

Probing new physics mechanisms in neutrinoless double-beta decay with SNO+

Ashley Robert Back

School of Physics and Astronomy
Queen Mary, University of London

Submitted in partial fulfilment of the requirements of the Degree of Doctor of Philosophy.

2017

Declaration

I, Ashley Robert Back, confirm that the research included within this thesis is my own work or that where it has been carried out in collaboration with, or supported by others, that this is duly acknowledged below and my contribution indicated. Previously published material is also acknowledged below.

I attest that I have exercised reasonable care to ensure that the work is original, and does not to the best of my knowledge break any UK law, infringe any third party's copyright or other Intellectual Property Right, or contain any confidential material.

I accept that the College has the right to use plagiarism detection software to check the electronic version of the thesis.

I confirm that this thesis has not been previously submitted for the award of a degree by this or any other university.

The copyright of this thesis rests with the author and no quotation from it or information derived from it may be published without the prior written consent of the author.

Signature: Ashley R. Back

Date: 2016-09-22

Details of collaboration and publications:

Andringa, S. & others, "Current Status and Future Prospects of the SNO+ Experiment", *Adv. High Energy Phys.*, **2016**, 6194250 (2015)

Acknowledgements

Many people have been part of the journey towards the completion of this thesis, but I only have limited space, so I apologise in advance if you are reading this and I have not thanked you personally.

First, I would like to thank my supervisors and advisers to the project. Francesca, thank you for all your help and for always making time for me when I needed it. Jeff, your work ethic has been inspiring; thank you for being patient. Frank, thank you for all your advice on the project and for helping with the theory.

I was fortunate to call a few places “home”, throughout my PhD. I would like to thank everyone at Queen Mary, University of London, for their help, particularly Jeanne, Matt, Phil and Billy. Thank you to everyone at Sussex, for making me feel so welcome, while I spent nine months there. Particularly thanks to Simon, Lisa, Ed and Mark. Finally thank you to everyone in Sudbury, who made my three months stay so enjoyable, particularly to Erica for helping host a fantastic party in celebration of 50 years of the Doctor!

Thank you to all the other SNO+ UK students, Rob, John, Chris and Jack.

Thank you to all my friends and housemates, particularly Alex, Jenny, James and the Nolans.

To Christina, thank you for your help proof reading various sections of this thesis, as I completed them.

Special thanks must go to Steffi for your endless German efficiency, particularly at lunch and tea breaks.

To James, thank you for the enormous amount of work you have put into echidna, I could never have done it without you.

To Evelina, it has been a privilege to share so many experience with you, from meeting bears to painting mine walls and mastering GitHub! All I can say is “good job!”.

Finally, to my mum, dad and sister, thank you for your love and support, I honestly could not have done it without you.

Contents

List of figures	11
List of tables	17
1 Neutrinos and neutrinoless double-beta decay	21
1.1 The Standard Model neutrino	21
1.2 Neutrinos beyond the Standard Model	24
1.2.1 Motivation for neutrinos beyond the Standard Model	24
1.2.2 Adding Dirac neutrino masses to the Standard Model	29
1.2.3 The Majorana mass term	31
1.3 Probing Majorana neutrinos	32
1.4 Experimental status of neutrino mass searches	37
1.5 Majoron-emitting modes of $0\nu 2\beta$	38
1.6 Experimental searches for Majoron-emitting modes	42
2 The SNO+ detector	47
2.1 Overview of the detector	47
2.2 Detector upgrades	50
2.2.1 Electronics upgrades	50
2.2.2 photomultiplier tube (PMT) repairs	50
2.3 $0\nu 2\beta$ in SNO+	52
2.3.1 Tellurium loading	52
2.3.2 Background considerations	53
2.3.3 Current sensitivity estimates	53
2.4 SNO+ physics program	56
2.4.1 Water-phase physics	56
2.4.2 Solar neutrinos	56
2.4.3 Reactors, supernovae and geoneutrinos	57

2.5	Detection and data acquisition	60
2.5.1	A hit	60
2.5.2	An event	61
2.5.3	A run	64
3	High-level checks for SNO+ data quality	65
3.1	Overview of the software	66
3.1.1	The parent processor	66
3.1.2	General format of the processors	69
3.2	The run processor	69
3.3	The PMT processor	76
3.4	The timing processor	81
3.5	The trigger processor	86
3.6	Usage of the processors	88
3.7	HLDQ nearline	88
3.8	Conclusion	89
4	Setting limits and the echidna software package	91
4.1	Formalism	91
4.2	Overview of echidna	93
4.3	Spectra creation and manipulation	94
4.3.1	The Spectra class	94
4.3.2	Creation and configuration	95
4.3.3	Filling spectra	96
4.3.4	Spectra scaling and integration	97
4.3.5	Convolutions	98
4.3.6	Data storage	100
4.4	The Poisson likelihood chi-squared	100
4.4.1	Likelihood	102
4.4.2	Extended likelihood	102
4.4.3	Binned maximum likelihood	103
4.4.4	The Poisson likelihood χ^2	105
4.5	A fixed background fit	107
4.6	Investigating signals	110
4.7	Floating parameters in the fit	113
4.8	Minimisation	114

5	Verification of echidna as a limit setting package	117
5.1	Overview of the KamLAND-Zen and its results	117
5.2	Digitising the KamLAND-Zen spectra	119
5.3	Methodology for verification study	121
5.4	Estimating digitisation uncertainty	122
5.5	Measurement of the double-beta rate	125
5.6	Setting limits on Majoron-emitting neutrinoless double-beta decay ($0\nu 2\beta$) modes	125
5.7	Comparison of limits	127
6	Estimating SNO+ sensitivity to Majoron-emitting modes	135
6.1	General assumptions	135
6.2	Backgrounds	136
6.2.1	^{130}Te double-beta decay	136
6.2.2	Solar neutrinos	138
6.2.3	Internal backgrounds	138
6.2.4	AV and leaching backgrounds	139
6.2.5	α, n backgrounds	140
6.2.6	External backgrounds	140
6.3	Energy systematics	141
6.4	A fixed background fit	142
6.5	Floating the double-beta rate	145
6.6	Floating the detector energy resolution	149
6.7	Floating other systematics	153
6.8	Discussion	155
7	Conclusions	163
	Bibliography	167

List of figures

1.1	A comparison of the fluxes measured by each of the solar neutrino experiments, with the prediction based on the Standard Solar Model and electroweak theory.	28
1.2	Energy-level diagram showing the binding energy, relative to the most bound nucleus, for isotopes where $A = 136$. Green bands represent even-even nuclei, whilst red bands mark odd-odd nuclei.	33
1.3	Feynman diagram showing $\nu_e \rightarrow \bar{\nu}_e$ transition, through the effective “black box” operator for neutrinoless double-beta decay.	34
1.4	Feynman diagram showing neutrinoless double-beta decay via the exchange of a light Majorana neutrino.	35
1.5	The phase space for $0\nu 2\beta$. Plots the value of $m_{\beta\beta}$ as a function of the lightest neutrino mass. The plot indicates the quasi-degenerate region (QD) and allowed regions for the NH and IH, as well as limits from the $0\nu 2\beta$ experiments and cosmology.	36
1.6	Feynman diagram showing neutrinoless double-beta decay via the exchange of a light Majorana neutrino, with the emission of two Majorons.	39
1.7	Feynman diagram showing neutrinoless double-beta decay via the exchange of a light Majorana neutrino, with the emission of a single Majoron.	40
1.8	Expected spectral shapes for spectral index $n = 1, 2, 3, 5, 7$ for the different Majoron-emitting modes of $0\nu 2\beta$ in ^{136}Xe . The spectra have been convolved using KamLAND-Zen’s detector response function, including energy resolution and energy-scale non-linearities. The resolution-limited line at the Q -value represents $0\nu 2\beta$ without Majoron emission.	42

2.1	Wire-frame diagram of the SNO+ detector, showing the cavity, PSUP (<i>dark green</i>), hold-up ropes (<i>magenta</i>), hold-down ropes (<i>red</i>) and AV (<i>blue</i>).	48
2.2	(a) The set-up used for dark-box testing. (b) Cleaning the lower part of PMT being in mineral solution. (c) Potting the hub with the two-part SilGel mix.	51
2.3	Stacked histogram spectral plot, for 0.5 % loading using the Te-diol method (390 NHit/MeV). Includes the main background contributions and an example signal scaled to $m_{\beta\beta} = 200$ meV.	54
2.4	Expected fluxes for the four main solar neutrino signals and background spectra (scaled to target levels), with a 5.5 m fiducial volume.	58
2.5	Expected visible anti-neutrino energy spectrum, where the Terrestrial Neutrino Unit (TNU) is a measure of one inverse beta decay per 10^{32} protons. Including stacked (Bruce—blue, Pickering and Darlington—red—and others—yellow) oscillated spectrum, non-oscillated (dashed) and geoneutrino spectrum(solid)	59
2.6	Concise overview of the SNO trigger system	61
2.7	An overview of the SNO timing system.	63
3.1	Flow diagram showing the high-level DQ (HLDQ) processors in the two processing and how they fit in relation to the other FEC processes	66
3.2	Summary of status of the HLDQ checks. Red shading identifies the applied checks, the blue histogram shows the check result (pass/fail), for run 8614.	68
3.3	Crate coverage maps showing the occupancy of (<i>left</i>) all hit PMTs and (<i>right</i>) the occupancy of channels where calibration always succeeded (status 3 in Table 3.3)—for run 8614	79
3.4	Coverage map plotting the θ and ϕ coordinates of usable PMTs—for run 8614.	81
3.5	Distributions showing (<i>left</i>) all positive Δt and (<i>right</i>) positive Δts that are within or slightly outside of the trigger window—for run 8614.	85
3.6	Histogram showing the calculated rate for each trigger signal in the trigger word—for run 8614.	87

-
- 4.1 Spectral plot, showing the spectral shapes of the Poisson background, flat background and signal, that we have simulated to use in the toy model. . . . 92
- 4.2 Overlaid histograms showing the effect of the four convolutions we have implemented in echidna. In each case the input is the blue curve and the convolved spectrum is the red curve. The convolutions are: a. Gaussian smearing, with mean energy 2.4 MeV and energy resolution of 5 %; b. Poisson smearing, also with mean energy 2.4 MeV and energy resolution of 5 %; c. Shift by -0.2 MeV; d. Scale by factor of 0.97. 101
- 4.3 Spectral plot showing the “data” generated for the toy model, as well as the Poisson and flat backgrounds, in the 2.3 to 2.7 MeV ROI 107
- 4.4 A comparison of hypothetical χ^2_λ curves for three different scenarios. A sensitivity study (*red*) estimates an upper limit; a fit to data that is consistent with H_0 (*blue*), gives an upper limit on the signal scale; and a fit to data that significantly rejects H_0 (*green*), calculates a two-sided confidence interval for the signal scale. 111
- 4.5 χ^2_λ curve showing the limit calculated for the fixed background fit, to the toy model data. 112
- 5.1 Visible energy spectra for the different Majoron-emitting modes of $0\nu 2\beta$ in KamLAND-Zen. The spectra have been convolved using a detector response function, including energy resolution and energy-scale non-linearities. The resolution-limited line at the Q -value represents $0\nu 2\beta$ without Majoron emission. 119
- 5.2 The data points show the visible energy spectrum of candidate events. The curves show the combined best-fit background spectrum (grey dashed line), $2\nu 2\beta$ spectrum (purple) and 90 % confidence upper limits on $0\nu 2\beta$ and Majoron-emitting modes, for spectral indices $n = 1, 2, 3, 7$. The red line shows the sum of the $2\nu 2\beta$ and best-fit background spectra. The best-fit has $\chi^2/\text{d.o.f} = 100.4/87$ 120
- 5.3 Spectral plot using the digitised spectra to reproduce Figure 5.2. 121

