{aligned}\tag{A.18}$$

which is solved numerically for y_l .

Interestingly, because of the area-preserving nature of the distribution, the value of y_l is identical to that which I would obtain if I had used the distribution $f(x)$ to compute the value x_l required to obtain the same limit σ_l , and then transformed x_l to the distribution $g(y)$ via equation A.2. The most significant difference between these two distributions is that $y_{max} \neq c/x_o$, as noted above, where (choosing the positive root solution) $y_{max} < c/x_o$, in general. This means that, if I use the most likely value of x to compute a “most likely” value of y using equation A.2, then I obtain a different value than simply taking the most likely value of y from the distribution $g(y)$, where I define “most likely” to be the value of the argument at the maximum of the distribution.

In terms of a half-life limit measurement, pulling the most likely value of the half-life limit from the distribution $g(y)$ actually leads to a worse limit than computing it from the median of the limit of the number of observed events (the value is lower, and we want to push the lower limit on $\tau_{1/2}^{0\nu}$ as high as we can). Given the distribution $f(x)$, what we see is that small changes in the measured x to the right of the median ($x > x_o$) quickly lead to a measurement of $\tau_{1/2}^{0\nu}$ that approaches 0, whereas small changes in x in the other direction while still remaining above 0 ($0 < x < x_o$), so we remain in the physical region,

cause the value of $\tau_{1/2}^{0\nu}$ to slowly approach ∞ . Thus, it seems the maximum of the distribution $g(y)$ is a good representation of the half-life limit, since small statistical fluctuations of the measured number of events away from the true value of x_o may have a large effect on the measured half-life limit.

Moving on to consider the effective neutrino mass, I may look at the distribution of 90% confidence limits of m_ν^{eff} in a similar way, where, letting m_ν^{eff} be the random variable z , I have

$$\begin{aligned} z &= \pm b/\sqrt{y} \\ &= \pm b/\sqrt{c/x} \equiv \pm \sqrt{x/a} \\ h(z) &= \frac{2a|z|}{\sqrt{2\pi\sigma^2}} e^{-(\pm az^2 - x_o)^2/2\sigma^2}. \end{aligned} \quad (\text{A.19})$$

Here, I have chosen to force the mass to be negative when the 90% confidence limit on the number of observed $0\nu\beta\beta$ decays ($\alpha_{0\nu} + \Delta_{\alpha_{0\nu}}$) and, subsequently, $\tau_{1/2}^{0\nu}$ are negative. Thus, the positive signs correspond to $z \geq 0$ (positive half-lives) and the negative signs to $z < 0$ (negative half-lives). Equations 3.4 and 3.5 supply the constant b (it is approximately 35.73×10^{-2}). This function is extremal at

$$\begin{aligned} z_{min} &= 0, \pm \infty \\ z_{max} &= \pm \sqrt{\frac{\pm x_o + \sqrt{x_o^2 + 2\sigma^2}}{2a}}, \end{aligned} \quad (\text{A.20})$$

with the same sign conventions as above.

As with $g(y)$, I may use the cumulative distribution for $h(z)$ to compute a

limit σ_l on the most likely value of z , z_{max} . Like before, I have

$$\begin{aligned}
H(z \leq 0) &= \int_{-\infty}^0 h(z') dz' \\
&= \int_{-\infty}^0 \frac{2a|z'|}{\sqrt{2\pi\sigma^2}} e^{-(\pm az'^2 - x_o)^2/2\sigma^2} \\
&= \frac{1}{2} - \frac{1}{2} \operatorname{erf} \left(\frac{x_o}{\sqrt{2\sigma^2}} \right). \tag{A.21}
\end{aligned}$$

Similarly,

$$\begin{aligned}
\int_0^{z_m} h(z') dz' &= \frac{1}{2} \operatorname{erf} \left(\frac{x_o}{\sqrt{2\sigma^2}} \right) + \frac{1}{2} \operatorname{erf} \left(\frac{az_m^2 - x_o}{\sqrt{2\sigma^2}} \right) \\
\int_{z_m}^{\infty} h(z') dz' &= \frac{1}{2} - \frac{1}{2} \operatorname{erf} \left(\frac{az_m^2 - x_o}{\sqrt{2\sigma^2}} \right) \\
\Rightarrow \left(\int_{-\infty}^0 + \int_0^{z_m} + \int_{z_m}^{\infty} \right) h(z') dz' &= 1, \tag{A.22}
\end{aligned}$$

as expected, where z_m is the median of the distribution $h(z)$. If $z_m = \sqrt{x_o/a}$, the transformed value of the median of $f(x)$ under equation A.19, then I obtain

$$\int_{z_m}^{\infty} h(z') dz' = \frac{1}{2}, \tag{A.23}$$

showing that the transformed median of $f(x)$ is equal to the median of $h(z)$.

To compute σ_l , I must determine the value z_l such that

$$\begin{aligned}
\int_{-\infty}^{z_l} h(z') dz' &= \sigma_l \\
\Rightarrow \sigma_l &= \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{az_m^2 - x_o}{\sqrt{2\sigma^2}} \right), \tag{A.24}
\end{aligned}$$

which, again, is solved numerically.

Remembering that this defines the limit on the measured effective neutrino

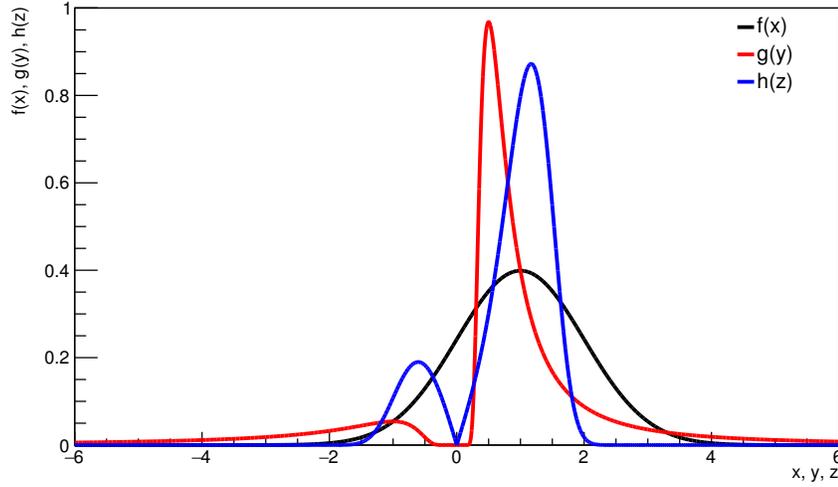


Figure A.4: The Gaussian distribution $f(x)$ is shown along with its transformation under equation A.3, $g(y)$, and equation A.19, $h(z)$. In all three distributions, $x_o = \sigma = a = b = c = 1$. With these parameters, $f(x)$ is maximal at $x_{max} = 1$, $g(y)$ is maximal at $y_{max} = 0.5$, which corresponds to $x = 2$, and $h(z)$ is maximal at $z_{max} = 1.17$, corresponding to $x = 1.37$.

mass, if I choose the most likely value of m_ν^{eff} to be at the maximum of the distribution $h(z)$, then I get a worse limit than computing m_ν^{eff} using the value of x_o and equation A.19 (the former value is higher, and we want to push the effective mass limit as low as possible). However, I choose this as a reasonable representation of the most likely m_ν^{eff} limit as per my above argument for using the value y_m for the most likely limit of $\tau_{1/2}^{0\nu}$.

Figure A.4 shows the functions given by equations A.1, A.3 and A.19, assuming $x_o = \sigma = a = b = c = 1$. Note that the maxima y_{max} and z_{max} , after transforming back to the random variable x , do not equal x_{max} .

Appendix B

^{60}Co calibration source

These appendices contain additional information relevant to the ^{60}Co calibration source, including the engineering drawings of the physical source (appendix B.1), the detailed procedures for the source construction and testing (appendix B.2), and additional information regarding the source simulation in RAT (appendix B.3).

B.1 Source drawings

This section contains the engineering drawings of the ^{60}Co calibration source. These include drawings for the copper encapsulation (figures B.1, B.2, B.3 and B.4), Delrin[®] container (figure B.5) and stem (figure B.6), and the quick connect (figure B.7). Figure B.8 shows the complete assembly of the source.

Note that these drawings may be subject to minor alterations as the design for the stem evolves to accommodate changes in how it is secured to the quick connect. Also, there may be small changes in the event the source isotope changes from ^{60}Co to ^{46}Sc , say.

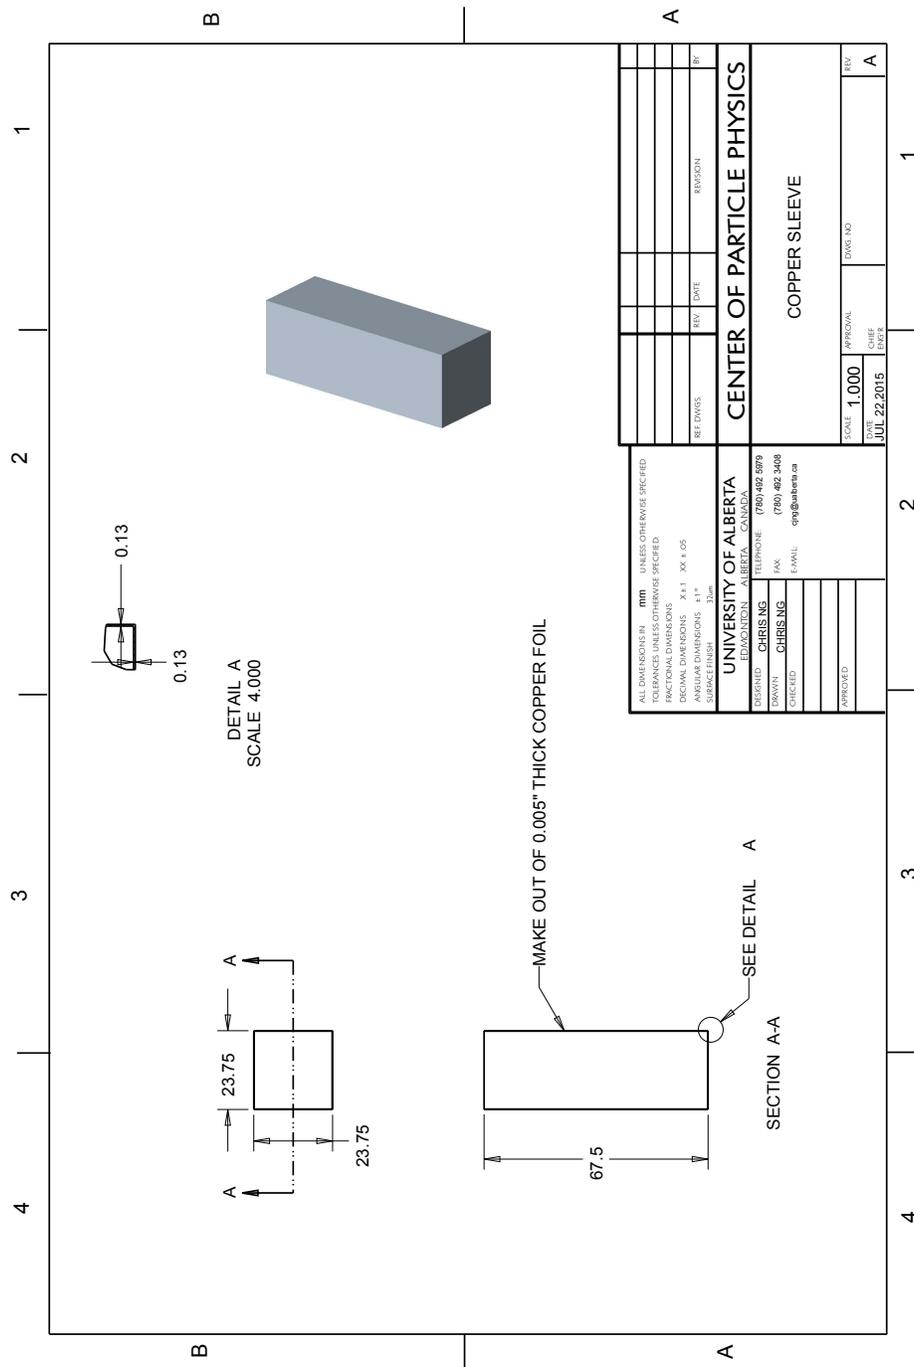


Figure B.1: The 5/1000" copper sleeve is formed by bending a sheet of copper into a rectangle and soldering a small copper overlap along one edge. The bottom is soldered along all four edges, with an overlap to accommodate the soldering on each.