- 5.4 Spectral plot, showing the visible energy for the KamLAND-Zen dataset of candidate events, along with the dominant $2\nu 2\beta$ background contribution (magenta) and the combined contribution from all other radioactive backgrounds (grey). The red curve is the best fit summed spectrum, assuming no additional signal contributions. The light red shaded region shows a $\pm 1\sigma$ Poisson error band. The lower panel plots the ratio of the deviation of the data points from the best fit curve compared to $\pm 1\sigma$ error bands, MC samples with a factors of 5, 10 and 20 times the exposure of the data. 123
- 5.5 Spectral plot, showing the visible energy for the KamLAND-Zen dataset of candidate events, along with the dominant $2\nu 2\beta$ background contribution (magenta) and the combined contribution from all other radioactive backgrounds (grey). The red curve is the best fit summed spectrum, assuming no additional signal contributions. The light red shaded region shows a $\pm 1\sigma$ Poisson error band. The lower panel plots the ratio of the deviation of the data points from the best fit curve compared to $\pm 1\sigma$ error bands (blue). The red band on the lower panel represents the size of the digitisation uncertainty in relation to the Poisson error band. 129
- 5.6 Spectral plot showing the fixed best-fit of the digitised spectra to the digitised data array. Underneath we include the residuals from MC, in units of σ . . . 130
- 5.7 Quartet of χ^2 curves for the KamLAND-Zen fit, with floating $2\nu 2\beta$ rate and floating combined background, showing χ_λ^2 as a function of signal decays. With (*top-left* to *bottom-right*) spectral indices $n = 1, 2, 3, 7$. The dashed lines mark the position of the 90 % confidence upper limit on the rate. 131
- 5.8 Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the $2\nu 2\beta$ rate and signal scale axes, with (*top-left* to *bottom-right*) spectral indices $n = 1, 2, 3, 7$ 132
- 5.9 Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the combined background rate and signal scale axes, with (*top-left* to *bottom-right*) spectral indices $n = 1, 2, 3, 7$ 133

- 5.10 Comparison of echidna's limits for the digitised spectra, with KamLAND-Zen's limits. The dashed line at 1 signifies the KamLAND-Zen, limit. The purple lines show the ratio of echidna's limit divided by the corresponding KamLAND-Zen limit, with and shaded purple band representing the statistical uncertainty on the calculated limit. The red lines represent the corresponding sensitivity estimate for each mode, calculated using echidna 134
- 6.1 Spectral plot, showing contributions to the fixed background spectrum, in the energy range 1.0 to 3.5 MeV. 142
- 6.2 Quartet of χ^2 curves for the fixed background fit, showing χ_λ^2 as a function of signal decays, with (*top-left* to *bottom-right*) $n = 1, 2, 3, 7$. The dashed lines mark the position of the 90 % confidence upper limit. 144
- 6.3 Spectral plot showing the fixed background spectrum and $n = 1$ signal spectrum, at the 90 % confidence limit. Underneath the histogram shows the contribution from each bin to the full sensitivity. 145
- 6.4 Spectral plot showing the fixed background spectrum and $n = 2$ signal spectrum, at the 90 % confidence limit. Underneath the histogram shows the contribution from each bin to the full sensitivity. 146
- 6.5 Spectral plot showing the fixed background spectrum and $n = 3$ signal spectrum, at the 90 % confidence limit. Underneath the histogram shows the contribution from each bin to the full sensitivity. 146
- 6.6 Spectral plot showing the fixed background spectrum and $n = 7$ signal spectrum, at the 90 % confidence limit. Underneath the histogram shows the contribution from each bin to the full sensitivity. 147
- 6.7 Spectral plot, showing the four signals scaled to their limits, and the floated $2\nu 2\beta$ spectrum at its best fit value. 148
- 6.8 Quartet of χ^2 curves for the fixed background fit, with floating $2\nu 2\beta$ rate, showing χ_λ^2 as a function of signal decays. With (*top-left* to *bottom-right*) $n = 1, 2, 3, 7$. The dashed lines mark the position of the 90 % confidence upper limit. 150
- 6.9 Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the $2\nu 2\beta$ rate and signal scale axes, with (*top-left* to *bottom-right*) $n = 1, 2, 3, 7$ 151

-
- 6.10 Spectral plot showing the four fitted Majoron spectra, scaled to the 90% confidence upper limit, that we set. Also shown is ^{130}Te $2\nu 2\beta$ background, that we float in the fit. We scale this to its best-fit rate. We also show all floating spectra convolved to their best-fit energy resolution. 154
- 6.11 Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the $2\nu 2\beta$ rate and signal scale axes. With (*top-left* to *bottom-right*) $n = 1, 2, 3, 7$ 155
- 6.12 Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the energy resolution and signal scale axes. With (*top-left* to *bottom-right*) $n = 1, 2, 3, 7$. 156
- 6.13 Quartet of χ^2 curves for the fixed background fit, with floating $2\nu 2\beta$ rate and energy resolution, showing χ_λ^2 as a function of signal decays. With (*top-left* to *bottom-right*) $n = 1, 2, 3, 7$. The dashed lines mark the position of the 90% confidence upper limit. 157
- 6.14 Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the energy scale and signal scale axes. With (*top-left* to *bottom-right*) $n = 1, 2, 3, 7$ 158
- 6.15 Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the energy shift and signal scale axes. With (*top-left* to *bottom-right*) $n = 1, 2, 3, 7$ 159
- 6.16 Spectral plot showing the four fitted Majoron spectra, scaled to the 90% confidence upper limit, that we set. Also shown is ^{130}Te $2\nu 2\beta$ background, that we float in the fit. We scale this to its best-fit rate. We also show all floating spectra convolved to their best-fit energy resolution. 161

List of tables

1.1	Summary of the properties of common $0\nu 2\beta$ isotopes.	37
1.2	Summary of current experimental 90 % confidence upper limits on $0\nu 2\beta$, via the MM. We calculate a range of values for $m_{\beta\beta}$, based on the range of values for the NME, as Table 1.1 outlines.	39
1.3	Summary of Majoron-emitting modes of $0\nu 2\beta$	41
1.4	Nuclear information for Majoron-emitting modes of $0\nu 2\beta$	43
1.5	Summary of current limits on $T_{1/2}^{0\nu 2\beta \chi^0(\chi^0)}$ for Majoron emitting modes of $0\nu 2\beta$, for all modes listed in Table 1.3.	45
2.1	Comparison of the $0\nu 2\beta$ activity for three different $^{\text{Nat}}\text{Te}$ loading scenarios: 0.3 %, 0.5 % and 5.0 % $^{\text{Nat}}\text{Te}$, by mass.	55
2.2	Expected precision of flux measurements, for a global fit to the solar neutrino spectrum.	57
3.1	Bit-wise description of the run-type bitmask.	71
3.2	Bit-wise description of the trigType bit-mask.	72
3.3	Definitions of the four statuses we use to identify usable channels	76
3.4	Summary of the criteria HLDQ applies for a standard physics run	90
5.1	The KamLAND-Zen limits on the Majoron-emitting $0\nu 2\beta$ half-lives $\left(T_{1/2}^{0\nu 2\beta \chi^0(\chi^0)}\right)$ and effective Majoron-neutrino coupling constant g_{ee} , at 90 % C.L., for ^{136}Xe . The model notation follows the same form as Table 1.3. The third, fourth and fifth columns indicate whether the mode is a Goldstone boson, its leptonic charge L and its spectral index n	118

6.1	Summary of the expected background contributions, over 5 y, from the α, n backgrounds and corresponding delayed 2223 keV- γ	140
6.2	Summary of the number of counts in the ROI 1.0 to 3.5 MeV, for the main (grouped) background contributions.	141
6.3	Summary of the estimated 90 % confidence lower limits on the half life ($T_{1/2}^{0\nu 2\beta\chi^0(x^0)}$), for the fixed background fit.	143
6.4	Summary of 90 % confidence limits for the fit floating the two-neutrino double-beta decay ($2\nu 2\beta$) rate.	149
6.5	Summary of the signal parameter values used in the fit.	152
6.6	Summary of 90 % confidence limits for the fit floating the $2\nu 2\beta$ rate.	153
6.7	Summary of limits.	160
6.8	Table summarising the best fit values of the systematic uncertainties.	160

Abstract

In this thesis, I present the theory of neutrinoless double-beta decay ($0\nu 2\beta$), particularly the theory of exotic modes of $0\nu 2\beta$ involving the emission of one or two Majorons. Alongside this, I summarise the most recent results in the experimental search for $0\nu 2\beta$, including limits on the rate of these exotic processes. I describe the SNO+ experiment and its physics goals, which include the search for $0\nu 2\beta$. As part of the SNO+ collaboration, I have made a significant contribution towards the development of the data quality software that is essential for ensuring SNO+ can achieve its physics goals—including in $0\nu 2\beta$ searches. I describe how I developed a software package that performs the high-level data quality checks.

Continuing with the software theme, I then describe a python-based limit-setting and fitting software package called echidna. I have been a lead developer of echidna as part of my PhD, so I describe the software in detail and how it can be used to set limits on $0\nu 2\beta$ signals. By reproducing the sensitivity results of the KamLAND-Zen experiment, in four key Majoron-emitting $0\nu 2\beta$ modes, I verify the use of echidna as a limit-setting tool for this type of search. Finally, I present the results of a comprehensive sensitivity study, where I determine the potential sensitivity of SNO+ to the same set of Majoron-emitting modes that KamLAND-Zen and other $0\nu 2\beta$ experiments have already investigated.

Chapter 1

Neutrinos and neutrinoless double-beta decay

1.1 The Standard Model neutrino

When presenting the historical context of neutrino physics, authors often begin with W. Pauli’s famous address, to a physics conference in Tübingen, where he began his letter “Dear Radioactive Ladies and Gentlemen” [1]. Yet, our story begins with the discovery of radioactivity in Uranium. By placing metallic salts in a dark box, with a further layer of shielding outside, and then exposing a photographic plate to the salts, inside the box, H. Becquerel demonstrated that radiation was coming from the salts [2]. He concluded that the radiation was due to the presence of Uranium in the salts, and began investigating uranium-rich salts as well as metallic uranium. Using a gold-leaf electrometer, Becquerel studied the electric discharge of the radiation, concluding that uranium exhibited an “invisible phosphorescence”. In 1899, E. Rutherford proved there were two distinct types of radioactive decay— α and β^1 [3]. Decays via α and β have a fixed energy loss, but β -emissions, as J. Chadwick highlighted in 1914 [4], follow a continuous visible energy spectrum. This was a surprising result because it seemed to violate energy conservation. The consensus was that the process emitted a single β particle, so if this did not always carry the total (fixed) energy, released by the decaying nucleus, then there was some “missing energy”.

Pauli’s solution to this problem, was to introduce a new weakly-interacting, neutral fermion, which he called a *neutron* [1]. With a new particle, the two decay products share the energy released by the nucleus, conserving total energy, but since the new particle is

¹We now know that high-energy photons (γ) are a third type.

neutral, ionisation detectors cannot detect its energy contribution. Pauli proposed an upper limit on the mass of his *neutron*, at 0.01 proton masses, but conceded that one should have detected these *neutrons*, if they exist. E. Fermi renamed Pauli's particle to the neutrino, after the discovery of what we now know as the neutron, by J. Chadwick, in 1932 [5]. In the first published reference to the neutrino [6], Fermi concluded that it could be massless; a conclusion F. Perrin also reached independently [7].

Confirmation of the existence of Pauli's particle required an observation of the neutrino interacting with matter, at a location separate from its source. With this goal in mind, the approach taken by F. Reines and C. Cowan, was to try and detect antineutrino emissions from a nuclear reactor, via inverse beta decay ($\bar{\nu}_e + p \rightarrow n + e^+$) [8]. Two tanks containing 200 l² water, provided the proton target, and also contained up to 40 kg of CdCl₃. The experimental signature was a prompt pulse from the positron annihilation and a coincident pulse, some microseconds later, due to γ emissions, from a neutron capture in cadmium. Reines and Cowan placed the water tanks between three tanks of liquid scintillator, monitored by 55 photomultiplier tubes (PMTs), to record pulses from the γ emissions. The source of potential antineutrinos, was radiation from the Savannah River nuclear reactor, 11 m from the detector. With 100 live days of data, Reines and Cowan measured the cross section of $\bar{\nu}_e$ on protons, as $\sigma_{exp} = 12_{-4}^{+7} \times 10^{-44} \text{ cm}^2$, based on a rate of 3.0 ± 0.2 events per hour [9]. This compares favourably with predicted cross sections $\sigma_{th} = (5 \pm 1) \times 10^{-44} \text{ cm}^2$. Note the predicted cross section at the time, was before a factor of two increase due to parity non-conservation in the handedness of the neutrino. Reines and Cowan present a convincing argument for the observation of the free antineutrino, supporting their result with measurements proving that: the first pulse was due to positron annihilation, the second pulse was due to neutron capture in cadmium and that the signal rate was proportional to the number of target protons. They also ruled out any other known fission product, as the source of the signal, by showing that the measured rate remained constant with shielding between the reactor and detector.