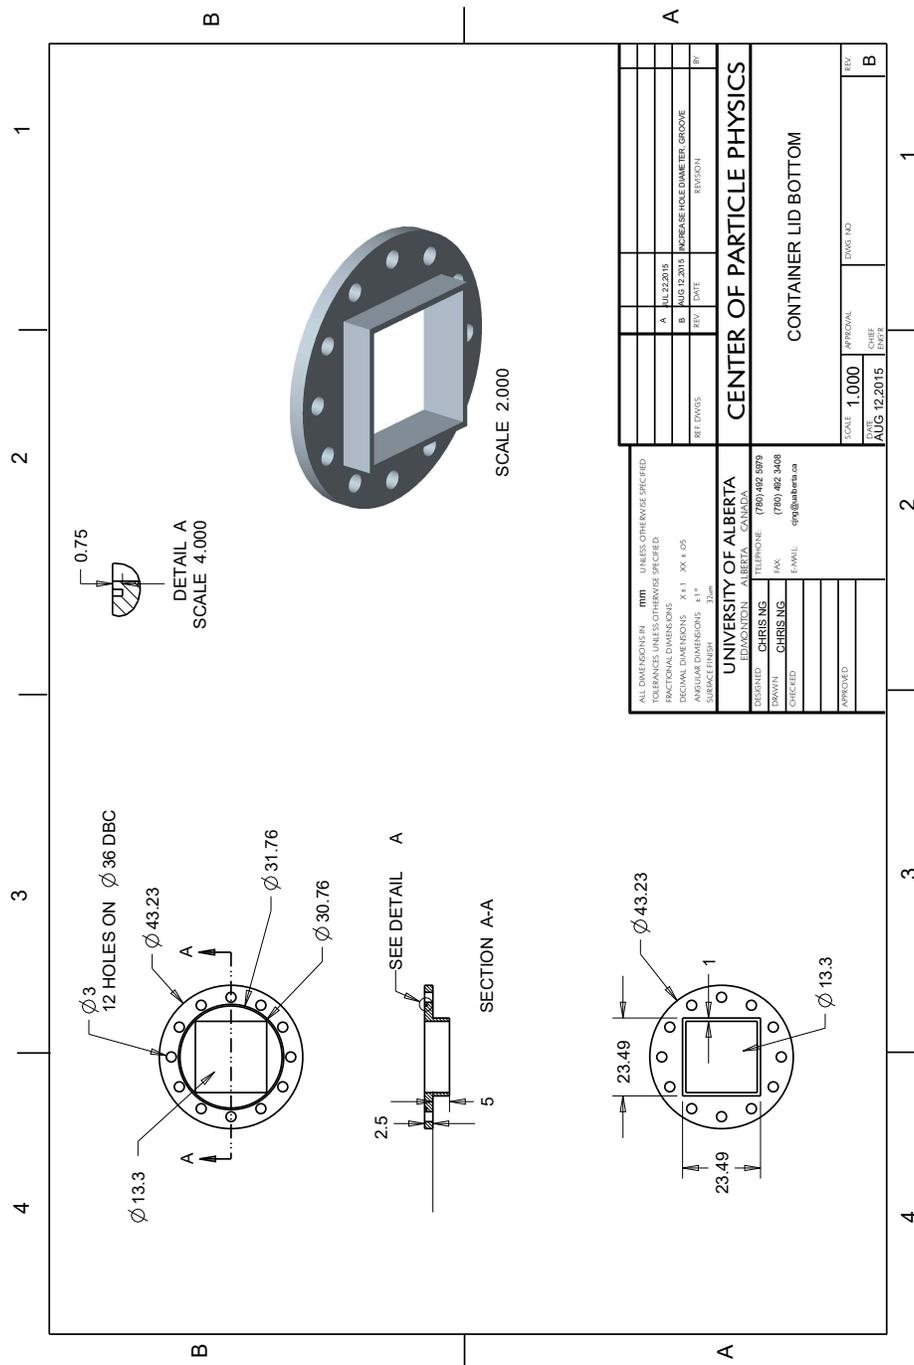


Figure B.2: The bottom section of the copper flange is machined with a lip protruding from the underside onto which the copper sleeve is soldered. The upper side contains a small groove to accommodate the indium wire that forms a helium-tight seal between the two halves.

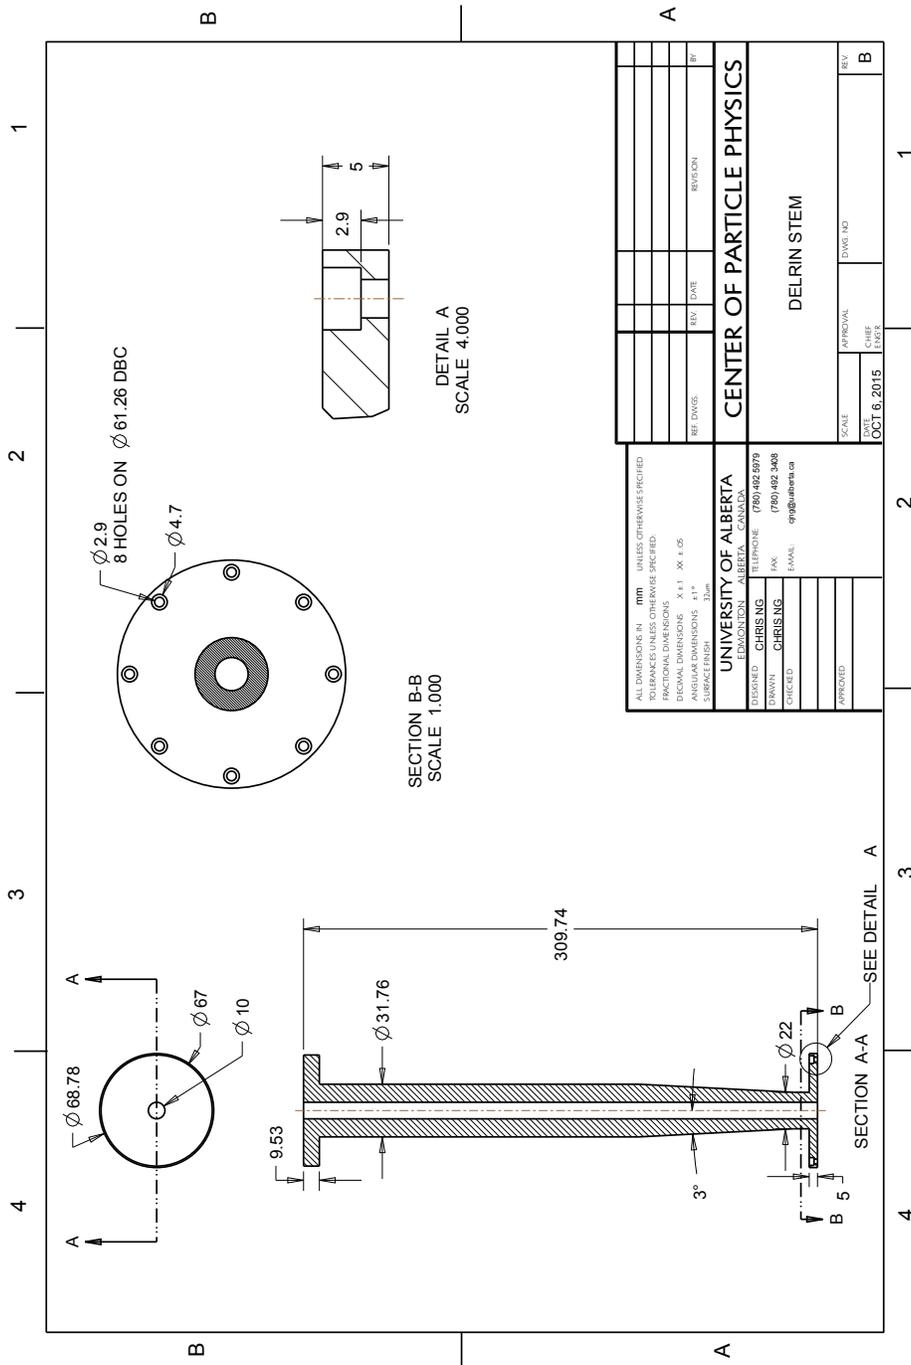


Figure B.6: The stem connects directly to the Delrin® container, acting as the upper flange. A bore hole through the stem enables the passage of PMT wiring.