Another experiment aiming to detect the free $\bar{\nu}_e$ produced by beta emissions, from a nuclear reactor, was an experiment conducted at Brookhaven, by R. Davis [10]. Davis sought to detect neutrinos, by inverting the electron capture decay of ³⁷Ar, the reaction:



Whether neutrinos were distinct from their antiparticles, was unclear at the time, so Davis was experimenting to see if an *antineutrino* source, might induce the *neutrino* interaction in

²We use l to represent the unit of one litre.

Equation (1.1). Davis used a 3900 l tank of CCl_4 , exposed to radiation from the Brookhaven reactor, as well as a 200 l tank to measure the cosmic ray background. Since, ^{37}Ar is inert, Davis calculated the activity using a novel method of extracting and counting ^{37}Ar , after a period of exposure. The number of argon nuclei observed, was not enough to suggest an observation of antineutrino capture in ^{37}Cl , but Davis was able to set a limit of $\sigma = 2 \times 10^{-42} \text{ cm}^2$, on the cross section of the reaction. One can attribute the failure to observe the process, to the fact that the neutrino and antineutrino interact differently and that Lepton number is a conserved quantity, as E. J. Konopinski and H. M. Mahmoud proposed in 1953 [11].

Also in Brookhaven, in 1962, L. M. Lederman, M. Schwartz and J. Steinberger, performed the first neutrino beam experiment [12]. They directed a beam of protons at a beryllium target, causing a shower of decaying pions and kaons. The pions decay into neutrinos. They directed the beam at 13.5 m iron shielding, which would stop all decay products except the neutrinos. The aim of the experiment was to see if the beam induced electrons in a detector 21 m from the base of the iron shielding. The experiment saw no events, in the detector, that were consistent with ν_e interactions, confirming that:

1. Pions decay into an anti-muon, accompanied by a muon-type neutrino (Equation (1.2)), but cannot decay into an electron-type neutrino (Equation (1.3)):

$$\pi^+ \rightarrow \mu^+ + \nu_\mu, \quad (1.2)$$

$$\pi^+ \rightarrow \mu^+ + \nu_e. \quad (1.3)$$

2. Neutrons cannot capture muon-type neutrinos (Equation (1.5)), but can capture electron neutrinos, via inverse beta decay (Equation (1.4)):

$$\nu_\mu + n \rightarrow p + e^-, \quad (1.4)$$

$$\nu_e + n \rightarrow p + e^-. \quad (1.5)$$

These conclusions confirmed the existence of the muon neutrino and that Lepton flavour was also a conserved quantity. The DONUT collaboration discovered the tau neutrino in 2000 [13].

The Glashow-Weinberg-Salam Standard Model (SM) [14, 15, 16] provides a description of the weak interaction and beta decay. The goal of the model was to unify the electromagnetic and weak interactions, in a single local gauge group. The simplest gauge group that achieves this, is $SU_L(2) \times U_Y(1)$, where Y is the weak hypercharge. The SM contains three generations.

In the lepton sector: the electron, muon and tau, and their accompanying neutrinos ν_e , ν_μ and ν_τ . There are also three generations in the quark sector, with six quarks in total. All twelve fermions have a corresponding antiparticle. In the lepton sector, the SM requires conservation of Lepton number and also Lepton flavour.

The Brout-Englert-Higgs (BEH) mechanism generates fermion masses in the SM. In this mechanism, when one breaks electroweak symmetry, the Higgs field acquires a vacuum expectation value $v = 246$ GeV. This defines the scale of the fermion masses. The mechanism requires a doublet of complex Higgs fields—four degrees of freedom—predicting a single, neutral Higgs boson. Data from the Large Hadron Collider (LHC), agrees well with this prediction. The Yukawa interactions, under $SU_L(2) \times U_Y(1)$, generate Dirac mass terms, which generate the masses of the charged leptons and quarks. Since, unification of the electromagnetic and weak interactions, does not require right-handed neutrino fields, neutrinos do not develop a Dirac mass term. After spontaneous, electroweak symmetry breaking, neutrinos remain massless, two-component, Weyl particles [17].

1.2 Neutrinos beyond the Standard Model

1.2.1 Motivation for neutrinos beyond the Standard Model

Despite his failure to observe free neutrinos in his Brookhaven experiment (as we saw in Section 1.1), R. Davis noted the potential to study electron neutrinos from the Sun, using the interaction in Equation (1.1). The Standard Solar Model (SSM), gave a clear theory of the series of interactions that regulate the Sun's power output, and accurately calculated the flux of neutrinos released by these interactions [18]. J. Bahcall et al. noted that Davis' experiment should be sensitive to the high-energy part of the solar neutrino spectrum—produced by the following reactions:

$${}^8\text{B} \rightarrow {}^8\text{B}^* + e^- + \nu_e \quad (1.6)$$

$$e^- + {}^7\text{Be} \rightarrow {}^7\text{Li} + \nu_e \quad (1.7)$$

$$(1.8)$$

Initial results from Davis' experiment, now buried underground ((4200 ± 100) mwe) in Homestake Mine, to reduce cosmic ray muons, showed a solar neutrino flux below the prediction of the SSM [19]. The Homestake experiment collected 108 further measurements of the solar

neutrino flux, over a period of 24 years, to determine an average rate. The measured rate was less than one third of the rate predicted by the SSM [20]—a discrepancy of more than 3σ , after considering statistical and systematic uncertainties. This discrepancy in the solar neutrino flux gained the title “the Solar Neutrino Problem”.

Atmospheric neutrino experiments saw a similar problem—known as the “atmospheric neutrino anomaly”. These experiments measure neutrinos produced by the $\pi \rightarrow \mu^-$ decay, in the upper atmosphere, which yields two muon neutrinos and one electron neutrino, so the expectation was for double the flux of ν_μ , compared to ν_e . The convention was to quote the double ratio:

$$R = \frac{(N_\mu/N_e)_{data}}{(N_\mu/N_e)_{model}}, \quad (1.9)$$

where N_μ is the number of muon neutrinos and N_e the number of electron neutrinos [21]. One considers any value of R , straying considerably from unity, an anomalous result. Kamiokande, a water Cherenkov detector, containing over 2 kt ultra-pure water (UPW), measured values for the ratio that suggested more equal fluxes of ν_μ and ν_e . As did the 8 kt Irvine-Michigan-Brookhaven, water Cherenkov experiment.

B. Pontecorvo proposed a solution to the Solar Neutrino Problem and the atmospheric neutrino anomaly, in the form of neutrino oscillations [22]. Pontecorvo predicted that the solar ν_e -flux could be up to a factor of two smaller than the SSM predicted. He discussed the possibility of bi-directional mixing between electron and muon neutrino flavours, and also oscillations to sterile neutrinos, which do not interact via the electromagnetic, weak or strong interactions.

One can express the neutrino flavour states $|\nu_\alpha\rangle$ as a coherent superposition of mass states [17]:

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* |\nu_k\rangle, \quad \alpha = e, \mu, \tau. \quad (1.10)$$

In Equation (1.10), $|\nu_k\rangle$ is a neutrino state with mass m_k , $|\nu_\alpha\rangle$ is the state of the neutrino flavour and U is a unitary mixing matrix. The neutrino mass states are eigenstates of the Hamiltonian:

$$\mathcal{H} |\nu_k\rangle = E_k |\nu_k\rangle, \quad (1.11)$$

with energy eigenvalues given by the dispersion relation:

$$E_k = \sqrt{p^2 + m_k^2}. \quad (1.12)$$

In the ultra-relativistic limit ($p = E$), one can approximate this as:

$$E_k \simeq E + \frac{m_k^2}{2E}. \quad (1.13)$$

The neutrino mass states must also be solutions of the Schrödinger equation, so evolve over time as a plane wave. One can re-write Equation (1.10) to give the time-dependent flavour states:

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* e^{-iE_k t} |\nu_k\rangle. \quad (1.14)$$

Equation (1.14) is valid if one cannot resolve neutrino mass differences in weak decays. From Heisenberg's uncertainty principle, one could resolve a mass difference with a neutrino energy orders of magnitude larger than it. For a neutrino with energy uncertainty ΔE , whose state changes over the time interval Δt , the uncertainty relation is:

$$\Delta E \Delta t \gtrsim 1. \quad (1.15)$$

Substituting for ΔE using Equation (1.13), and assuming $\Delta t \simeq L$, the distance travelled by the neutrino, between two points in space-time, gives:

$$|\Delta m_{kj}^2| \frac{L}{2E} \gtrsim 1. \quad (1.16)$$

For example, for ${}^8\text{B}$ solar neutrinos, with an energy 5 MeV, travelling 149.6×10^6 km from the Sun, one could expect to be sensitive to a squared mass difference of the order $\Delta m^2 \gtrsim 7 \times 10^{-5} \text{ eV}^2$.

From the unitarity of the matrix \mathbf{U} , one can invert Equation (1.10), to express the mass states as a function of flavour states. This allows one to write:

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_k U_{\alpha k}^* e^{-iE_k t} U_{\beta k} \right) |\nu_\beta\rangle, \quad (1.17)$$

which describes a pure flavour state, at $t = 0$, that is a superposition of mass states, becoming a superposition of flavour states, at time $t > 0$, if \mathbf{U} is not diagonal. From this, one must be able to define the probability of a transition from flavour state $|\nu_\alpha\rangle$, at time $t = 0$, to another pure flavour state $|\nu_\beta\rangle$ at time $t > 0$. The derivation of this probability is beyond the scope of this thesis, see e.g. [17, 23], but the crucial point is that a non-zero probability of transition between two pure flavour states, requires a non-zero value for Δm_{kj}^2 . Thus, if oscillations between pure flavour states occur, the flavour states must be a superposition of at least two massive neutrino states (one state could remain massless).

The resolution of the Solar Neutrino Problem, came in the form of the Sudbury Neutrino Observatory (SNO) experiment, a 1 kt, heavy water, Cherenkov detector. SNO was able to detect solar neutrinos via three interactions:

1. The Charged Current (CC) interaction:

$$\nu_e + d \rightarrow p + p + e^-, \quad (1.18)$$

which is sensitive only to electron neutrinos.

2. The Neutral Current (NC) interaction:

$$\nu_\ell + d \rightarrow n + p + e^-, \quad (1.19)$$

which is sensitive to all three neutrino flavours.

3. The Elastic Scattering (ES) interaction:

$$\nu_\ell + e^- \rightarrow \nu_\ell + e^-, \quad (1.20)$$

which is sensitive to all flavours, but six times more sensitive to ν_e than to the other two.

SNO measured the NC flux as $\phi_{NC}^{SNO} = 6.42_{-1.57}^{+1.57}(\text{stat.})_{-0.58}^{+0.55}(\text{syst.})$ [24], which agreed with the predictions from the SSM ($\phi_{SSM} = 5.05_{-0.81}^{+1.01}$). This proved that one could resolve the Solar Neutrino Problem by incorporating neutrino flavour transitions into the model. Figure 1.1, shows a comparison of the measured solar neutrino fluxes, against the SSM prediction, for all solar neutrino experiments, highlighting SNO's solution to the Solar Neutrino Problem.

Transitions between flavour states also proved a viable solution to the atmospheric neutrino anomaly. In 1998, the 50 kt, water Cherenkov Super-Kamiokande detector, found measure-

Total Rates: Standard Model vs. Experiment
Bahcall–Serenelli 2005 [BS05(OP)]




Figure 1.1: A comparison of the fluxes measured by each of the solar neutrino experiments, with the prediction based on the Standard Solar Model and electroweak theory. Taken from [25].

ments of the flux ratio R and the distribution of μ -like events, supported the hypothesis of two-flavour neutrino oscillation [26]. Neutrino oscillation experiments have since made precision measurements of most of the oscillation parameters. Yet, the sign of the large mass difference remains unknown, so one cannot be sure of the mass ordering. The two possibilities are the normal hierarchy (NH), where $m_3 > m_2 > m_1$, or the inverted hierarchy (IH), where $m_2 > m_1 > m_3$. Neutrino oscillations remain the sole experimental result that the SM cannot explain, suggesting that the SM is an effective theory of a more general model that incorporates neutrino masses.