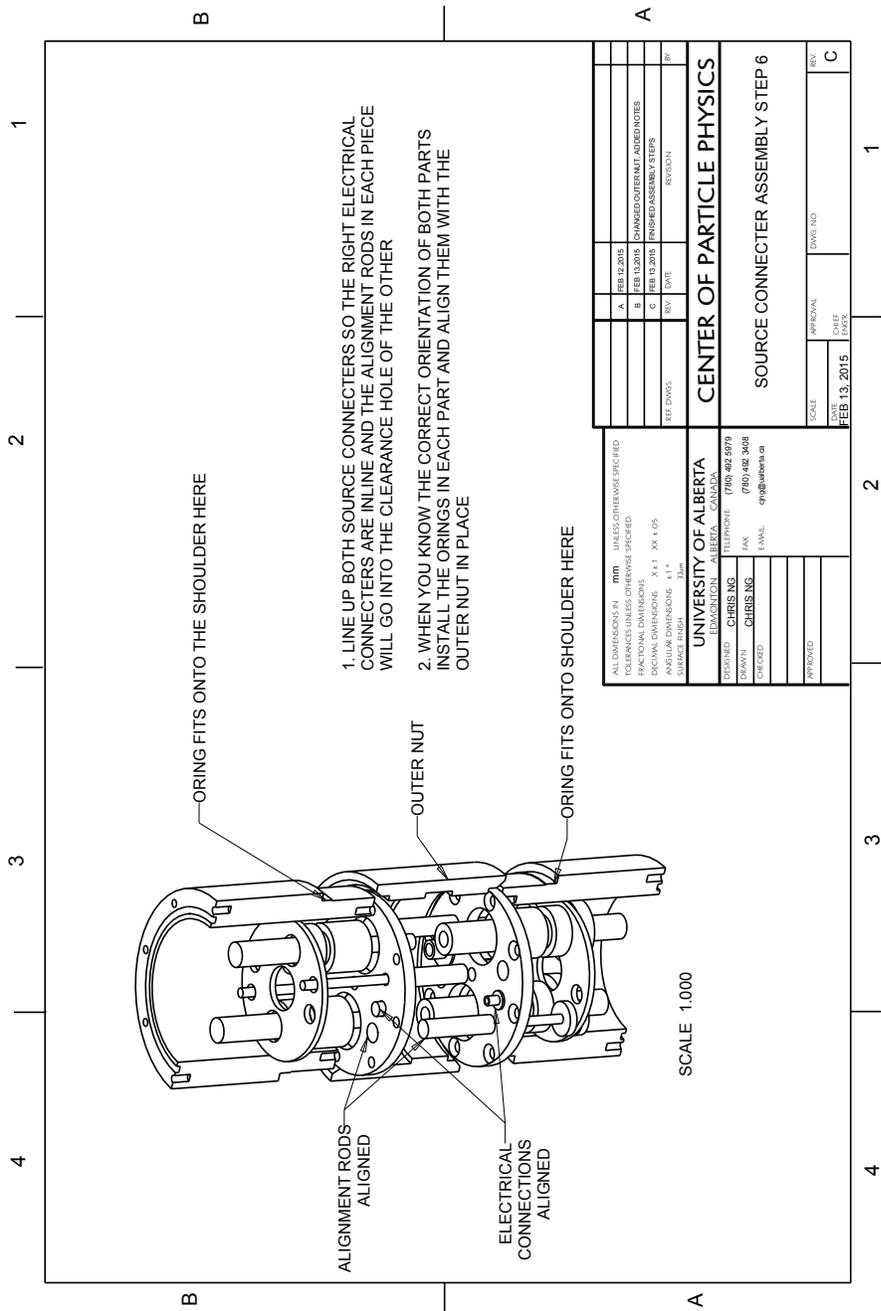


Figure B.7: The quick connect is a custom-designed coupler between the calibration sources and the SNO+ umbilical. A central nut connects the two connector halves, which have opposite threading.

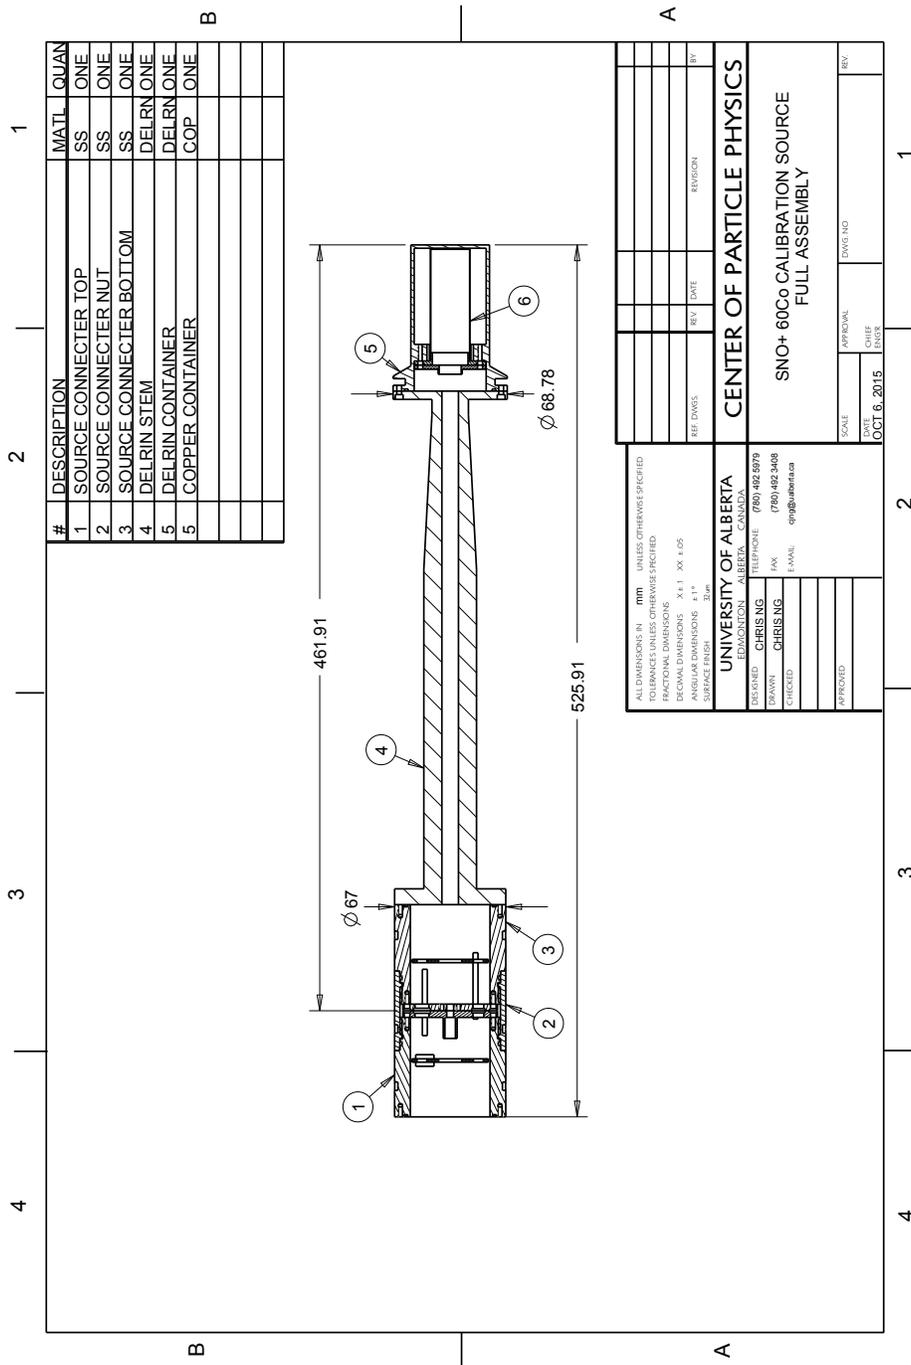


Figure B.8: The complete assembled ^{60}Co calibration source, including both halves of the quick connect.

B.2 Calibration source procedures

B.2.1 Deposition procedure

Liquid $^{60}\text{CoCl}_2$ deposition procedure (October 31, 2013)

Notes: All work will be done in CCIS L2-235A, which is permitted for open source work. An assistant will be available to open doors, record any relevant information and take any photos required, but should at no time handle any radioactive source, or any equipment suspected to be contaminated. The assistant's main role is to ensure the risk of spreading contamination is minimized.

Project: 5 μL droplets of $^{60}\text{CoCl}_2$ (in HCl) will be deposited on an 8mm diameter, 2mm thick disk of plastic scintillator, then left until the HCl evaporates, leaving behind a layer of $^{60}\text{CoCl}_2$ crystals on the scintillator surface.

Materials:

- Spill tray
- Spill tray liner
- Absorbent cloth
- Microscope slides (x3)
- Micropipette capillary
- Micropipette plunger
- Tweezers (straight and curved)
- Petri dishes (x5 for button pieces)
- Screw-cap bottle (for contaminated micropipettes) with radioactive label
- Plastic bottle holder
- Liquid $^{60}\text{CoCl}_2$ source
- Plastic scintillator pieces (x14)
- Lab coat
- Lab goggles
- Gloves (vinyl or nitrile)
- Plastic waste bag
- Radioactive warning signage

Procedure:

- 1) All persons in the lab don lab coat, goggles and double gloves.
- 2) Turn on fume hood light. Place spill tray in fume hood and insert liner. Place the plastic bottle holder on the left side of the spill tray. Ensure absorbent cloth is inside fume hood in case of spill.
- 3) Make a measurement of the background radiation using the Geiger counter. Particular note of the background should be made for the
 - Spill tray
 - Micropipette plunger
 - Bottle
 - Fume hood base
 - Fume hood door
 - Door handles (L2-235A exit and L2-235A1 entrance)
 - Counter near fume hood
 - Sink (including taps)
 - Floor near fume hood

- 4) Place the microscope slides in a line along the front right of the spill tray.

Place the empty screw-cap bottle (remove the lid and store in the cupboard beside the fume hood) to be used for contaminated micropipette storage in the fume hood in the appropriate position on the plastic bottle holder. Lay it on its side, noting the 'up' direction, on the plastic holding tray. Always place the bottle that way, being careful not to upset any contents. The bottle should be labelled with radioactive signage, as they will be used to store contaminated equipment.

Assemble the micropipette by placing the plunger into the capillary (insert on the side closest to the white band), then place the micropipette in the spill tray next to the microscope slides.

Place the waste bag next to the fume hood.

- 5) Using the tweezers, place 14 pieces of plastic scintillator in a line on the microscope slides (5, 4 and 5, with one control button on each slide containing 5 buttons), noting the position of each particular piece (whose dimensions were pre-measured). The scintillator pieces are stored in the cupboard next to the fume hood in petri dishes, each labelled with a scintillator identification number. The pieces labelled 'bottom' should be used. Ensure the petri dishes are returned to the cupboard after removing these pieces. Plastic scintillator should always be handled with the tweezers. Leave tweezers on the spill tray when not in use.
- 6) Remove the liquid $^{60}\text{CoCl}_2$ source from the source storage cabinet and place in spill tray. An assistant should open the door for you to prevent the spread of contamination.
- 7) Very slowly and carefully open the source vial, keeping it upright and on the spill tray to ensure no spillage. Set lid aside on spill tray. Never lift the vial when it is open or in the opening process (it should always sit on the spill tray).
- 8) Hold the pre-assembled micropipette. One hand should be on the plunger, the other on the capillary above the black $5\mu\text{L}$ calibration line. Put end of capillary into source liquid and use plunger to retract $5\mu\text{L}$ of solution.
- 9) Deposit $5\mu\text{L}$ of solution onto the centre of plastic scintillator (by eye). Liquid will likely form a droplet at bottom of capillary. Press the droplet onto the scintillator to deposit the source. Repeat until required number of droplets for appropriate activity is reached. Each button should receive 3 drops per cycle, which is the maximum number allowed owing to the diameter of the droplet that will form on the button when the 3 droplets combine.

- 10) Repeat steps 8-9 for all 14 scintillator pieces. The two control pieces will receive no deposition. A maximum of 3 drops can be dispensed on a button at any time, owing to the diameter of the droplet that will form when the drops combine. Allow at least 15 minutes evaporation time between 3-droplet-depositions on a single button. When all buttons are complete, put the micropipette in its storage bottle using the curved tweezers. The number of drops for each button is as follows:
 - 1A 10 (~300Bq)
 - 1B 11 (~330Bq)
 - 1C 12 (~360Bq)
 - 2A 13 (~400Bq)
 - 2B 15 (~450Bq)
 - 2C 16 (~480Bq)
 - 3A 18 (~540Bq)
 - 3B 20 (~600Bq)
 - 3C 22 (~660Bq)
 - 4A 7 (~210Bq)
 - 4B 9 (~270Bq)
 - 4C 24 (~720Bq)
 - C1 0 (0Bq)
 - C2 0 (0Bq)
- 11) Remove outer gloves and put in waste bag. Pick up the source vial lid and place very carefully on the source vial and tighten the lid, holding the vial with your free hand. Leave the source vial in the spill tray.
- 12) Use a filter to wipe the outside of the source vial (to ensure no radioactivity was transferred to the exterior of the source vial). Staple the filter into an envelope.
- 13) Put lid on the screw-cap bottle holding the contaminated micropipette, ensuring that you do not disrupt the bottle contents by keeping the bottle on its side on the bottle holder while putting the lid on. Use a single filter paper to wipe the complete surface of the bottle and staple the filter into an envelope. Leave the bottle on the bottle holder until the results of the wipe tests are known.
- 14) Return the source vial to the storage locker. Note on the source inventory sheet the amount of activity removed (5 μ L ~200Bq).
- 15) Decontaminate the curved tweezers in the radioactive decontamination sink. Run water from the tap into the sink with as high a flow rate that can be maintained with no splashing. Hold the tweezers under the running water for 1 minute to remove any residual radiation, then towel dry. Wipe the tweezers with a filter paper and staple the filter into the envelope. Store tweezers in a bag to await the results of the wipe test.
- 16) Close the fume hood door completely. The scintillator should remain untouched in the fume hood for approximately 24 hours to allow sufficient time for the HCl to evaporate completely.
- 17) Put up radioactivity warning signs and contact information for emergencies on the fume hood, which explicitly states that the fume hood is not to be opened unless authorized.
- 18) Remove gloves, place in waste bag, then don new gloves. Use the Geiger counter to recount all surfaces outside of the fume hood, as in step 3, noting the count rates in the lab's log book.
- 19) Remove personal protective equipment. Place coat and glasses on hooks and gloves in radioactive waste container, then leave the room.