1.2.2 Adding Dirac neutrino masses to the Standard Model

A possible extension to the SM, to incorporate neutrino masses, is to introduce the right-handed component of the neutrino field $\nu_{\alpha R}$, where $\alpha = e, \mu, \tau$. Often referred to as the minimally extended SM, this brings neutrinos in line with the charged leptons and quarks, which already have right-handed singlet components. The additional light right-handed component is almost undetectable. The minimally extended SM introduces three right-handed neutrino fields by adding an extra term to the SM Higgs-lepton Yukawa Lagrangian, to give:

$$\mathcal{L}_{H,L} = - \sum_{\alpha,\beta=e,\mu,\tau} Y_{\alpha,\beta}^{\ell\ell} \overline{L_{\alpha L}} \Phi \ell'_{\beta R} - \sum_{\alpha,\beta=e,\mu,\tau} Y_{\alpha,\beta}^{\nu\nu} \overline{L_{\alpha L}} \tilde{\Phi} \ell'_{\beta R} + \text{H.c.}, \quad (1.21)$$

where the second term has the same form as the corresponding term in the Higgs-quark Yukawa Lagrangian. In Equation (1.21), $Y_{\alpha,\beta}^{\ell\ell}$ and $Y_{\alpha,\beta}^{\nu\nu}$ are the matrices of yukawa couplings for the charged leptons (ℓ) and neutrinos (ν). The Higgs fields are represented by Φ and $\tilde{\Phi}$ and $\overline{L_{\alpha L}}$ and $\ell'_{\beta R}$ are the charged lepton and neutrino fields. One can write the Lagrangian in matrix form by collecting the fields in three-dimensional vectors in flavour space:

$$\boldsymbol{\nu}'_R = \begin{pmatrix} \nu'_{eR} \\ \nu'_{\mu R} \\ \nu'_{\tau R} \end{pmatrix}, \quad \boldsymbol{\nu}'_L = \begin{pmatrix} \nu'_{eL} \\ \nu'_{\mu L} \\ \nu'_{\tau L} \end{pmatrix}, \quad (1.22)$$

$$\boldsymbol{\ell}'_R = \begin{pmatrix} e'_{-R} \\ \mu'_{-R} \\ \tau'_{-R} \end{pmatrix}, \quad \boldsymbol{\ell}'_L = \begin{pmatrix} e'_{-L} \\ \mu'_{-L} \\ \tau'_{-L} \end{pmatrix}. \quad (1.23)$$

$$(1.24)$$

After the Higgs field acquires a vacuum expectation value

$$\langle \Phi \rangle = \left(\frac{v}{\sqrt{2}}, 0 \right)^T, \quad (1.25)$$

this yields the mass terms for the charged leptons and neutrinos:

$$\mathcal{L}_{H,L} = - \left(\frac{v}{\sqrt{2}} \right) \left[\overline{\boldsymbol{\ell}'_L} \cdot \mathbf{Y}^{\ell\ell} \cdot \boldsymbol{\ell}'_R + \overline{\boldsymbol{\nu}'_L} \cdot \mathbf{Y}^{\nu\nu} \cdot \boldsymbol{\nu}'_R \right] + \text{H.c.} \quad (1.26)$$

The Yukawa matrices \mathbf{Y}^{ℓ} and \mathbf{Y}^{ν} can be diagonalised by re-defining the leptonic fields. This leads to the bi-unitary diagonalisation:

$$V_L^{\nu\dagger} \mathbf{Y}^{\nu} V_R^{\nu} = \text{diag}(Y_1^{\nu}, Y_2^{\nu}, Y_3^{\nu}), \quad (1.27)$$

via two unitary mixing matrices Y_L^{ν} and Y_R^{ν} . Here, the eigenvalues y_k^{ν} are real and positive. Substituting the chiral charged lepton arrays and chiral massive neutrino arrays:

$$\boldsymbol{\ell}_L = V_L^{\ell\dagger} \cdot \boldsymbol{\ell}'_L, \quad \boldsymbol{\ell}_R = V_R^{\ell\dagger} \cdot \boldsymbol{\ell}'_R, \quad (1.28)$$

$$\mathbf{n}_L = V_L^{\nu\dagger} \cdot \boldsymbol{\nu}'_L, \quad \mathbf{n}_R = V_R^{\nu\dagger} \cdot \boldsymbol{\nu}'_R, \quad (1.29)$$

the diagonalised Higgs-lepton Yukawa Lagrangian is:

$$\mathcal{L} = \left(\frac{v}{\sqrt{2}} \right) \left[\sum_{\alpha=e,\mu,\tau} y_{\alpha}^{\ell} \overline{\ell}_{\alpha L} \ell_{\alpha R} + \sum_{k=1}^3 y_k^{\nu} \overline{\nu}_{kL} \nu_{kR} \right] + \text{H.c.} \quad (1.30)$$

Given the Dirac fields:

$$\ell_{\alpha} = \ell_{\alpha L} + \ell_{\alpha R}, \quad (\alpha = e, \mu, \tau), \quad (1.31)$$

$$\nu_k = \nu_{kL} + \nu_{kR}, \quad (k = 1, 2, 3), \quad (1.32)$$

one can write the mass terms in the Lagrangian as:

$$\mathcal{L}_{HL} = - \sum_{\alpha=e,\mu,\tau} \frac{y_{\alpha}^{\ell} v}{\sqrt{2}} \overline{\ell}_{\alpha} \ell_{\alpha} - \sum_{k=1}^3 \frac{y_k^{\nu} v}{\sqrt{2}} \overline{\nu}_k \nu_k. \quad (1.33)$$

The second term is a Dirac mass term for the neutrinos giving the neutrino masses:

$$m = \frac{y_k^{\nu} v}{\sqrt{2}} \quad k = 1, 2, 3. \quad (1.34)$$

Note, from Equation (1.34), the mass is proportional to the Higgs vacuum expectation value $v = (246 \text{ GeV})$. One can estimate the mass of the lightest neutrino state using cosmological constraints. Since light, massive neutrinos constitute hot dark matter, experiments such as the Planck experiment can place a constraint on the sum of the three neutrino mass states. In a recent analysis the Planck Collaboration estimate a 95 % Bayesian credible interval on the upper bound, as $\sum_{k=1}^3 m_k < 0.23 \text{ eV}$ [27]. From this, taking the mass of the lightest neutrino state as $\sim 0.1 \text{ eV}$, puts the yukawa couplings (y_k^{ν}) at $\sim 10^{-12}$. The Dirac mechanism for

generating a neutrino mass, provides no explanation for the smallness of the Higgs-neutrino Yukawa couplings.

1.2.3 The Majorana mass term

An alternative possibility for massive neutrinos is to introduce a Majorana mass term. One can express a Majorana field ψ by chiral decomposition to give:

$$\psi = \psi_L + \psi_R, \quad (1.35)$$

where ψ_L and ψ_R are the left-handed and right-handed components. But in the Majorana interpretation the charge conjugate gives the condition that:

$$\psi_R = (\psi_L)^C = (\psi^C)_R, \quad (1.36)$$

which we can substitute back into the chiral decomposition, to give:

$$\psi = \psi_L + \eta(\psi^C)_R = \psi_L + e^{i\phi}(\psi_L)^C. \quad (1.37)$$

The phase factor $\eta = e^{i\phi}$ is arbitrary, meaning that if we take the charge conjugate of the Majorana field:

$$\psi^C = e^{i\phi}\psi, \quad (1.38)$$

we can see that the charge conjugate of the Majorana field coincides with itself, up to an arbitrary phase factor. This means that neutrinos coincide with their own antiparticle and can be thought of as the fermionic equivalent to the photon.

With a Majorana field there are now two possibilities to generate a Majorana mass term for neutrinos. The left-left Majorana mass term takes the form $\nu_L^T C M \nu_L$ and the right-right Majorana mass term: $\nu_R^T C M \nu_R$. The right-right term requires the addition of a right-handed neutrino field ν_R (not found in the SM) that is independent of the left-handed neutrino field and its antiparticle $(\nu_L)^C$. The problem issue with including just the right-right Majorana mass term, is that it is not protected by the electroweak symmetry, and so would yield masses much larger than the expected lightest neutrino mass (~ 0.1 eV). One would expect this term to produce a mass on the scale of new physics, for example m_{GUT} at the scale of Grand Unified Theories (GUT). This seems to suggest the left-left term as our best choice to extend

the standard model. Yet, whilst the right-right term is $SU(2) \times U(1)$ invariant, the left-left term transforms as a triplet under $SU(2)$, so requires us to introduce an isotriplet Higgs field to the SM.

What would happen then if we were able to introduce both terms to the SM. This is what see-saw mechanisms attempt to do. One can add to the SM lagrangian, one left-left mass term that gives a Dirac mass m_D and one right-right Majorana mass-term. After spontaneous symmetry breaking, we can write the neutrino mass term as:

$$\bar{\nu}_L m_D \nu_R + \frac{1}{2} \nu_{Ri}^T C M_{ij} \nu_{Rj}. \quad (1.39)$$

This means the neutrino mass m_ν depends on both a heavy Majorana mass M and a light Dirac mass m_D :

$$m_\nu = \frac{m_D^2}{M}. \quad (1.40)$$

If we set the Dirac contribution at the scale of electroweak symmetry breaking— $m_D \sim m_{EW} \sim 100$ GeV—and set the Majorana contribution at the GUT scale— $M \sim 1 \times 10^{15}$ GeV—we get a neutrino mass of the order $m_\nu \sim 0.01$ eV. This an order of magnitude smaller than current estimates on the mass of the lightest neutrino state.

1.3 Probing Majorana neutrinos

Two-neutrino double-beta decay ($2\nu 2\beta$) is a rare process that has a typical lifetime, depending on the nucleus, of $T_{1/2}^{2\nu 2\beta} \sim 10^{18}$ to 10^{20} y [28]. The process occurs when a nucleus, with atomic number Z , undergoes two simultaneous beta decays:

$$(Z, A) \rightarrow (Z + 2, A) + 2e^- + 2\bar{\nu}_e. \quad (1.41)$$

The SM allows it, because it conforms to all known conservation laws, including Lepton number conservation. First, observed in 1950, in ^{130}Te , using a geochemical technique [29] $2\nu 2\beta$ is now a known decay mode in naturally occurring nuclei. Its lifetime is well measured for most of these nuclei. The reason for the decay lies in the pairing force, which causes nuclei with an even number of both protons (Z) and neutrons (N), to have a tighter binding than a nucleus with the same $A = Z + N$, where both N and Z are odd. Where the initial nucleus is less bound than the final nucleus, and both states are more bound than the intermediate,

odd-odd nucleus, $2\nu 2\beta$ is energetically favourable. Figure 1.2 shows an energy level diagram for $2\nu 2\beta$, in ^{136}Xe .




Figure 1.2: Energy-level diagram showing the binding energy, relative to the most bound nucleus, for isotopes where $A = 136$. Green bands represent even-even nuclei, whilst red bands mark odd-odd nuclei. Taken from [28].

neutrinoless double-beta decay ($0\nu 2\beta$) is a similar process to $2\nu 2\beta$, except it emits no neutrinos:

$$(Z, A) \rightarrow (Z + 2, A) + 2e^-, \quad (1.42)$$

violating Lepton number conservation, hence the SM prohibits it. W. H. Furry proposed it in 1939, as a test of Majorana’s theory [30]. J. Schechter and J. W. F. Valle, expressed the interaction in Equation (1.42), as an effective operator [31]. In the virtual process shown in Figure 1.3, the “black box” may contain any mechanism, but $0\nu 2\beta$ will be the net result. Connecting the vertices together, as in Figure 1.3, to propagate a $\nu_e \rightarrow \bar{\nu}_e$ transition, generating a mass for the process. Thus an observation of $0\nu 2\beta$, regardless of the underlying physics mechanism, would yield a non-zero Majorana mass for the electron neutrino. This makes $0\nu 2\beta$ a useful probe of the Majorana nature of the neutrino.

The neutrinoless mode and $2\nu 2\beta$, share some common features. In both cases, one can assume negligible nuclear recoil—the emitted leptons carry the majority of the energy released.




Figure 1.3: Feynman diagram showing $\nu_e \rightarrow \bar{\nu}_e$ transition, through the effective “black box” operator for neutrinoless double-beta decay.

Most transitions are from ground state to ground state. The four weak vertices mean that the rate is proportional to G_F^4 , where G_F is the Fermi constant, hence both processes have a low rate. A key difference, is that for $2\nu 2\beta$, the two neutrons, in the parent nucleus that end up as protons, are uncorrelated, where as for $0\nu 2\beta$, they are correlated, since they are linked by light Majorana neutrino exchange.

The standard mechanism, encapsulated in the “black box”, is the decay via a light Majorana neutrino exchange—often referred to as the “mass mechanism”. Figure 1.4 shows this in a Feynman diagram. In this mechanism, the decaying nucleus emits two virtual W bosons, which exchange a Majorana neutrino, emitting two electrons. In a decay via this mechanism, the two electrons take all the energy released by the nucleus, ignoring nuclear recoil. If the kinetic energies of the two electrons are T_1 and T_2 , the visible energy will be $T_1 + T_2$. In contrast to Chadwick’s continuous β -spectrum, whose end-point ($Q_{\beta\beta}$) is the Q -value of the decay, the visible energy spectrum for this mechanism, is a delta function at $Q_{\beta\beta}$. The spectrum smearing of the spectrum, is due to the energy resolution of the detector.

The non-zero Majorana mass, one probes with $0\nu 2\beta$, is the effective Majorana neutrino mass ($m_{\beta\beta}$). One defines it as [23]:

$$m_{\beta\beta} = \left| \sum_{k=1}^3 U_{ek}^2 \right| \quad (1.43)$$

where the summation is over all three neutrino mass states, with masses m_k . Here U is the same unitary mixing matrix that appears in Equation (1.10). Expanding the summation, one can show explicitly the components that determine the value of $m_{\beta\beta}$:

$$m_{\beta\beta} = |U_{e1}|^2 m_1 + e^{i\alpha_2} |U_{e2}|^2 m_2 + e^{i\alpha_3} |U_{e3}|^2 m_3, \quad (1.44)$$




Figure 1.4: Feynman diagram showing neutrinoless double-beta decay via the exchange of a light Majorana neutrino.

where:

$$\alpha_2 \equiv 2\lambda_2, \quad \alpha_2 \equiv 2(\lambda_3 - \delta_{13}). \quad (1.45)$$

The parameters λ_2 , λ_3 and δ_{13} are the two Majorana phases and the CP violating phase, in the parametrisation of the unitary mixing matrix. The next generation of long-baseline neutrino oscillation experiments, should determine their values.

Equation (1.44), also highlights the dependence on the mass states, whose values are also unknown, but the mass squared differences can help to constrain them. As a consequence, the value of $m_{\beta\beta}$ is highly sensitive to the mass ordering. Figure 1.5 identifies the allowed regions of the phase space, defined by $m_{\beta\beta}$ and the mass of the lightest neutrino mass state ($m_{\text{lightest}} - m_1$ if the ordering is normal and m_3 for IH). The large bands in the allowed regions account for the uncertainty in the values of α_2 and α_3 . To the top-right of the phase space, is a quasi-degenerate region (QD), where the phase space is similar for both IH and NH. The NH scenario has a region of complete cancellation, where $m_{\beta\beta}$ could have a value of zero. The IH regime sees no complete cancellation. The $0\nu 2\beta$ experiments constrain the phase space, by excluding values of $m_{\beta\beta}$, from the top of Figure 1.5, moving downwards, whilst limits from cosmology constrain m_{lightest} from the right. Figure 1.5 has the potential to offer conclusions, based on different scenarios. One such example, consider the case where oscillation experiments yield a result consistent with IH, but $0\nu 2\beta$ experiments exclude the

entire IH region of the phase space. In this instance, one could rule out the standard Mass Mechanism (MM) as the effective operator for $0\nu 2\beta$.




Figure 1.5: The phase space for $0\nu 2\beta$. Plots the value of $m_{\beta\beta}$ as a function of the lightest neutrino mass. The plot indicates the quasi-degenerate region (QD) and allowed regions for the NH and IH, as well as limits from the $0\nu 2\beta$ experiments and cosmology. Taken from [32].

The expression:

$$\left[T_{1/2}^{0\nu 2\beta}\right]^{-1} = \mathcal{G}_{0\nu}^{\mathcal{N}} |\mathcal{M}_{0\nu}^{\mathcal{N}}|^2 \frac{|m_{\beta\beta}|^2}{m_e^2}, \quad (1.46)$$

relates the decay rate for $0\nu 2\beta$ to $m_{\beta\beta}$, where $T_{1/2}^{0\nu 2\beta}$ is the isotope dependent half life of the decay and m_e is the electron mass. Equation (1.46), also includes two isotope dependent parameters, the nuclear matrix element for $0\nu 2\beta$ ($\mathcal{M}_{0\nu}$) and the phase space factor ($\mathcal{G}_{0\nu}$), which the following sections discuss in more detail.

1.4 Experimental status of neutrino mass searches

The search for neutrino masses, is a vibrant experimental field, with different approaches to probing neutrino masses. This section will review a selection of current $0\nu 2\beta$ searches, and discuss some of the experimental challenges these experiments face.

From the previous section, one of the key factors in designing an experiment to probe $0\nu 2\beta$ is the choice of isotope. As we saw, Equation (1.46) parametrises the nuclear effects in the nuclear matrix element ($\mathcal{M}_{0\nu}$) and phase space factor ($\mathcal{G}_{0\nu}$). The phase space factor is a numerical solution to the Dirac wave function and the Thomas-Fermi equation. Each isotope has different solutions for both single and double-beta decays. The nuclear matrix element incorporates the transitions between energy levels that the decay process involves. There are different methods for modelling these and calculating $\mathcal{M}_{0\nu}$, which use an approximation of some form. Some common models are the Interacting Shell Model (ISM), which considers a subset of orbits around the Fermi level; the Quasi-particle Random Phase Approximation (QRPA) employs a large valence band which limits the configurations it can include; in the Interacting Boson (IBM-2) model restricts the calculation to 0^+ and 2^+ neutron pairs. Table 1.1 we summarise the phase space factors and range in nuclear matrix elements we consider in this thesis.

Table 1.1: Summary of the properties of common $0\nu 2\beta$ isotopes. We calculate the range in $m_{\beta\beta}$, by selecting the minimum and maximum, for each isotope, from the references [33, 34, 35]. We took the phase space factors from [36].

Isotope	$G_{0\nu}(10^{-15}y^{-1})$	$M_{0\nu, min}$	$M_{0\nu, max}$
^{76}Ge	2.363	2.300	6.258
^{82}Se	10.160	2.180	5.571
^{96}Zr	20.580	2.717	3.469
^{100}Mo	15.920	4.220	6.739
^{124}Sn	9.040	2.100	4.504
^{128}Te	0.588	2.340	5.552
^{130}Te	14.220	2.120	4.742
^{136}Xe	14.580	1.760	3.050
^{150}Nd	63.030	2.670	2.670

In addition to nuclear effects and their impact on the value of $m_{\beta\beta}$, when choosing a double-beta isotope one should also consider the experimental design. A crucial element is the relative abundance of the isotope and cost as both of these will affect the amount of isotope in the target mass. Each isotope also has a specific Q-value, the energy at the end-point of the double-beta spectrum. A higher Q-value isotope is often beneficial to move the region of interest for $0\nu 2\beta$ to a higher energy region, where radioactive backgrounds are typically lower.

After considering the double beta isotope, some the next key considerations are in the construction of the detector. Experiments strive for low radioactive backgrounds, good energy resolution at the end point of the $2\nu 2\beta$ spectrum and the ability to load a large mass of the target isotope. Table 1.2 summarises the limits on $m_{\beta\beta}$ from some current $0\nu 2\beta$ experiments. The GERDA experiment, shown in Table 1.2 uses germanium diodes to probe $0\nu 2\beta$ in ^{76}Ge . CUORICINO and its successor CUORE rely on tellurium oxide bolometers to study $0\nu 2\beta$ in ^{130}Te . Both of these approaches offer a high purity detector, and so low radioactive backgrounds and good energy resolution but there are some challenges with scalability, which could cause problems achieving greater target masses in the future. The most stringent limit in Table 1.2 comes from KamLAND-Zen, a liquid scintillator experiment sensitive to $0\nu 2\beta$ in ^{136}Xe . One can dissolve a large amount of material in liquid scintillators enabling a high target isotope mass. Purification also means that such experiments can achieve low levels of radioactive backgrounds but sometimes they suffer from poorer energy resolution. A final category of experiments are those with a separate source and detector, such as NEMO-3 and its successor SuperNEMO. In Both of these experiments a foil source contains the target isotope, which is surrounded by detectors. This approach leads to excellent reconstruction of the event topology but there are significant challenges with scalability.

1.5 Majoron-emitting modes of $0\nu 2\beta$

The introduction of Majoron-emitting modes of $0\nu 2\beta$ came from Gelmini and Roncadelli, in 1981 [45]. They proposed a model for $0\nu 2\beta$ with an extra vertex, when compared with Figure 1.4. As Figure 1.7 shows, the decay is now propagated by light Majorana neutrino and the emission of a scalar—labelled χ^0 —called a Majoron. This is the simplest model of Majoron-emitting $0\nu 2\beta$ and is now disfavoured because it requires excessive fine-tuning to coincide with precise measurements of the Z boson resonance at the Large Electron-Positron collider [46]. The Majoron introduced here is a light, spin-0 particle that can couple to

Table 1.2: Summary of current experimental 90 % confidence upper limits on $0\nu 2\beta$, via the MM. We calculate a range of values for $m_{\beta\beta}$, based on the range of values for the nuclear matrix element (NME), as Table 1.1 outlines.

Expt	Isotope	Status	$T_{1/2}^{0\nu} (10^{25} y)$	$m_{\beta\beta min} (eV)$	$m_{\beta\beta max} (eV)$	Ref.
GERDA	^{76}Ge	current	0.021	2.276	6.193	[37]
NEMO-3	^{82}Se	past	0.010	1.787	4.567	[38]
NEMO-3	^{100}Mo	past	0.110	0.356	0.568	[39]
CUORICINO	^{128}Te	past	0.011	7.108	16.864	[40]
CUORICINO	^{130}Te	past	0.280	0.335	0.750	[41]
EXO-200	^{136}Xe	current	1.100	0.260	0.450	[42]
KamLAND-Zen	^{136}Xe	current	17.000	0.066	0.115	[43]
NEMO-3	^{150}Nd	past	0.002	3.528	3.528	[44]

all fermions. It can be thought of a Goldstone boson, a scalar particle produced during spontaneous symmetry breaking, producing one Goldstone boson per symmetry broken. In the case of Majoron emission, the broken symmetry is the conservation of classical lepton number. Unlike the Goldstone boson though, Majorons do not necessarily have to be massless, though we only consider massless Majorons. The rest of this subsection will explore further Majoron-emitting mechanisms that are not disfavoured.




Figure 1.6: Feynman diagram showing neutrinoless double-beta decay via the exchange of a light Majorana neutrino, with the emission of two Majorons.




Figure 1.7: Feynman diagram showing neutrinoless double-beta decay via the exchange of a light Majorana neutrino, with the emission of a single Majoron.

Another model for $0\nu 2\beta$ by Majoron-emission was proposed in 1994 by C. P. Burgess and J. M. Cline, in which Majorons could carry classically unbroken lepton number [47]. These so called “charged” Majoron modes, suggest the potential for a mechanism for $0\nu 2\beta$ that preserves the conservation of lepton number. Yet the neutrino coupling to the Majoron, still requires a Majorana mass term for the neutrino, and so an observation of $0\nu 2\beta$ via the emission of charged Majorons would imply a nonzero Majorana mass for the neutrino. Note, in such mechanisms, the Majoron strays from the Goldstone boson interpretation due to the preservation of classical lepton number. A year later, P. Bamert, C. P. Burgess and R. N. Mohapatra, proposed further “multi-Majoron” modes, where—as Figure 1.6 portrays— $0\nu 2\beta$ could emit two Majorons [48]. These modes could be both un-charged or charged, again allowing for a mechanism that conserves classical unbroken lepton number. Between the models presented so far, there we can see that there are three properties of model for Majoron emission that one can change: the number of Majorons emitted (one or two), whether they are charged or uncharged and whether one can class them as a Goldstone boson.