Completed by _____

B.2.2 Bonding procedure

Plastic scintillator bonding procedure (October 31, 2013)

Notes: All work will be done in CCIS L2-235A, which is permitted for open source work. An assistant will be available to open doors, record any relevant information and take any photos required, but should at no time handle any radioactive source, or any equipment suspected to be contaminated. The assistant's main role is to ensure the risk of spreading contamination is minimized.

Project: The plastic scintillator pieces prepared in the liquid $^{60}\text{CoCl}_2$ deposition procedure will be heat bonded to clean 8mm diameter, 2mm thick disks of plastic scintillator in an oven using an alignment jig. The heat bonding is a slow process that requires approximately 24 hours to cycle completely. The $^{60}\text{CoCl}_2$ will be fully encapsulated in the plastic scintillator after the bonding is complete.

Materials:

- Oven
- Bonding jig (3 pieces)
- Absorbent cloth
- Tweezers
- Plastic scintillator pieces with evaporated $^{60}\text{CoCl}_2$ source (x14)
- Clean plastic scintillator pieces (x14)
- Lab coat
- Lab goggles
- Gloves (vinyl or nitrile)
- Plastic waste bag
- Radioactive warning signage

Procedure:

- 1) All persons int the lab don lab coat, goggles and double gloves.
- 2) Bring oven into room, leave on cart near sink, and plug in.
- 3) Make a measurement of the background radiation using the Geiger counter. Particular note of the background should be made for the
 - Oven interior
 - Oven door exterior
 - Fume hood door
 - Door handles (L2-235A exit and L2-235A1 entrance)
 - Counter near fume hood
 - Floor near fume hood
- 4) Open the fume hood door. The fume hood contains the 14 pieces of plastic scintillator that have had the source deposited on them (2 with none). Separately place the bonding jig pieces in the fume hood next to the spill tray.
- 5) Retrieve one piece of clean scintillator labelled 'top' from the cupboard and, by hand, depress into the indent in the jig weight cylinder. Using the tweezers, pick up the corresponding 'bottom' piece of scintillator with deposition from the spill tray and place in the indent in the jig base.
- 6) Set the weight cylinder on top of the jig base, plastic scintillator sides facing together, ensuring the edges of the base and the weight cylinder are flush. Carefully slide the jig retaining cylinder over the weight cylinder, minimizing any movement of the weight cylinder in the process, and allow the retaining cylinder to come to rest on the jig base.
- 7) Securely holding the jig base, transport the entire jig to the oven. Place the jig on the metal grating near the oven centre, with the thermal indicator strip facing outward (so it is visible through the oven door). Close the oven door.

□□□□ □
□□□□
□□□□ □

8) Put up radioactivity warning signs and contact information for emergencies, which explicitly states that the oven is not to be opened unless authorized, near the oven.

□□□□ □
□□□□
□□□□ □

9) Go through the bonding heat cycle (approximately 28 hours):
Turn the oven on to heat setting H and 'temperature' 28 (start in afternoon)
The oven temperature will rise to approximately 80°C in 3 hours on the oven thermometer or 20 hours on the temperature indicator (more correct). This will be the maximum temperature the oven attains
Allow the oven to sit for approximately 14 hours (overnight) after turning it on to ensure sufficient time for the oven to heat and the scintillator halves to fully bond
Turn of the oven completely and allow to cool for approximately 8 hours, or to about 30°C (on the oven thermometer), at which point the bonding jig is safe to remove.

Note that the bonding heat cycle requires the oven be left on overnight. Always remove gloves and place in waste bag, and remove lab coat and goggles before leaving the room. Also, always re-don all safety equipment upon re-entering the room, even for a task as simple as turning off the bonding oven.

□□□□ □
□□□□
□□□□ □

10) When the bonding heat cycle is complete, open the oven door and remove the bonding jig, holding securely to its base, to the fume hood. Place it beside the spill tray.

□□□□ □
□□□□
□□□□ □

11) Slide the retaining cylinder off of the jig base, then remove the scintillator puck from the jig. Remove the corresponding petri dish from the cupboard next to the fume hood and place the puck in it.

□□□□ □
□□□□
□□□□ □

12) Place the petri dish containing the completed puck into the locked source storage box.

□□□□ □
□□□□
□□□□ □

13) Use the Geiger counter to recount all surfaces outside of the fume hood, as in step 3, noting the count rates.

□□□□ □
□□□□
□□□□ □

14) Repeat steps 4-13 for all pieces.

- 15) When bonding is complete for all scintillator pieces, a wipe test of the room must be done. The wipe test consists of wiping a filter paper (in the cupboard next to the fume hood) over surfaces in the lab to collect any possible contamination. Each filter should cover an area of approximately 1m². The papers should be placed in envelopes and sent to the UofA's Radiation Safety group to be counted in their scintillation counter. The surfaces to be wiped include the:
 - Spill trays
 - Screw-cap bottles
 - Fume hood base and door
 - Door handle exiting L2-235A
 - Door handle entering L2-235A1
 - Oven door and controls
 - Oven interior
 - Bonding jig
 - Counter top
 - Floor
 - Sink

Mail the filter papers (multiple stapled inside each envelope) to:
Cory McHugh
Radiation Protection Advisor, Department of Environmental Health & Safety
Risk Management Services
University of Alberta
3107 Research Transition Facility
Edmonton, AB T6G 2R5

Radiation Safety will notify you if contamination is found. Wipe tests are more sensitive than the Geiger counter, but the latter can alert you early on if there is contamination in the room. The Geiger counter must be used each day that radioactive work is done.

- □ 16) If the counting results from EH&S show the wipes of the buttons are consistent with background, then move the buttons into sealable baggies for storage. The baggies are labelled with the tests that were/are to be performed on the button, as well as the source type and measured activity (once known). These baggies should be stored in the locked source storage box.

Completed by _____

B.2.3 Button counting and testing procedure

⁶⁰Co button source counting procedure (February 5, 2014)

Notes: All work will be done in CCIS L2-114, CCIS L2-116 and CCIS L2-235A, which are permitted for closed source work. The ⁶⁰Co button sources should only be handled either using tweezers or with gloved hands, and should be stored in their respective containers in the locked source storage box when not in use.

Project: The ⁶⁰Co button sources will be counted on the HPGe detector in L2-114 to determine the absolute rate of each source, as well as to determine if there is any surface contamination or if the sources leak any radioactivity.

Materials:

- Filter paper
- ⁶⁰Co button sources (x14)
- Tweezers
- 20 mL Nalgene bottle (x2)
- Paraffin sealing tape for bottle
- Reverse osmosis (RO) lab water from L2-116
- Graduated cylinder
- Digital scale
- Polystyrene collar
- 1L glass beaker
- 250 mL Nalgene bottle (x2), one clean and one containing contaminated capillary
- Scotch tape
- Calibrated ⁶⁰Co source (old) from L2-235A1
- 1/2 thickness plastic scintillator button
- Baggies to store used filter papers
- Gloves (vinyl or nitrile), lab coat and safety glasses
- Kimwipes
- Large ultrasonic cleaner in L2-116
- Radioactive warning signage

Procedure:

- 1A 1) Each button source should be counted on the HPGe detector in L2-114 to determine their absolute count rates.
- 1B
- 1C
- 2A Put a filter paper on the centre of the detector, then, using the tweezers, place one of the button sources on the centre of the filter paper. The filter paper is simply acting as a non-intrusive barrier between the source and the detector to ensure no contamination is transferred between them.
- 2B
- 2C
- 3A
- 3B
- 3C Count each of buttons 1A-4C until 10,000 counts are received under the 1.33 MeV gamma peak, which is defined to be between channels 13815 and 13864 (some counts from the peak tails will be missing, but should be compensated for by any background counts under the peak, so the total number of counts of the 1.33 MeV gamma will be close to 10,000). Approximately the same number of counts should be observed under the 1.17 MeV gamma peak, with a small sum-peak observed at 2.50 MeV. Buttons C1 and C2 are control pieces with no radiation deposited in them. As such, these pieces should be counted overnight (approximately 24 hours) to ensure no ⁶⁰Co counts from the sources are observed.
- 4A
- 4B
- 4C
- C1
- C2

- Calibration Source 2) To get a measure of the absolute rate of the button sources, it is possible to compare the HPGe spectrum of each to the spectrum of a calibrated ^{60}Co source. We have two calibrated sources, one 'old' and one 'new'. To get the absolute count rate, we simply take a ratio of the peak heights of the unknown to the calibrated source, since we expect the peak shape to be the same. In this way, all unknown systematics and backgrounds will simply cancel out (except for the sum-peak, the height of which depends on the source rate).
- Put a filter paper on the centre of the detector, then place a cleaned 1/2 plastic scintillator button on the centre of the filter paper. The thickness of this button is $1.95 \pm 0.01\text{mm}$. Next, place the old calibration source on the button such that the active part of the source is directly over the centre of the button. The purpose of the button is to have the ^{60}Co at approximately the same height above the detector as it is in the button sources. Count the source until 100,000 counts are received under the 1.33 MeV gamma peak, which is defined to be between channels 13815 and 13864.
- Filter 3) To estimate the background from the filter paper, count a filter paper overnight (approximately 24 hours). This will be used as the background run for the wipe tests.
- Bottle + Water 4) As all of the wipe test counting and soak test counting will be done using the RO water supply in L2-116, the background of this water source should be counted. A 60mL screw-cap Nalgene bottle will be used to hold the water.
- First, a cleaned bottle containing water should be counted for several days to obtain an estimate of the background. Measure the mass of the bottle (including the lid). Next, place approximately 15 mL of the RO water in the bottle (use the graduated cylinder). Again, measure the mass of the bottle plus RO water to give us a measure of the total mass of the water being counted (the volume is really only accurate to approximately 1 mL, whereas the mass is accurate to 0.01 g). Secure the cap with sealing tape. Place the bottle (with its lid on) on the centre of the detector.
- Count the RO water sample for approximately 4 days (96 hours) to get a good estimate of the backgrounds in the water. With the lid and tape on the bottle, the evaporation rate of the water should be minimal.