The next key question is how can one distinguish between these modes experimentally. Given the relative smallness of the Q -value of $2\nu 2\beta \sim 2$ MeV, compared to the magnitude of the typical three-momentum of a nucleon in the parent nucleus (~ 60 MeV), one can ignore all but the first-order term in the expansion for the decay rate expression [48]:

$$\frac{d\Gamma}{d\varepsilon_1 d\varepsilon_2} = C(Q - \varepsilon_1 - \varepsilon_2)^n [p_1 \varepsilon_1 F(\varepsilon_1)] [p_2 \varepsilon_2 F(\varepsilon_2)]. \quad (1.47)$$

Equation (1.47) shows that one can completely categorise the predicted energy spectra for different Majoron-modes, by a single integer parameter n —the spectral index. In Equation (1.47), ε_1 and ε_2 are the energies of the final-state electrons and p_1 and p_2 are their three-momenta. The constant C is independent of both the electron energies and momenta. The factors $F(\varepsilon_i)$ represent the Fermi function of the final-state electrons and account for spectral distortions due to nuclear effects in the decay.

In Table 1.3, we summarise the Majoron-emitting modes discussed in this section, specifying for each one its decay mode (how many final-state Majorons), whether one can consider the Majorons as Goldstone bosons, whether they are charged with lepton number L and their corresponding spectral index n . Figure 1.8 indicates the spectral shape Equation (1.47) yields for different values of n . The larger the spectral index value, the further the peak in the spectrum is pushed towards lower energies. The spectra in Figure 1.8 have been convolved using a detector response function for the KamLAND-Zen experiment, to demonstrate the particular visible energy spectra in this detector, but the general spectral shapes and trends would be consistent across all $0\nu 2\beta$ detectors. An important point to note is that the spectral shape for $n = 5$ coincides with the $2\nu 2\beta$ spectrum, so if a Majoron-emitting mode of $0\nu 2\beta$ was predicted to have a spectral index $n = 5$, it would be experimentally indistinguishable from the $2\nu 2\beta$ spectrum. As yet, no mode with spectral index $n = 5$ has been predicted.

Table 1.3: Summary of Majoron-emitting modes of $0\nu 2\beta$.

	Decay mode	Goldstone boson	L	n
Model				
IB	$0\nu 2\beta \chi^0$	no	0	1
IC	$0\nu 2\beta \chi^0$	yes	0	1
ID	$0\nu 2\beta \chi^0 \chi^0$	no	0	3
IE	$0\nu 2\beta \chi^0 \chi^0$	yes	0	3
IIB	$0\nu 2\beta \chi^0$	no	-2	1
IIC	$0\nu 2\beta \chi^0$	yes	-2	3
IID	$0\nu 2\beta \chi^0 \chi^0$	no	-1	3
IIE	$0\nu 2\beta \chi^0 \chi^0$	yes	-1	7
IIF	$0\nu 2\beta \chi^0$	gauge boson	-2	3
bulk	$0\nu 2\beta \chi^0$	bulk field	0	2




Figure 1.8: Reproduced from [49]. Expected spectral shapes for spectral index $n = 1, 2, 3, 5, 7$ for the different Majoron-emitting modes of $0\nu 2\beta$ in ^{136}Xe . The spectra have been convolved using KamLAND-Zen’s detector response function, including energy resolution and energy-scale non-linearities. The resolution-limited line at the Q -value represents $0\nu 2\beta$ without Majoron emission.

1.6 Experimental searches for Majoron-emitting modes

In this section we review the current experimental limits on Majoron-emitting modes of $0\nu 2\beta$. Before comparing the limits for different experiments, we need an isotope-independent metric for an experiment’s sensitivity to $0\nu 2\beta$, so we can draw a direct comparison between two experiments using different $0\nu 2\beta$ isotopes, in the same way $m_{\beta\beta}$ lets us directly compare experiments probing $0\nu 2\beta$ via the standard mechanism. A suitable metric is the effective coupling (g_{ee}), which defines the strength of the coupling of a Majorana neutrino to a Majoron. In the same way Equation (1.46) relates $T_{1/2}^{0\nu 2\beta}$ and $m_{\beta\beta}$ in terms of the nuclear matrix element ($\mathcal{M}_{0\nu}^N$) and phase space factor ($\mathcal{G}_{0\nu}^N$), defining g_{ee} in as a function of the rate of Majoron-emitting $0\nu 2\beta$ decays, also includes ($\mathcal{M}_{0\nu}^N$) and ($\mathcal{G}_{0\nu}^N$) as factors. The exact definition depends

on the number of Majorons in the final state. For modes that emit a single Majoron:

$$g_{ee}^2 = \frac{1}{(\mathcal{G}_{0\nu}^{\mathcal{N}}) |\mathcal{M}_{0\nu}^{\mathcal{N}}|^2 T_{1/2}^{0\nu 2\beta \chi^0(\chi^0)}}, \quad (1.48)$$

where $T_{1/2}^{0\nu 2\beta \chi^0(\chi^0)}$ is the half life for a particular Majoron-emitting mode of $0\nu 2\beta$, and for modes that emit two Majorons:

$$g_{ee}^4 = \frac{1}{(\mathcal{G}_{0\nu}^{\mathcal{N}}) |\mathcal{M}_{0\nu}^{\mathcal{N}}|^2 T_{1/2}^{0\nu 2\beta \chi^0(\chi^0)}}. \quad (1.49)$$

As we saw for experiments probing $0\nu 2\beta$ via the standard mechanism, one of the key factors that can affect the sensitivity of an experiment is the choice of isotope. For some Majoron-emitting modes, different methods for calculating the NME numerically, introduce uncertainty in the value of $\mathcal{M}_{0\nu}^{\mathcal{N}}$, whilst for others numerical modelling provides an approximation. Table 1.4 summarises the phase space factors and range in NMEs for all Majoron-emitting modes of $0\nu 2\beta$ in ^{130}Te . One can produce similar tables for the other isotopes (not included here), which we did to calculate the ranges in g_{ee} in Table 1.5.

Table 1.4: Nuclear information for Majoron-emitting modes of $0\nu 2\beta$ in ^{130}Te . We took the range of NMEs from [50, 34, 35, 33], for $n = 1$, and [50] for the remaining modes. We use the phase-space factors outlined in [51], for $n = 1$ and [50], for modes $n = 3, 7$. No NME or phase-space factor, is available for $n = 2$.

	$G^{0\nu \chi^0(\chi^0)}$	$M_{min}^{0\nu \chi^0(\chi^0)}$	$M_{max}^{0\nu \chi^0(\chi^0)}$
Model			
IB	5.940e-16	2.120	4.742
IC	5.940e-16	2.120	4.742
ID	1.060e-17	0.001	0.001
IE	1.060e-17	0.001	0.001
IIB	5.940e-16	2.120	4.742
IIC	4.970e-18	0.120	0.120
IID	1.060e-17	0.001	0.001
IIE	4.830e-17	0.001	0.001
IIF	4.970e-18	0.120	0.120

Table 1.5 summarises the current experimental limits on all the Majoron-emitting modes of $0\nu 2\beta$ from Table 1.3. In Table 1.5, we consider only experiments that have published limits on the full set of Majoron-emitting modes from Table 1.3. Other experiments have published limits on, for example, all the $n = 1$ modes, but not the higher-order modes. As we saw when comparing experiments probing $0\nu 2\beta$ via the standard mechanism, the liquid-scintillator based KamLAND-Zen experiment has the most stringent limits on g_{ee} , using ^{136}Xe . EXO-200, another experiment using ^{136}Xe as its double-beta isotope, also has strong limits on g_{ee} . The germanium diode experiment (GERDA) has limits on all Majoron-emitting modes of $0\nu 2\beta$ in ^{76}Ge . Notable by its absence, is the double-beta isotope ^{130}Te . As of 2017 there are no published limits for Majoron-emitting modes of $0\nu 2\beta$ in ^{130}Te , a popular isotope due to its high natural abundance. An experiment such as SNO+, with a similar design to KamLAND-Zen, could provide competitive limits on all Majoron-emitting modes of $0\nu 2\beta$, in ^{130}Te .

Table 1.5: Summary of current limits on $T_{1/2}^{0\nu 2\beta\chi^0}(x^0)$ for Majoron emitting modes of $0\nu 2\beta$, for all modes listed in Table 1.3.

Model	KamLAND-Zen $T_{1/2}^{0\nu 2\beta\chi^0}(x^0)$ (y)	g_{ee}	[49]	$T_{1/2}^{0\nu 2\beta\chi^0}(x^0)$ (y)	EXO-200 g_{ee}	[52]	$T_{1/2}^{0\nu 2\beta\chi^0}(x^0)$ (y)	GERDA g_{ee}	[42]
IB	$>2.6 \times 10^{24}$	$< (0.8 - 1.6)10^{-5}$	$>1.2 \times 10^{24}$	$< (0.8 - 1.7)10^{-5}$	$>4.2 \times 10^{23}$	$< (3.4 - 8.7)10^{-5}$			
IC	$>2.6 \times 10^{24}$	$< (0.8 - 1.6)10^{-5}$	$>1.2 \times 10^{24}$	$< (0.8 - 1.7)10^{-5}$	$>4.2 \times 10^{23}$	$< (3.4 - 8.7)10^{-5}$			
ID	$>4.5 \times 10^{23}$	< 0.68	$>2.7 \times 10^{22}$	$0.6 - 5.5$	$>0.8 \times 10^{23}$	< 2.1			
IE	$>4.5 \times 10^{23}$	< 0.68	$>2.7 \times 10^{22}$	$0.6 - 5.5$	$>0.8 \times 10^{23}$	< 2.1			
IIB	$>2.6 \times 10^{24}$	$< (0.8 - 1.6)10^{-5}$	$>1.2 \times 10^{24}$	$< (0.8 - 1.7)10^{-5}$	$>4.2 \times 10^{23}$	$< (3.4 - 8.7)10^{-5}$			
IIC	$>4.5 \times 10^{23}$	< 0.013	$>2.7 \times 10^{22}$	< 0.06	$>0.8 \times 10^{23}$	4.7×10^{-5}			
IID	$>4.5 \times 10^{23}$	< 0.68	$>2.7 \times 10^{22}$	$< 0.6 - 5.5$	$>0.8 \times 10^{23}$	< 2.1			
IIE	$>1.1 \times 10^{22}$	< 1.2	$>6.1 \times 10^{21}$	$< 0.5 - 4.7$	$>0.3 \times 10^{23}$	< 2.2			
IIF	$>4.5 \times 10^{23}$	< 0.013	$>2.7 \times 10^{22}$	< 0.06	$>0.8 \times 10^{23}$	4.7×10^{-5}			
bulk	$>1.0 \times 10^{24}$	-	$>2.5 \times 10^{23}$	-	1.8×10^{23}				

Chapter 2

The SNO+ detector

2.1 Overview of the detector

Creighton mine is an active nickel mine, South-West of the city of Sudbury, in Northern Ontario. The year 1999 saw the completion of the construction of Sudbury Neutrino Observatory, more than two kilometres below the surface, on the 6800' level of Creighton mine's shaft number nine [53]. The granite rock overburden is the same as almost 6000 mwe (metre water equivalent) [54] and reduces the muon flux, through the main part of the detector, to a rate of 70 per day. Sudbury Neutrino Observatory (SNO) completed data-taking in 2006. After its completion, an extension program led to the SNOLAB underground facility we use today. The SNO+ collaboration re-purposed and upgraded parts of the SNO detector to build their experiment. SNO+ benefits from the shielding its location provides and the high radio-purity levels afforded by maintaining SNOLAB as a class 2000 clean room.