- | | | |
|---------|--------------------------|--|
| 1A | <input type="checkbox"/> | |
| Wipe 1A | <input type="checkbox"/> | |
| 1B | <input type="checkbox"/> | |
| Wipe 1B | <input type="checkbox"/> | |
| 1C | <input type="checkbox"/> | |
| Wipe 1C | <input type="checkbox"/> | |
| 2A | <input type="checkbox"/> | |
| Wipe 2A | <input type="checkbox"/> | |
| 2B | <input type="checkbox"/> | |
| Wipe 2B | <input type="checkbox"/> | |
| 2C | <input type="checkbox"/> | |
| Wipe 2C | <input type="checkbox"/> | |
| 3A | <input type="checkbox"/> | |
| Wipe 3A | <input type="checkbox"/> | |
| 3B | <input type="checkbox"/> | |
| Wipe 3B | <input type="checkbox"/> | |
| 3C | <input type="checkbox"/> | |
| Wipe 3C | <input type="checkbox"/> | |
| 4A | <input type="checkbox"/> | |
| Wipe 4A | <input type="checkbox"/> | |
| 4B | <input type="checkbox"/> | |
| Wipe 4B | <input type="checkbox"/> | |
| 4C | <input type="checkbox"/> | |
| Wipe 4C | <input type="checkbox"/> | |
| C1 | <input type="checkbox"/> | |
| Wipe C1 | <input type="checkbox"/> | |
| C2 | <input type="checkbox"/> | |
| Wipe C2 | <input type="checkbox"/> | |
- 5) A wipe test of each button should be performed to ensure there is no surface contamination. This has been done using dry filter papers that were then counted by the Radiation Safety Office. We will repeat this ourselves using the HPGe counter.
- Don gloves. Place a dry filter paper on the centre of the detector. Using a reagent bottle filled with RO water, moisten another filter paper. Remove one button from its storage container using a gloved hand and wipe with the moistened filter paper. Rub both faces and the cylindrical side of the button source. Place the moistened filter paper on a Kimwipe. Remove any excess water from the button using another Kimwipe (this should be minimal), then place the button back in its storage container. Repeat for all buttons, stacking subsequent wipes on top of the first wipe, and allow the wipes to dry overnight.
- Count the filter papers overnight (approximately 24 hours) to ensure the count rate is consistent with background. If the count rate is not consistent with background, then it could indicate that surface contamination was removed by the wipe. When the counting is done, remove the filter papers together and put into a storage bag until it is confirmed that no contamination is present on the wipes. If there is any indication that contamination is present, then the wipe tests will need to be re-counted using a bifurcate search method to determine which wipe(s) contain contamination. In this case, the wipe tests should also be re-done to determine if the first wipe removed all the contamination.
- When the filter papers are done being counted, re-count each of the buttons, as outlined in step 1. If the count rate is not consistent with that measured for the same button in step 1, then it could indicate surface contamination was present in the initial count and was removed by the wipe. The sum of the buttons plus the wipes should be consistent with the sum of the buttons before the wipe test.
- | | | |
|----|--------------------------|--|
| 1A | <input type="checkbox"/> | |
| 1B | <input type="checkbox"/> | |
| 1C | <input type="checkbox"/> | |
| 2A | <input type="checkbox"/> | |
| 2B | <input type="checkbox"/> | |
| 2C | <input type="checkbox"/> | |
| 3A | <input type="checkbox"/> | |
| 3B | <input type="checkbox"/> | |
| 3C | <input type="checkbox"/> | |
| 4A | <input type="checkbox"/> | |
| 4B | <input type="checkbox"/> | |
| 4C | <input type="checkbox"/> | |
| C1 | <input type="checkbox"/> | |
| C2 | <input type="checkbox"/> | |
- 6) A soak test of each button should be performed to ensure no radioactive material leaches out (off) of the button (surface) when submerged in water. If no radiation leaches into the water, then it implies the ⁶⁰Co is completely encapsulated in the plastic.
- Fill the large US cleaner in L2-116 with RO water and heat it to 45 °C. The US cleaner will require several hours to heat the water, which is initially at room temperature, but will be able to adequately hold the temperature of the water at 45 °C for the duration of the soak tests.
- Fill a clean 1 L beaker 3/4 full with RO water. Place the beaker in the polystyrene collar and float on the surface of the RO water in the US cleaner in L2-116 (the collar simply prevents the beaker from tipping, where the beaker is there to contain any spillage).
- Don gloves. Measure the mass of each clean 60mL Nalgene bottle (with lid on), then fill them with approximately 15 mL of RO water (use the graduated cylinder). Measure the mass of the bottle (with lid on) plus RO water. Then, using the tweezers, place buttons 1A-4C in one bottle (15 mL of RO water is sufficient to ensure every button is completely submerged) and C1,C2 in the other. Re-measure the masses of the bottles, water and buttons. Secure the lid with paraffin tape, then measure the mass once again. Next, float the Nalgene bottle in the water in the 1 L beaker already in the US cleaner. Leave the bottle in the beaker for approximately 4.5 hours, which allows sufficient time for the water in the bottle to heat and remain at 45 °C.

- 1A
- 1B
- 1C
- 2A
- 2B
- 2C
- 3A
- 3B
- 3C
- 4A
- 4B
- 4C
- C1
- C2
- Bottle

7) Remove the bottles from the water bath (the beaker can remain) and place on a Kimwipe (to remove external water). Measure the masses of the bottles, water and buttons to test whether any water has evaporated from inside the bottle, then remove the tape from the bottle. Take the buttons out of the bottle and remove any excess water from the buttons using a Kimwipe, placing the buttons back in their respective storage bags when dry. Put the lid back on the bottle and re-secure with tape.

Count the dried buttons according to the procedure in step 1.

When the buttons are done being counted, place the bottle containing the water on the centre of the detector and count for several days (approximately 100 hours). Ideally, the resulting spectrum will be consistent with the background measured in step 4.

- Empty Bottle
- Micropipette

8) To ensure the activity of each button is consistent with the amount of activity removed from the source vial, the contaminated micropipette should be counted. The sum of the micropipette and buttons should be equal to the expected amount of removed activity, which is the difference in the vial activity before and after the deposition.

Cover the endcap of the HPGe counter with a Kimwipe and secure a clean and empty 250 mL Nalgene bottle to the wipe with tape. Count overnight (approximately 24 hours) to obtain a background sample for the bottles.

Remove the empty bottle from the detector and replace with a bottle containing the capillary and plunger. Ensure that the bottle is not jostled or rotated, so that the capillary/plunger does not move on the inside (this limits how far an contamination will spread inside the bottle). Count the bottle until 10,000 counts are received under the 1.133 MeV gamma peak, which is defined to be between channels 13815 and 13864.

Because the source in this case is assumed to be uniformly distributed across the entire inner (outer) surface of the capillary (plunger, sans handle), the absolute count rate will need to be adjusted by obtaining a geometric correction factor from Monte Carlo.

- 9) After the micropipette has been counted, it should be decontaminated in L2-235A. Don gloves, safety glasses and lab coat, then measure the background activity in the room with the Geiger counter.

Run water from the tap into the radioactive decontamination sink with as high a flow rate that can be maintained with no splashing. Remove the micropipette from the bottle and pull out the plunger. Hold the capillary under the running water for 1 minute to remove any residual radiation, then place in the waste bag. Do the same for the plunger, but place the dried plunger in plastic bag and towel used to dry it in waste bag. Rinse bottles under water for 2 minutes, then towel dry and place in cupboard and towel in waste bag.

The plunger and bottles need to be wipe tested to check if they were sufficiently decontaminated. A wipe test consists of wiping a filter paper over surfaces in the lab to collect any possible contamination. Each filter should cover an area of about 1m². The papers should be placed in envelopes and sent to the UofA's Radiation Safety group to be counted in their scintillation counter. The surfaces to be wiped include the:

- Micropipette plunger
- Interior and exterior of the bottle
- Counter top
- Floor
- Sink

Mail the filter papers (multiple stapled inside each envelope) to:

- Cory McHugh
- Radiation Protection Advisor, Department of Environmental Health & Safety
- Risk Management Services
- University of Alberta
- 3107 Research Transition Facility
- Edmonton, AB T6G 2R5

Radiation Safety will notify you if contamination is found. Wipe tests are more sensitive than the Geiger counter, but the latter can alert you early on if there is contamination in the room. The Geiger counter must be used each day that radioactive work is done.

Completed by _____

B.2.4 Source assembly and leak testing, and Delrin[®] annealing procedures

⁶⁰Co source assembly and testing procedure (18 June 2016)

Notes: All work will be done in CCIS L2-114 and L2-116. The ⁶⁰Co button sources should only be handled either using tweezers or with gloved hands, and should be stored in their respective containers in the locked source storage box when not in use. Preventing the spread of contamination is critical. Assistants or experts should be used to assist with/do work, as required/specified. Work will be video recorded to enable a review of whether cleanliness procedures were followed.

Project: The ⁶⁰Co source will be assembled and tested to ensure both the inner copper and outer Delrin encapsulations are both liquid and helium leak-tight. Each level of encapsulation will be tested to ensure the source fully encapsulates the internal radioactivity.

Materials:

- Copper encapsulation parts (flange/box), stainless steel 4-40 screws/locknuts (x8 each), indium wire, wrenches
- Delrin container/stem, Viton o-ring and stainless steel 4-40 screws/locknuts (x8 each)
- Source connector, Viton o-ring and stainless steel M3 screws (x8)
- Wired SMB (Pomona 72999) and LEMO (FAG.0B.304.CLA) plugs (x1 each)
- RG-174U coaxial cable and AWG26 wires (x4), 50 cm lengths,
- Shrink wrap tubing (1/8 " clear)
- Wire cutters
- Soldering iron, InterTAN 64-027 0.05" Pb-free, rosin core solder
- Multimeter
- PMT (Hamamatsu H10721-110P)
- ⁶⁰Co button source (x1)
- BC-630 optical grease and applicator (wooden stick)
- 8 cm x 8 cm sheet of aluminized mylar (x1)
- Kapton tape and scissors
- Hot glue gun and glue
- Digital scale
- Helium leak-check equipment (ASI 20 MD leak detector and tubing)
- Acrylic vacuum chambers and pump
- Plastic tank for helium-atmosphere source assembly
- 250 mL screw-cap Nalgene bottles (x2), 1 L beaker (x1)
- Reverse osmosis (RO) water from L2-116
- Tweezers (plastic, disposable)
- Kimwipes
- Baggies (zip-sealed)
- Alconox solution and ethanol
- Gloves (vinyl or nitrile), lab coat and safety glasses
- Geiger counter, filter papers, envelopes for radiation monitoring
- Ultrasonic cleaner in L2-116 (and Alconox)

Procedure:

Work site prepared 1) Don gloves. Prepare a work area by cleaning the lab bench surface with a solution of diluted Alconox+RO water followed by a Kimwipe wetted with ethanol.

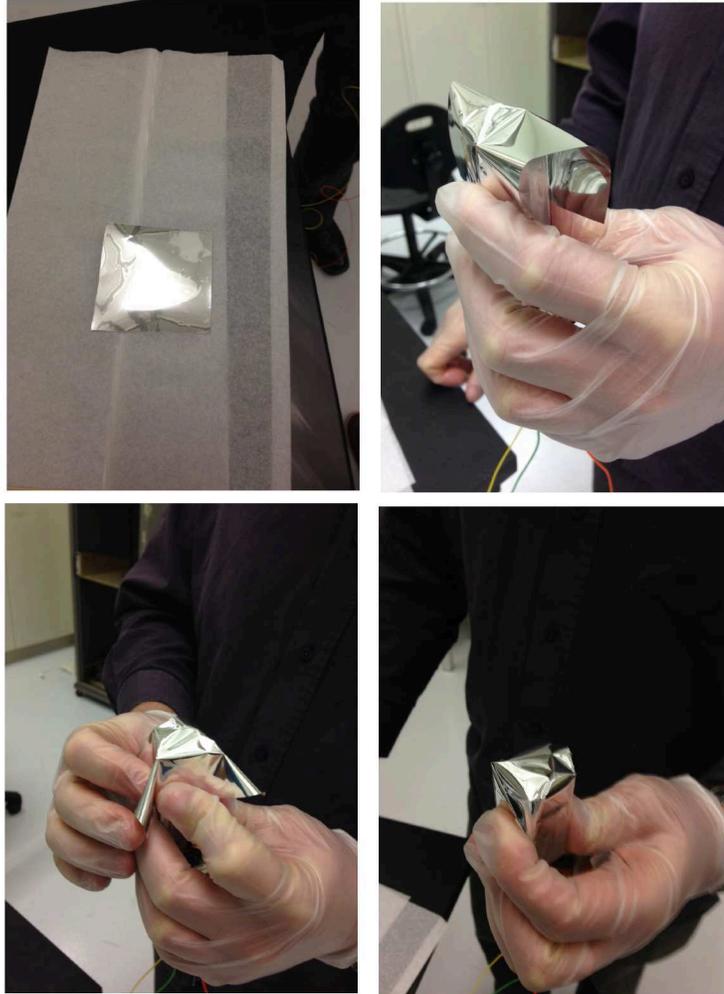
Counting complete Use the Geiger counter to ensure there is no pre-existing radiation in the work area, including the lab bench, floor near the bench and door to the lab.

Surface	Rate pre-work (CPM)	Rate post-work (CPM)
Background		
Lab bench		
Floor		
Door		

2) The inner copper encapsulation consists of two parts: the box that holds the PMT and the flange that allows the PMT wiring to pass through the copper via a glass vacuum feedthrough. These two parts will be joined using an indium wire crush seal. The series of steps below concludes with the joining of the copper parts.

- | | | |
|----------------------------|--------------------------|--|
| Copper clean | <input type="checkbox"/> | The copper box and flange must be cleaned. Don a new pair of gloves. Wipe copper box and flange with a Kimwipe wetted with Alconox+RO water, then place the copper box, flange and 12 4-40 screws in the ultrasonic bath containing RO water and Alconox at 65°C. Clean for 20 minutes, then repeat the ultrasonic clean with RO water only at 20°C for 20 minutes. Change gloves and remove items from the ultrasonic cleaner. Pat dry with a clean Kimwipe, then place on counter on another clean Kimwipe (with another covering the parts) to air dry for 2 hours. Move to a new zip-sealed baggie for storage once dry. |
| Copper stored | <input type="checkbox"/> | |
| Pins labelled | <input type="checkbox"/> | <p>The PMT wires must be soldered to the feedthrough pins – one pin for each of the supply voltage (red) wire, control voltage (white) wire, ground (black) wire, return voltage (blue) wire, coaxial cable and coaxial shield. The pins shall be labelled 1-6 in the order as above, with punched indicators on the flange.</p> <p>The black (ground), red (supply +5 V), white (control +0.5 V to +1.1 V), and blue (reference +1.2 V output) wires (AWG26) and a coaxial cable (signal) (RG-174U) on the PMT will be soldered to the corresponding pins on the vacuum feedthrough in the PMT flange.</p> |
| Wires clipped and stripped | <input type="checkbox"/> | Lay out the soldering equipment (iron, solder, wires, shrink wrap and wire cutters) on the lab bench. Soldering will be done by a Department of Physics electronics technician. Gloves will be worn during soldering. Wires on the PMT will be clipped to a length of 5 cm with the RG-174U cut to 6 cm, and ~5 mm of casing stripped from the ends of the AWG26 wires and 12 mm stripped from the outer part of the RG-174U and 5 mm from the inner conductor. An ~8 mm piece of shrink wrap tubing will be placed around each wire, except the coaxial which will be left non-shrink wrapped (for ease of assembly), to insulate each connection post-soldering. |
| Shrink wrap in place | <input type="checkbox"/> | |
| Soldering complete | <input type="checkbox"/> | Remove the copper flange from the storage baggie. The wires are then soldered onto the pins, coaxial cable shield first, then the inner conductor of the coaxial cable, then the other wires, using a standard soldering iron and the InterTAN 64-027 Pb-free solder. The coaxial shield should not be pre-tinned (for ease of assembly). The shrink wrap will be fitted using a heat gun. The heat gun will be directed away from the assembly until it has warmed up, and then directed toward the wires but away from the PMT to ensure the PMT does not overheat. |
| Shrink wrapped | <input type="checkbox"/> | |
| Continuity check | <input type="checkbox"/> | Test to ensure the soldering of the PMT wiring was successful. To make sure there are no stray wires from the shield, do a continuity check on the coaxial cable between the shield and core (using the multimeter). Next, put the PMT into a dark box on a clean Kimwipe and use a power supply to supply +5 V and +1.0 V to the PMT. Read out the PMT to an oscilloscope. If the PMT creates pulses that are readable, then the soldering will be deemed to have been successful. |
| Soldering successful | <input type="checkbox"/> | |
| Bench prepared | <input type="checkbox"/> | 3) Replace gloves. Clean the lab bench surface with a dilute solution of Alconox+RO water followed by a Kimwipe wetted with ethanol. Place a clean Kimwipe on the bench surface on which to do work. The button source, tweezers, Alconox solution and additional Kimwipes should be close at hand. |

- Button out Before it may be coupled to the PMT face, the button source must be cleaned. Don gloves. Wet a Kimwipe with a dilute solution of Alconox. Remove the button source from its storage baggie using tweezers and place on the Kimwipe. Dispose of the tweezers in the garbage.
- Alconox
RO water Rub the surface of the button for 30 seconds, then lay the button down on a dry Kimwipe. Place the used Kimwipe into a clean baggie. Repeat the process using a Kimwipe wetted with RO water. Dry the button source with a clean Kimwipe (store this Kimwipe with the others), then, using tweezers, put the button source into a clean baggie for storage. Dispose of the tweezers in the garbage. Save the Kimwipes for later counting on the HPGe detector to measure if any surface contamination was present. Dispose of the remaining bench Kimwipe in the garbage.
- Kimwipes stored
- Button clean Count the baggie containing the Kimwipes on the HPGe to determine whether any contamination was removed from the surface of the button. This is for reference only. The button is not considered a level of containment.
- Kimwipes counted
- Work site prepared Replace gloves. Clean the lab bench surface with a dilute solution of Alconox+RO water followed by a Kimwipe wetted with ethanol. Place a clean Kimwipe on the bench surface on which to do work. The PMT, optical grease (BC-630), aluminized mylar foil, Kapton tape, button source, scissors and tweezers should be close at hand.
- PMT/flange cleaned Wipe the flange with a Kimwipe wetted with ethanol to remove any residual flux from the soldering process. Wipe the aluminum body and glass face of the PMT with a Kimwipe wetted with ethanol. Lay the assembly on the dry Kimwipe on the bench surface and dispose of the soiled Kimwipe in the garbage.
- Button/PMT coupled Have an assistant open the BC-630 container. Use a wooden stick to place a small amount of the compound on the centre of the PMT face. Using tweezers, remove the button source from its storage baggie and place it onto the PMT face. Dispose of the tweezers in the garbage. Press the button, causing the compound to spread over the PMT face, visibly ensuring no air bubbles are present between the button and glass. Hold the PMT face-up throughout (so the button does not slide from the PMT face centre).
- Mylar applied Replace gloves. Wrap the button-end of the PMT with a reflective 8 cm x 8 cm aluminised mylar sheet (see pictures below). Place the centre of the mylar sheet onto the centre of the button source (top right). Fold opposite sides of the mylar down the PMT body, then fold the other 2 sides down the body (bottom left). The mylar should be taught over the PMT face, applying slight pressure to the plastic scintillator. The remaining mylar 'wings' should be folded toward each other on 2 sides of the PMT (bottom right).



- Kapton tape applied Have the assistant retrieve a 20 cm long strip of Kapton tape. Use the tape to seal the edges of the mylar onto the PMT body. This also secures the button in place, so the PMT may now be laid on its side.
- PMT inside box 4) The PMT is now ready to slide into the copper box. Replace gloves. Remove the clean copper box from its baggie. Slide the wrapped PMT into the copper box and press down until the mylar contacts the bottom of the box.
- Glue seam applied To ensure the PMT will not move inside the copper encapsulation, place a seam of hot glue (from a hot glue gun) around the bottom edge of the PMT (wire-side) between the PMT and the copper. Wait several minutes for the glue to set. Prepare the helium environment while waiting.

- Helium environment created The plastic tank will be used to create the helium-rich environment in which the copper assembly will occur. Secure the tank raised on blocks with its opening toward the lab bench. Place a small stand within the tank to create a working surface, and cover with a clean Kimwipe. Secure the gas line to the inside bottom of the tank with tape, then fill the tank with helium gas from the bottle of compressed helium gas in L2-114 while holding the He sniffer next to tank near the lip. When the sniffer detects He, then stop filling.
- Copper laid out Move the copper box/flange with PMT, indium wire, and 8 4-40 screws, nuts and locknuts into the helium-filled tank and set on the Kimwipe. Set the wrenches and wire cutters down on the lab bench on a clean Kimwipe.
- Indium wire in place Replace gloves. Place the indium wire in the groove on the bottom flange and cut so there is minimal overlap of the ends (there should be no gap between the wire ends, however). Place top flange onto the bottom so that the ridge goes into the groove. Ensure the wires from the PMT are securely tucked inside the copper box.
- Nuts finger-tight
- Nuts wrench-tight Place eight of the screws into the holes nearest the corners of the box (the four remaining holes are for securing the copper into the Delrin enclosure). The screw heads will be on the box-side of the flange, with a nut holding the screw in place on the opposite side. Using the plain nuts, finger tighten each of the nuts, tightening opposite nuts each time. Then, using a wrench, apply a half turn to each nut, successively tightening opposite nuts, until each is wrench-tight. This will ensure the indium wire is sufficiently crushed to create a leak-tight seal. One screw at a time, replace the nut with a nylon-insert lock nut, again wrench-tight. This will prevent any future loosening of the copper flange, ensuring the integrity of the seal.
- Nuts replaced with locknuts
- Seal complete
- 5) The copper encapsulation is now completely constructed. It must be tested to ensure it is leak-tight. The interior of the source is filled with helium because it was assembled in a helium-rich environment. There are 2 sets of tests that should be completed to establish confidence that the enclosure is not leaking. Always wear gloves when handling the copper.
- Acrylic chamber prepared Prepare the acrylic vacuum vessel for leak-checking. The helium leak checker is connected inline with the gas flow circuit of the vessel/pump system. Test the system by pumping the vessel down to 20 mbar and monitoring the level of helium present in the stream. This constitutes a background count for the helium leak-check of the source assembly. Release the vacuum.
- Background rate noted
- Note that, if at any point a leak in the copper encapsulation is found, then it will need to be re-sealed with a new length of indium, as per step 4, then re-leak-checked.
- Sniffer smells no He 1 – Remove the source from the plastic tank. Use the sniffer to check for helium escaping from the indium seal of the copper. Next, move the copper into the acrylic vacuum vessel. Pump the vessel down to 20 mbar and monitor the gas stream for helium. If no helium above background is observed in that time, then the source assembly will be deemed to not be leaking helium.
- Not leaking