The design specification for SNO+, as with other similar neutrino detectors (see for example Borexino [55] or KamLAND [56]), was to achieve low background levels through layers of shielding. Figure 2.1 shows a schematic of the detector, highlighting the key components. The rock overburden and rock surrounding the cavity, comprises the first layer. The cavity, as originally excavated for SNO, has a height of 34 m and a diameter of 22 m [53]. The outer layer of water shielding, contained by the cavity walls, has a mass of 5700 t. An 8 mm Urylon layer, applied to the cavity walls, prevents extra radon contamination entering from the rock [53].

Continuing towards the centre of the detector (in Figure 2.1), the next layer does not serve any shielding purpose, but contains an array of photomultiplier tubes (PMTs). The empirical data SNO+ measures, is the number of photons and the charge deposited in each




Figure 2.1: Wire-frame diagram of the SNO+ detector, showing the cavity, PMT support structure (PSUP) (*dark green*), hold-up ropes (*magenta*), hold-down ropes (*red*) and Acrylic Vessel (AV) (*blue*).

PMT, due to light in the detector. The PSUP, a geodesic sphere formed from stainless steel struts, is the foundation of this layer. SNO fixed 9438 inward-facing 8" Hamamatsu R1408 PMTs, to the PSUP [53]. Although, ~ 500 PMTs failed during SNO, SNO+ has an ongoing PMT repair program, so hopes still to have over 9000 inward-facing PMTs. They housed the PMTs, along with a concentrator to maximise detector coverage, in a hexagonal plastic casing. The hexagons tessellate together to form panels of between five and 27 PMTs (in different configurations). Once joined together on the PSUP, the panels form a near-impermeable barrier to contamination in the outer water shielding. Nevertheless, radiation from the inner surface of the PMTs and support structure, forms a significant component of the external background in the inner detector. Some outward-looking PMTs (OWLs) and PMTs on the neck are also available for providing a veto.

Contained within the PMT array and PSUP is a second layer of water shielding (~ 1700 t) [53], in which floats the AV. The SNO collaboration constructed the 6 m radius sphere, by bonding acrylic panels to form a 56 mm thick, hollow sphere. The purpose of the AV is to contain the target material. For SNO, this was 1 kt D_2O , which is heavier than water, so they installed 10 rope loops and corresponding groove plates, at the AV equator, to keep it suspended. The liquid scintillator (LS) target that SNO+ will use, is buoyant in water, so we have added a hold-down rope system, with anchor points on the cavity floor. The hold-down ropes are 38 mm diameter and made from polyethylene (Tensylon) [54]—chosen for its high radiopurity. Whilst installing the new rope system, we took the opportunity to upgrade the existing hold-up system to 19 mm diameter Tensylon ropes. Despite choosing high radiopurity materials, we predict notable external background contributions from both rope systems and the AV itself, notably from radon-rich dust on its outer surface.

For SNO+, the target material that the AV contains, will define three distinct phases of the experiment. During the *water phase*, we will fill the AV with 905 t of ultra-pure water (UPW), for a few months of data-taking [54]. SNOLAB has upgraded the SNO light-water, underground plant, which will supply the UPW to fill the AV, as well as the water shielding—some of which is already in place. The plant will also periodically re-circulate the UPW, to maintain purity levels. Following the water phase, we will gradually displace the AV-UPW with 780 t of *unloaded LS*. We will describe the composition of both the loaded and unloaded scintillator cocktails in Section 2.3.1. SNO+ has a dedicated scintillator purification plant (see [57] for more details), which we will use to fill the AV and then for re-circulation, to keep radiopurity. The third (*tellurium-loaded scintillator*) phase, will see the scintillator loaded with natural tellurium. We will elaborate on the loading technique in Section 2.3.1.

The AV features a 6.8 m tall neck, shown in Figure 2.1, for inserting calibration sources into the detector. We will calibrate the detector using both radioactive sources, deployed through the neck opening, and a set of optical calibration systems, both deployed sources and systems of optical fibres mounted on the PSUP. SNO+ has seven potential radioactive sources [54], whose primary function is to characterise the detector response—mainly the energy resolution, energy scale and linearity of the scale. The all-encompassing name for the optical fibre system is embedded LED/laser light injection entity (ELLIE). The timing module of ELLIE (TELLIE), uses LED pulses to calibrate the timing response of each PMT. TELLIE, along with the deployed laserball source, provide input for the PMT calibrations (PCA) group. The other modules are scattering module of ELLIE (SMELLIE) and attenuation module of ELLIE (AMELLIE).

2.2 Detector upgrades

2.2.1 Electronics upgrades

To cope with the higher data rate expected for SNO+ we have upgraded much of the electronics inherited from SNO that Section 2.5 describes. Testing during two short periods of (air-filled) detector running, once in February 2012 and then in October 2012, revealed over 200 (of about 9600) electronics channels that were not functioning as expected. In the majority of cases the problem was with the board that prepares the raw PMT signal before it enters the trigger electronics—the PMT interface card (PMTIC). We were usually able to identify a blown resistor or capacitor as the problem and replace them. The electronics repairs performed in autumn 2013 saw the inspection and successful repair of around 90% of these bad channels, which the commissioning run in December 2013 confirmed.

2.2.2 PMT repairs

As with much of the electronics, SNO+ has inherited its PMTs from SNO. SNO was designed to have close to 10,000 Hamamatsu R1408 PMT (see for a detailed schematic), but about 800 PMTs failed during its lifetime. We removed some PMTs to accommodate the hold-up rope-net, yet the design for SNO+ still includes nearly 9500 PMTs. Since increasing the PMT coverage of the AV, directly improves our energy resolution, there was a significant effort to repair some of the failed ones. As of early 2014, we had repaired about 300 PMTs at Laurentian

University (and before that at Queens University), with the ability to remove and repair the remaining failed PMTs by boat, during water-fill.

The rest of this section aims to provide a brief overview of the PMT repair process. Work began with some basic tests using the dark-box set-up pictured in Figure 2.2, predominantly to confirm that the PMTs actually required repairs and to begin isolating the problem. We measured the resistance and capacitance and gradually ramped up the potential to the operating voltage of the PMT; to check for breakdowns. In most cases we saw signs of failure by this point, but if we tested the trigger rate at different thresholds.




Figure 2.2: (a) The set-up used for dark-box testing. (b) Cleaning the lower part of PMT being in mineral solution. (c) Potting the hub with the two-part SilGel mix.

In 90% of cases during the PMT repair program replacing the base fixed the PMT. To replace the base, we disassembled and cleaned the PMT. We removed the base by meticulously slicing through the heat shrink wrap, being careful not to damage the hub, which we planned to reuse, or PMT itself. With the heat shrink wrap peeled off, we carefully prised away the hub from the silicone gel—used to create a watertight seal and protect the base and pins. Next, we de-soldered the base and discarded a large part of the gel seal, before soaking in mineral spirits (Figure 2.2) to break down any remaining gel. Once thoroughly cleaned, we soldered on a new base and repeated the dark-box testing to verify that PMT operated as normal.

We then mounted the original hub and secured in place using a new heat shrink wrap and glue tape, which are both heated steadily and evenly for a good seal. To further improve the seal, we applied a conformal coating to the inside of the hub and electronics and potted the PMTs. The potting process saw the hubs refilled with SilGel, a two-part silicon mix, using a nozzle that ensured the two components mix correctly (Figure 2.2). The SilGel shrinks as it

sets so we topped up the hubs before capping them. At this point, the PMT was repaired and ready for installation back in the detector.

2.3 $0\nu 2\beta$ in SNO+

^{130}Te is an isotope of tellurium, for which it is energetically favourable to decay via two simultaneous beta-emissions. The Q-value for this decay is (2527.518 ± 0.013) keV [58]. The collaboration chose ^{130}Te as its double-beta isotope predominantly for its high natural abundance, at 34.08%. Another key advantage for SNO+, is that the measured two-neutrino double-beta decay ($2\nu 2\beta$) lifetime ($T_{1/2}^{2\nu 2\beta}$) is 7.0 ± 0.9 (stat) ± 1.1 (syst) 10^{20} y [59], meaning that the $2\nu 2\beta$ rate is lower compared to other isotopes. This also means the sensitivity to neutrinoless double-beta decay ($0\nu 2\beta$), scales well with increased loading. Studies, from groups within SNO+ [54], have shown that we can load $^{\text{Nat}}\text{Te}$ to 0.3% by weight, and maintain a stable solution. The studies also show that loading to $\sim 5\%$ is possible.

2.3.1 Tellurium loading

When considering how to load tellurium, the main component is the LS itself. For its scintillator, SNO+ is using linear alkylbenzene (LAB), an aromatic hydrocarbon, consisting of a chain attached to a single benzene ring. The scintillator also acts as a solvent. The choice of LAB, as the scintillator for SNO+, is due to its stability and compatibility with acrylic [54]. It also offers long scattering and attenuation lengths. Lastly, it offers good energy resolution, through high light yield and an a linear response in energy. The scintillator component also includes a small amount of the fluor 2,5-diphenyloxazode (PPO).

The original method for loading the tellurium was to use a surfactant, a compound type commonly found in detergent that lowers the surface tension between two liquids or a liquid and a solid. One can dissolve 0.3%-0.5% tellurium in 5% (by mass) linear alkylbenzene sulphonate, which has the brand name PRS. Investigations into commercial PRS found it to be strongly coloured, and have a low radiopurity and high levels of cosmogenic ^{22}Na , so it requires extensive purification. A group in SNO+ discovered a novel loading technique, where one creates an organometallic complex that dissolves directly in LAB, eliminating the need for a surfactant. The group creates the complex by reacting telluric acid with 1,2-butandiol. A down-select review [60], recommended focusing development on the Te-diol method. With

this method, we expect to load 0.5 % tellurium, by mass. This means we will have 1330 kg of ^{130}Te in the AV, or 6.2623×10^{27} nuclei.

The final part of the loading cocktail is a wavelength shifter, which shifts the scintillation light to a wavelength region that better matches the quantum efficiency of the PMTs. Two options are under investigation at the moment: bis-MSB would give a light-yield of 200 NHit/MeV, whilst perylene corresponds to a light-yield of 300 NHit/MeV. We define an NHit, for a given event, as the number of PMTs for which we include a signal in the trigger sum. Section 2.5.2 explains how we define an event and how the trigger system sums PMT signals. The choice depends on the optics created by other components of the scintillator cocktail and will ultimately determine our energy resolution.

2.3.2 Background considerations


Located in Becancour, Québec, less than 900 km from SNOLAB, is the plant that will produce the LAB for SNO+. The short distance the LAB will travel above ground, reduces the risk of cosmogenic activation during transportation. A Chinese company is producing the telluric acid. Enough telluric acid for >0.1 % isotope-loading, is already at SNOLAB. Storing it underground, shielded from further cosmogenic rays, will give it a chance to radiochemically cool down. We hope to store the remaining telluric acid like this as well, for at least six months before deploying the isotope. We base our target levels, for internal radioactive backgrounds, on the levels attained by KamLAND [61] and Borexino [62, 63].

2.3.3 Current sensitivity estimates

The most recent sensitivity estimate assumes the Te-diol approach, for loading, with 0.5 % tellurium. The analysis predicts a Feldman-Cousins, 90 % confidence upper-limit, on the lifetime of $0\nu 2\beta$, via light Majorana-neutrino exchange, as $T_{1/2}^{0\nu 2\beta} \geq 1.96 \times 10^{26}$ y, with five years of data [64]. This corresponds to an effective Majorana-neutrino mass, accounting for uncertainty in the nuclear matrix element, within the range 38 to 92 meV. Figure 2.3, shows the stacked-histogram spectral plot, in the region of interest (ROI), with an example signal at 200 meV.