- Mass of enclosure 2 – Measure the mass of the copper enclosure (accurate to 0.01 g). Fill a beaker with approximately 500 mL RO water, then submerge the enclosure in the RO water (the flange side will sink). Place the beaker into the acrylic vacuum chamber then pump the chamber down to 20 mbar. Watch for bubbles to emerge from the solder joints/indium seal on the enclosure during pumping. Air bubbles indicate a leak in the enclosure (and may point to the location of the leak).
- Observe for bubbles
- Copper dried Return the vacuum chamber to atmospheric pressure. Remove the copper enclosure from its water bath. Dry the copper using a Kimwipe, then leave on the lab bench on a clean Kimwipe to air dry for 2 h (covered with another Kimwipe). Measure the mass of the enclosure once the air dry is complete. If there is a leak present in the system, we would expect the mass to be higher (as the enclosure would fill with water).
- Mass of enclosure
- Not leaking If no bubbles were observed and the measured masses post- and pre-soak are consistent, then the enclosure will be deemed to not be leaking liquid.
- Soaked 4 hours 6) Finally, the copper enclosure must undergo a soak test to ensure it fully encapsulates the internal radioactivity. Place the enclosure in a water bath at 45 °C for 4 h. Collect a sample (~50 mL) of water from the bath into a screw-cap Nalgene container and count the water sample on the HPGe detector. Remove the enclosure from the bath and dry using a Kimwipe. Leave on the lab bench on a clean Kimwipe (and covered with another Kimwipe) to air dry.
- Water sample collected
- Water counted If no contamination is present in the water after a count of 24 h, then the assembly procedure may continue. Wipe the copper enclosure with a Kimwipe wetted with a dilute Alconox+RO water solution, followed by an ethanol wipe. Place the enclosure into a clean zip-sealed baggie for storage.
- No contaminants detected If the water sample shows signs of contamination, it is likely the surface of the container was contaminated during assembly. In this case, the procedure should be re-done completely with a new copper enclosure.
- 7) The PMT electronics now need to be connected through the Delrin source stem to the source connector (which will allow the source to connect to the SNO+ electronics). The source connector contains two electronics connectors, a 4-pin LEMO and an SMB for the coaxial cable.
- Parts clean and bagged First, all parts must be cleaned. Don gloves. Clean the Delrin stem and enclosure, and the source connector with a Kimwipe wetted with a dilute Alconox+RO water solution. The Delrin parts, quick connect, stainless steel screws and lock nuts, and Viton o-rings should then be washed in an ultrasonic bath containing Alconox+RO water at 65°C for 20 minutes, followed by an ultrasonic bath using only RO water at 20°C for 20 minutes. Lightly wipe the all parts with Kimwipes to dry, allowing a further 2 h of air drying time on a clean Kimwipe on the lab bench (with another Kimwipe covering the parts). The Delrin enclosure, screws, lock-nuts and o-ring should then be secured in a baggie and set aside for later use.

- The copper enclosure must now be connected through the Delrin stem to the quick connect. Soldering will be done by a Department of Physics electronics technician.
- Wiring soldered to copper Don gloves. Strip 8 mm of casing from each end of each AWG26 wire and 12 mm from the outer part of each end of the RG-174U cable and 5 mm from the core of each end. Place an 8 mm length of shrink wrap tubing over one end of each wire. Have an assistant remove the copper enclosure from its storage baggie. Using a standard soldering iron and the InterTAN 64-027 Pb-free solder, solder the coaxial shield to pin 6, then the coaxial core to pin 5, followed by the AWG26 wires to the appropriate pins on the glass feedthrough on the copper enclosure. The supply voltage wire is pin 1, control voltage wire is pin 2, ground wire is pin 3, and reference voltage wire is pin 4. Solder the RG-174U coaxial cable core to pin 5 and shield to pin 6. The coaxial shield should not be pre-tinned (for ease of assembly).
 - Shrink wrap The shrink wrap will be fitted using a heat gun, directed (as much as possible) away from the copper enclosure to prevent it from overheating.
 - PMT tested Test to ensure the soldering was successful. Use a power supply to supply +5 V and +1.0 V to the PMT. Read out the PMT to an oscilloscope. If the PMT creates pulses that are readable, then the soldering will be deemed to have been successful.
 - O-ring in place Don gloves. Pass the wires through the bore hole in the stem, ensuring the tapered end of the stem is nearer the Delrin enclosure. Fit the Viton o-ring (size 033) into the groove on the top flange of the stem, then pass the wires into the source connector. The wires will now be soldered onto the glass feedthrough embedded in the source connector. Soldering will be done by a Department of Physics electronics technician. Gloves will be worn during soldering.
 - Wiring soldered to source connector Wires on the PMT will be clipped to a length of 5 cm with the RG-174U cut to 6 cm, and ~5 mm of casing stripped from the ends of the AWG26 wires and 12 mm stripped from the outer part of the RG-174U and 5 mm from the inner conductor. An ~8 mm piece of shrink wrap tubing will be placed around each wire, except the coaxial which will be left non-shrink wrapped (for ease of assembly), to insulate each connection post-soldering.
 - Heat shrink Using a standard soldering iron and the InterTAN 64-027 Pb-free solder, the wires are then soldered onto the feedthrough pins, coaxial cable shield first, then the inner conductor of the coaxial cable, then the other wires. The coaxial shield should not be pre-tinned (for ease of assembly). The shrink wrap will be fitted using a heat gun.
 - Stem assembled The 6 pins of the glass feedthrough will have been pre-soldered to the 4-pin LEMO connector, and the SMB connector.
 - PMT test **THE PINNING OF THESE CONNECTORS IS NOT YET DEFINED**
Once the PMT electronics are fully connected, the source connector may be secured to the stem. Press the stem onto the source connector, aligning the holes through the stem to the threaded holes on the source connector (do not rotate the stem more than 1/8 turn relative to the quick connect to prevent over-stressing of the wires). Secure the M3 screws opposite from each other in succession, tightening half turns each time until the o-ring is fully compressed.
 - PMT test The stem is now fully assembled. The connections should be tested to ensure a correct PMT signal is obtained. Connect the two LEMO and SMB connectors on the quick connect to their mating ends, with the corresponding pins matching those in the source connector. The soldering is deemed successful if PMT signals are seen on an oscilloscope.

- Tank filled 8) The remainder of the source may now be assembled. The plastic tank will be used to create the helium-rich environment in which the source assembly will occur. Secure the tank raised on blocks with its opening toward the lab bench. Place a small stand within the tank to create a working surface, and cover with a clean Kimwipe. Also place a stand with clamp in the tank (to be used to hold the stem during assembly). Secure the gas line to the inside bottom of the tank with tape, then fill the tank with helium gas from the bottle of compressed helium gas in L2-114 while holding the He sniffer next to tank near the lip. When the sniffer detects He, then stop filling.
- Parts in tank Don gloves. Take all of the clean and dried parts into tank (after removing those in storage from their baggies). Place the Delrin container, o-ring, screws and nuts onto the work surface. Clamp the stem/quick connect in the tank with the PMT hanging downward (to protect the wires from breaking). Set the wrench down on the lab bench on a clean Kimwipe.
- Stem clamped
- O-ring in place Slide the Viton o-ring into the groove in the Delrin container. Rotate the stem to about 45° (ensuring the copper is completely engulfed in the helium) while supporting the copper enclosure, then slide the copper enclosure into the Delrin container. Secure it in place using 4 of the screws.
- Copper secured
- Stem attached Ensuring the wiring is safely tucked inside the container, fasten the stem to the container using the 8 4-40 screws turning into the lock nuts. Secure the screws opposite from each other in succession, tightening half turns each time until the o-ring is fully compressed.
- Source assembled The source is now assembled, with helium gas contained in the inner vacant volumes. Leave the source in the tank until it is ready to be helium leak-checked.
- 9) Now that the source is completely assembled, it must be helium leak-checked. The interior of the source is filled with helium because it was assembled and stored in a helium-rich environment. There are 2 sets of tests that should be completed to establish confidence that the source is not leaking. Always wear gloves when handling the source.
- He leak-check device tested Prepare the acrylic vacuum vessel for helium leak-checking. The helium leak checker is connected inline with the gas flow circuit of the vessel/pump system. Test the system by pumping the vessel down to 20 mbar and monitoring the level of helium present in the stream. This constitutes a background count for the helium leak-check of the source assembly. Release the vacuum.
- Note that, if at any point a leak in the source is found, then it will need to be re-sealed with a new o-ring, as per step 8, then re-leak-checked.
- Sniffer smells no helium 1 – Remove the source from the plastic tank. Use the sniffer to check for helium escaping from the o-ring seal of the Delrin, or the source connector. Next, move the source into the acrylic vacuum vessel. Pump the vessel down to 20 mbar and monitor the gas stream for helium. If no helium above background is observed in that time, then the source assembly will be deemed to not be leaking helium.
- Not leaking

- Mass of source 2 – Measure the source mass (accurate to 0.01 g). Fill a container with RO water, then submerge the source in the RO water. Place the container in the acrylic vacuum chamber and pump down to 20 mbar, watching for bubbles to emerge from either of the o-ring seals on either end of the stem. Air bubbles indicate a leak in the enclosure (and may point to the location of the leak).
- Observe for bubbles
- Source dried Return the vacuum chamber to atmospheric pressure and remove the source from the RO water. Dry with a Kimwipe and place the source on a clean Kimwipe on the lab bench (covered by another Kimwipe) and allow to air dry for 2 h. Re-measure its mass once the air dry is complete (a higher mass may indicate that RO water entered the source).
- Mass of source
- Not leaking If no bubbles were observed and the measured masses post- and pre-soak are consistent, then the enclosure will be deemed to not be leaking liquid.
- Soaked 4 hours 10) Once the source is shown to be leak-tight, a final soak test of the entire source is required to ensure the source fully encapsulates the internal radioactivity. Place the entire source in a water bath at 45 °C for 4 h. Collect a sample (~50 mL) of water from the bath into a screw-cap Nalgene container and count the water sample on the HPGe detector. Remove the source from the bath and dry using a Kimwipe. Leave on the lab bench on a clean Kimwipe (and covered with another Kimwipe) to air dry.
- Water sample collected
- Water counted If no contamination is present in the water after a count of 24 h, then the assembly of the source is deemed to have been successful. Place the enclosure into a clean baggie for storage.
- No contamination detected If the water sample shows signs of contamination, it is likely the surface of the container was contaminated during assembly. In this case, the procedure should be re-done completely with a new Delrin enclosure.
- Site monitored with Geiger counter 11) The lab area where the source was assembled should be tested for residual contamination. First, do a count of the area using the Geiger counter, noting discrepancies with the background measure in step 1. Wipe tests of the area are also required. A wipe test consists of wiping a filter paper over surfaces in the lab to collect any possible contamination. Each filter should cover an area of about 30cm². The papers should be placed in envelopes and sent to the UofA's Radiation Safety group to be counted in their scintillation counter. The surfaces to be wiped include the:
- Wipe tests complete
 - Lab bench top
 - Lab entry door
 - Floor

Mail the filter papers (multiple stapled inside each envelope) to:

Cory McHugh
 Radiation Protection Advisor, Department of Environmental Health & Safety
 Risk Management Services
 University of Alberta
 3107 Research Transition Facility
 Edmonton, AB T6G 2R5

Radiation Safety will notify you if contamination is found. Wipe tests are more sensitive than the Geiger counter, but the latter can alert you early on if there is contamination in the room. The Geiger counter must be used each day that radioactive work is done.

Completed by _____

⁶⁰Co source Delrin annealing procedure (March 23, 2015)

- Notes:** **All work will be done in CCIS L2-116**
- Project:** The Delrin used to build the ⁶⁰Co source container and stem must be annealed to relieve the container from any internal stresses induced during machining. Care must be taken to ensure flanges do not warp during the procedure and that the o-ring groove in the container maintains its shape. The melting point of Delrin is 347 °F.
- Materials:** Delrin container and stem
 Stem and container saddles
 Annealing oven (SOLO controller SL9696)
 Viton o-ring
 Gloves (vinyl or nitrile)
 Ultrasonic cleaner in L2-116 (with RO water and Alconox)
- Procedure:**
- Delrin cleaned 1) Don gloves. Clean the Delrin in the ultrasonic cleaner in a bath of RO water and Alconox for 30 minutes. Rinse with RO water. Dry with a Kimwipe then place parts on a clean Kimwipe on a bench and cover with an additional Kimwipe (to prevent dust from falling onto the surface). Allow to air dry for 2 hours.
- Stem and container in oven Anneal the Delrin. Saddles have been built for the stem and container so they remain horizontal (long side parallel to ground) during annealing, minimizing warping of the flange ends. Put the saddles with container and stem into the oven.
- Controller programmed Program the annealing cycle into the SOLO controller (following the manual).
- 2) Complete the annealing cycle:
- Heat oven to 300°F at a rate of 20°F per hour (12 hours total from room temperature)
 Hold the oven at 300°F for 30 min plus 15 minutes for every 1/8" thickness of Delrin.
 The thickest Delrin is in the stem at 10.88 mm (0.43") thick. Thus, hold the temperature at 300°F for 1 hour 25 minutes.
 Cool the oven to 150°F at a rate of 10°F per hour (15 hours total)
 Turn off oven and allow to come to room temperature before removing stem.
- Cycle begun The whole annealing cycle takes approximately 30 hours to complete.
- Cycle complete
- Delrin annealed Remove stem and container. Check to ensure the flanges are not warped (flat). Ensure the o-ring groove is round by checking it with the o-ring.

Completed by _____

B.3 Simulation

Section B.3.1 contains an example RAT macro that a user may utilize to produce a simulation of the ^{60}Co calibration source.

Table B.1 lists all of the `Co60Source` database fields. These are alterable in either a RAT macro or in the `Co60Source` database. Types are double and units are Bq, mm, MeV or degrees, unless otherwise specified. Values highlighted in boldface are those that a typical user may be required to modify.

B.3.1 Example RAT macro

```
#####  
#  
#   Co60Source.mac - Example Macro for the Co60 Source.  
#   Runs a 500 Bq Co60 source for 1 minute at z=100 mm.  
#  
#####  
  
/PhysicsList/OmitMuonicProcesses true  
/PhysicsList/OmitHadronicProcesses true  
  
/rat/db/load geo/Co60Source.geo  
  
/rat/db/set GEO[Co60Source] sample_position [0., 0., 100.]  
/rat/db/set GEO[Co60Source] ref_date ""  
/rat/db/set GEO[Co60Source] ref_activity 500.
```

```
/run/initialize
/rat/proc frontend
/rat/proc trigger
/rat/proc eventbuilder
/rat/proc calibratePMT
/rat/proc scintFitter
/rat/proc count
/rat/procset update 100
/rat/proclast outroot
/rat/procset file "Co60Source.root"

/generator/add co60source

/rat/run/duration 60 s
/rat/run/start

exit
```

B.3.2 Database fields

This table contains a list of all database fields relevant to both the simulation of the co calibration source in RAT and the running of the physical ^{60}Co calibration source in the SNO+ detector. All fields in boldface are those a user may need to change. Other fields pertain strictly to the simulation and are not typically changed by the user.

Field	Description	Default
version	Version of source	1.
sample_position [3]	Location in SNO+ of the centre of the scintillator button	[0., 0., 0.]
ref_date (string)	Date when source activity was calibrated	“18 Sep 2014 12:00:00”
ref_activity	Source activity on reference date	200.
ref_activity_err	Error in reference activity	1.
check_overlaps (bool)	Whether to check for volume overlaps when building source	1
mother (string)	Name of source mother volume	“inner_av”
container_radius	Outer radius of container base	23.7
container_height	Height of container	58.0
container_thickness	Wall thickness of container	2.0
container_collar_ height	Thickness of internal collar inside container	10.0
container_collar_ hole_width	Side length of square hole in collar	25.0
container_collar_ hole_rad	Size of holes in collar in collar	14.0
container_slope_ height	Height of tapered container section	6.17
container_flange_ radius	Radius of container flange	34.39
container_flange_ thickness	Thickness of container flange	9.0
container_flange_ base_thickness	Thickness of flange base below nut groove	1.5
container_screw_ hole_radius	Radius of screw hole in flange	1.45
container_nut_ groove_height	Height of nut groove (to fit nut thickness)	4.0
container_nut_ groove_width	Width of nut groove (to fit nut radius)	7.35
container_gap_ copper	Gap between copper box and container	0.5
container_material (string)	Material making up container	“G4_ POLYOXY METHYLENE”
container_colour [4]	RGB α of container colour	[0, 0, 0, 0.5]
⋮	⋮	⋮

Field	Description	Default
stem_flange_thickness	Thickness of flange on stem	5.0
stem_screw_head_hole_radius	Radius of hole to hold screw head	2.35
stem_screw_head_hole_depth	Depth of hole to hold screw head	2.90
connect_radius	Radius of connector end of stem	29.0
connect_thickness	Thickness of connector end of stem	9.53
bore_radius	Thickness of bore through stem centre	5.0
stem_flange_end_radius	Radius of stem at flange connection	11.0
stem_flange_end_length	Length of stem section at flange connection	8.7376
stem_connect_end_radius	Radius of stem at connector connection	15.88
stem_length	Length of stem, including flanges	309.74
stem_taper_angle	Angle of tapered stem section away from stem axis	3.0
stem_material (string)	Material making up stem	"G4_POLYOXY METHYLENE"
stem_colour [4]	RGB α of stem colour	[0, 0, 0, .5]
oring_groove_height	Depth of o-ring groove	1.27
oring_groove_width	Width of o-ring groove	2.26
oring_groove_inner_radius	Inner radius of o-ring groove	25.26
oring_material (string)	Material making up o-ring	"G4_VITON"
oring_color [4]	RGB α of o-ring colour	[1, 0, 0, .5]
copper_material (string)	Material making up copper	"G4_Cu"
copper_colour [4]	RGB α of copper colour	[1, 0, 1, .5]
indium_material (string)	Material making up indium wire	"G4_In"
indium_colour [4]	RGB α of indium colour	[0, .5, 1, .5]
glass_material (string)	Material making up glass feedthrough	"glass"
glass_colour	RGB α of glass colour	[0, .5, 1, .5]
:	:	:

Field	Description	Default
copper_gap	Gap between copper box and container	0.5
copper_height	Height of copper box	67.5
copper_width	Side length of copper box	23.75
copper_thickness	Thickness of copper box	0.0127
copper_flange_rad	Radius of copper flange	21.615
copper_flange_height	Thickness of copper flange	2.5
copper_flange_lip_height	Height of copper flange bottom lip	5.0
copper_flange_lip_width	Side length of copper flange bottom lip	23.49
copper_flange_lip_thickness	Thickness of copper flange bottom lip	1.0
copper_glass_rad	Radius of glass feedthrough	6.65
copper_metal_rad	Outer radius of metal ring around glass	7.25
copper_glass_height	Thickness of glass feedthrough	3.0
copper_oring_inner_rad	Inner radius of copper protrusion	15.43
copper_oring_outer_rad	Outer radius of copper protrusion	15.83
indium_depth_bottom	Depth of indium groove	0.75
indium_depth_top	Height of copper protrusion	0.5
screws_enable (bool)	Turn on screws/nuts/holes	1
number_of_screws (int)	Number of screws in container flange	8
screw_distance_from_center	Distance from source centre to screw centre	30.715
screw_head_radius	Radius of screw head	2.27
screw_radius	Radius of screw body	1.415
screw_head_length	Length of screw head	2.8
screw_length	Length of screw, including head	15.0
screw_material (string)	Material making up screws	“stainless steel”
screw_colour [4]	RGB α of screw colour	[.5, .5, .5, .5]
:	:	:

Field	Description	Default
nut_radius	Outer radius of nut	3.545
nut_insert_thickness	Thickness of nylon nut insert	1.0
nut_thickness	Thickness of nut	3.78
nut_material (string)	Material making up nut	“stainless_steel”
nut_insert_material (string)	Material making up nut insert	“nylon ”
nut_colour [4]	RGB α of nut colour	[.5, .5, .5, .5]
nut_insert_colour [4]	RGB α of nut insert colour	[1, 1, 1, .5]
number_of_screws_collar	Number of screws in collar	4
number_of_screws_copper	Number of screws in copper flange	8
screw_distance_from_centre_collar	Distance to collar screws	23.49
screw_head_radius_collar	Radius of collar screw head	2.27
screw_radius_collar	Radius of collar screw body	1.415
screw_head_length_collar	Length of collar screw head	2.8
screw_length_collar	Length of collar screw, including head	12.0
screw_length_copper	Length of copper screw, including head	12.0
nut_radius_collar	Outer radius of collar nut	3.545
nut_insert_thickness_collar	Thickness of collar nylon nut insert	1.0
nut_thickness_collar	Thickness of collar nut	3.78
quick_connect_radius	Radius of quick connect	46.0
quick_connect_thickness	Thickness of quick connect	9.5
quick_connect_height	Height of quick connect	123.5
quick_connect_plate_thickness	Thickness of plate in quick connect	5.0
quick_connect_plate_height	Height of plate in quick connect	50.0
quick_connect_material	Material making up the quick connect	“stainless_steel”
quick_connect_colour [4]	RGB α of quick connect colour	[.5, .5, .5, 1.]
⋮	⋮	⋮

Field	Description	Default
pmt_window_radius	Radius of PMT window	5.0
pmt_active_radius	Radius of active part of PMT window	4.0
pmt_window_inset	Inset of PMT face from body face	1.5
pmt_length	Length of PMT body	50.0
pmt_face_length	Length of side of square PMT face	22.0
pmt_material (string)	Material making up PMT body	“aluminum”
pmt_active_material (string)	Material making up PMT face	“glass”
pmt_colour [4]	RGB α of PMT colour	[0, 0, 1, .5]
scintillator_radius	Radius of scintillator button	4.0
scintillator_thickness	Thickness of scintillator button	4.0
scintillator_material (string)	Material making up scintillator button	“G4_PLASTIC_SC_VINYLTOLUENE”
scintillator_colour [4]	RGB α of scintillator button colour	[0, 1, 1, .5]
sensitive_detector (string)	Name of sensitive detector type	“/mydet/pmt/calib”
lcn (int)	Logical channel number of source PMT	9191
source_efficiency	Absolute efficiency of source at observing a ^{60}Co decay	0.9
threshold_energy	Energy above which source efficiency is non-zero	0.0

Table B.1: Fields of the ^{60}Co source geometry, alterable in a RAT macro or the GEO database under index `Co60Source`. Values a typical user may need to change are in boldface. Changing other values may lead to failures when building the geometry.