Returning once more to the subject of loading, in this section I have mentioned three different loading scenarios: 0.3 %, 0.5 % and 5.0 % $^{\text{Nat}}\text{Te}$, by mass. We can compare these three modes and their impact on our sensitivity to $0\nu 2\beta$, by calculating the potential $0\nu 2\beta$

Figure 2.3: Stacked histogram spectral plot, for 0.5% loading using the Te-diol method (390 NHit/MeV). Includes the main background contributions and an example signal scaled to $m_{\beta\beta} = 200$ meV.



activity of each scenario. This calculation depends on factors including the mass of liquid scintillator (M_{LAB}) (about 780 kg) and a factor for the Fiducial Volume (FV) cut that we apply in the analysis. For this we use a ratio of the cubed radii of the FV (r_{FV}) and AV r_{AV} , which, assuming $r_{AV} = 3.5$ m, gives a factor of about 0.2. Along with the loading fraction L these tell us the mass of ^{130}Te . Given the relative abundance of ^{130}Te , we can determine the mass fraction of ^{130}Te in $^{\text{Nat}}\text{Te}$ ($f(^{130}\text{Te})$), which works out at 0.3470 and combines to give us the mass of ^{130}Te in SNO+—slightly over 160 kg. Then, given the molar mass of ^{130}Te (M_r) and Avagadro's number (N_A), we calculate the number of nuclei as 3.47×10^{26} . Finally, we multiply by $\log 2/T_{1/2}^{0\nu 2\beta}$, where $T_{1/2}^{0\nu 2\beta}$ is the expected half life of $0\nu 2\beta$. Equation (2.1) summarises this calculation:

$$A = \frac{\log 2}{T_{1/2}^{0\nu 2\beta}} \frac{N_A}{M_r} f(^{130}\text{Te}) \frac{r_{FV}^3}{r_{AV}^3} L M_{LAB}. \quad (2.1)$$

Table 2.1 compares the calculated activity, for the three loading scenarios discussed above, assuming the leading limit from Table 1.1 $T_{1/2}^{0\nu 2\beta} = 1.7 \times 10^{26}$ y.

Table 2.1: Comparison of the $0\nu 2\beta$ activity for three different $^{\text{Nat}}\text{Te}$ loading scenarios: 0.3 %, 0.5 % and 5.0 % $^{\text{Nat}}\text{Te}$, by mass.

Loading (percentage $^{\text{Nat}}\text{Te}$)	Activity (y^{-1})
0.3	3.05
0.5	5.08
5.0	50.76

From Table 1.1, it should be clear that the larger the loading fraction, the more counts we would observe for potential $0\nu 2\beta$ signal, which leads to increased sensitivity or a more precise measurement. The other option is to run the experiment for longer, increasing our livetime. We refer to the product of the active mass—the mass of ^{130}Te —and livetime as the exposure; the greater the exposure the greater our expected sensitivity. The preferred loading option for the first phase of $0\nu 2\beta$ data-taking with SNO+, the preferred loading option is 0.5 % and this is the value used for our main sensitivity estimates. The $0\nu 2\beta$ Phase 1 sensitivity estimates, presented in this thesis, assume 0.3 % $^{\text{Nat}}\text{Te}$, as they predate the decision for a default loading scenario of 0.5 %. In the second phase of $0\nu 2\beta$ data-taking, the aim is for few-percent loading, so at present we assume 5.0 % when calculating our sensitivity during this phase.

2.4 SNO+ physics program

Whilst the search for $0\nu 2\beta$ remains a priority for SNO+, the experiment has a wide range of physics goals, in all phases of the experiment. The staged approach to loading the tellurium, means we can investigate different physics processes in each phase. In this section, We will summarise the physics we can do in each phase and the key results we expect to see.

2.4.1 Water-phase physics

The primary goal of the *water-phase* physics program is to investigate invisible nucleon decay. This is where either a proton or neutron, in a nucleus, decays to some undetected final state, leaving the nucleus in an excited state. One can infer the decay through photons emitted as the nucleus de-excites. The list of potential mechanisms that have this signature includes almost 80 possible modes [65], an example is $n \rightarrow 3\nu$. SNO+ is sensitive to such processes, via the decay of ^{16}O nuclei (in water molecules). Were a neutron to decay, it would leave an excited ^{15}O nucleus, which de-excites via a 6.18 MeV photon 44 % of the time. For a proton decay, the excited nucleus would be ^{15}N , emitting a 6.32 MeV γ 41 % of the time [66].

SNO+ expects to improve on limits for these processes, published by SNO [67] and KamLAND [68], since using H_2O over D_2O reduces the neutral current backgrounds that SNO had, and the branching ratios for the ^{16}O decays are higher than those in KamLAND's ^{12}C search [69]. SNO+ expects high sensitivities as both of signals fall in an energy region that is low in radioactive backgrounds. One can also constrain the levels of radioactive backgrounds by a spectral fit, below 5 MeV. Using a Poisson counting method, the 90 % confidence upper-limits, on the neutron and proton lifetimes are: $\tau_n = 1.25 \times 10^{30}$ y and $\tau_p = 1.38 \times 10^{30}$ y [69]. These are a factor of two improvement on the KamLAND limits, with one month of data-taking.

2.4.2 Solar neutrinos

The *unloaded scintillator* phase will follow the *water phase*, lasting at least six months, but potentially up to a year. The key physics goal in this phase, is to build on the success of SNO in probing the solar neutrino spectrum. SNO+ hopes to make the first precision measurement of the CNO cycle. This could help solve the disagreement between the heavy ion (metal) levels predicted by 3D modelling and the results of helioseismic measurements—the so called

Table 2.2: Expected precision of flux measurements, for a global fit to the solar neutrino spectrum [70].

ν source	6 months data	12 months data
${}^8\text{B}$	10.0 %	7.1 %
${}^7\text{Be}$	5.1 %	3.3 %
pep	13 %	8.9 %
CNO + ${}^{210}\text{Bi}$	6.5 %	4.4 %

solar metallicity problem. Through a combined measurement of the pep flux and the ${}^8\text{B}$ spectrum from 1 to 5 MeV—below Sudbury Neutrino Observatory’s energy threshold—SNO+ can probe the ν_e survival probability, in the transition region between vacuum oscillations and the presence of matter effects. If levels of ${}^{14}\text{C}$ are low enough, SNO+ can make a measurement of the p p flux, which would test the solar luminosity constraint.

A recent SNO+ analysis estimated the precision to which SNO+ can expect to measure the flux of each of the solar neutrino signals, with one year of data. The study included each of the signals, as well as around 30 radioactive background contributions—scaled to the levels seen in Borexino (the target for SNO+). They used an extended likelihood fit to determine the expected precision and Table 2.2 shows the results. Figure 2.4 shows the visible energy spectra for each of the solar neutrino fluxes, as well as the key backgrounds, scaled to the target levels.

2.4.3 Reactors, supernovae and geoneutrinos

Another key physics area, which SNO+ can study with both *loaded* and *unloaded* scintillator, is reactor neutrinos. Three nuclear reactors surround SNO+: Bruce, 240 km away, and Pickering and Darlington, both at baselines of 350 km. Other reactors at longer baselines, but still within continental North America, also contribute to the reactor neutrino spectrum. By measuring the combined, oscillated spectrum (see Figure 2.5), SNO+ can measure Δm_{12}^2 to a precision of $0.2 \times 10^5 \text{ eV}^2$ [54], with seven years of data. This is similar precision to what KamLAND achieved [71]. Note the current global best fit is $\Delta m_{12}^2 = 7.54_{-0.22}^{+0.26} \times 10^5 \text{ eV}^2$ [65].

Two other areas where SNO+ can make measurements, are geoneutrinos and supernova neutrinos. Geoneutrinos, share the decay signature of reactor neutrinos, which is inverse β -decay. One can tag events by the characteristic delayed 2.22 MeV γ as a hydrogen atom




Figure 2.4: Expected fluxes for the four main solar neutrino signals and background spectra (scaled to target levels), with a 5.5 m fiducial volume [70].




Figure 2.5: Expected visible anti-neutrino energy spectrum, where the Terrestrial Neutrino Unit (TNU) is a measure of one inverse beta decay per 10^{32} protons. Including stacked (Bruce—blue, Pickering and Darlington—red—and others—yellow) oscillated spectrum, non-oscillated (dashed) and geoneutrino spectrum (solid) [54].

captures the neutron. One can measure the anti-neutrino energy, from the energy the positron deposits in the scintillator. SNO+ is on a homogeneous area of continental crust, so is well placed to study geoneutrinos. The detector should have enough sensitivity to separate out the contributions from the uranium and thorium chain decays, in both the crust and mantle. Inverse β -decay is also one of the channels SNO+ can use to probe the supernova neutrino spectrum. For an example supernova at 10 kpc, SNO+ would expect a burst of more than 500 neutrino events in less than 10 s. The detector also forms part of the supernova early warning system (SNEWS).

2.5 Detection and data acquisition

In the previous two sections, we outlined the physics processes that could occur in the SNO+-AV. Ultimately, each of these processes results in photons, the number of photons being proportional to the energy deposited in the scintillator. In this section we explore how we detect these photons and convert the information into a format we can analyse. First, comes detection itself, where we define a PMT hit and the information we know about each hit. Then we present an overview of the trigger system and how we define an event. Last we look at operating the detector, the data acquisition (DAQ) system and how we define a run.

2.5.1 A hit

Consider a physics process that happens in the SNO+-AV (one of the processes we described in the previous two sections). The raw observable we detect, to study this process, is each photon it produces. A photon, incident on any of the PMTs, may cause a cascade of electrons within the dynode. Whether a given γ , incident on the photocathode, results in a photoelectron, depends on the quantum efficiency of the PMT. The SNO+ PMTs typically have a quantum efficiency around 20 %. The cascade becomes a pulse that travels 35 m, from the base of the PMT, along a coaxial cable to enter the electronics, on the deck (see Figure 2.1).

The first part of the electronics the pulse meets, is the PMTIC, which handles signals from 32 PMTs. Each PMTIC feeds into a single front-end card (FEC). In the FEC daughter board (DB), a discriminator checks if the pulse crosses a certain threshold, and thus is not electronic noise. If the pulse is above threshold, it prompts the release of three signals, a 100 ns pulse (NHit100), a 20 ns pulse (NHit20) and a copy of the original pulse shape, now called ESum,

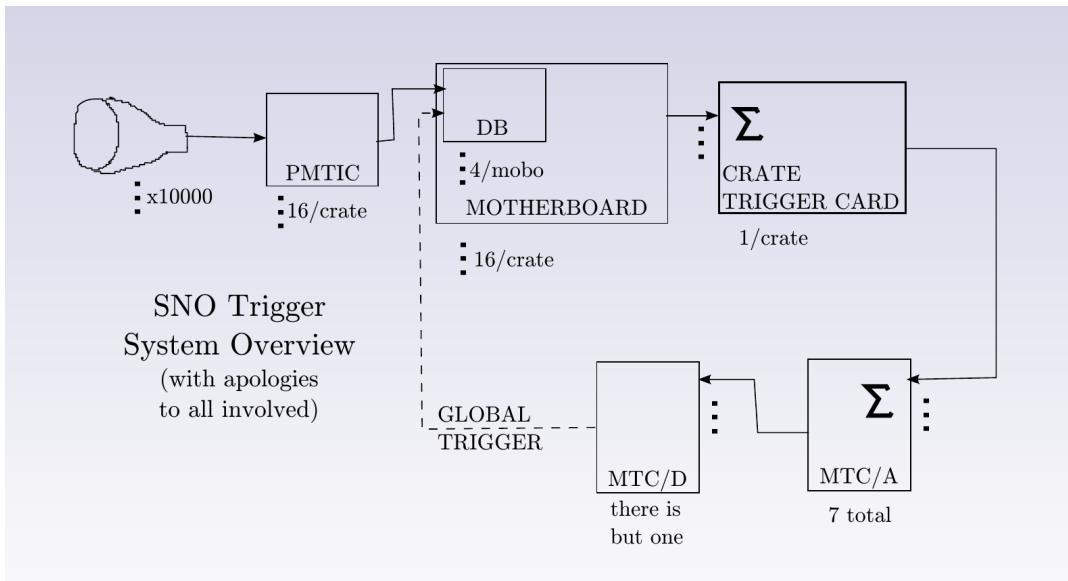


Figure 2.6: Concise overview of the Sudbury Neutrino Observatory trigger system [73].

and also starts a timer. The FEC integrates the pulse, over different time intervals, to produce four values:

- QHS—high gain, short integration charge
- QHL—high gain, long integration charge
- QLL—low gain, long integration charge
- TAC—timer value

These three charge values and the timer count constitute a single PMT hit.

2.5.2 An event

The goal of the trigger system is to decide what information we want to save as an event and what is background or electronic noise. Analogous to the discriminator in the DB, the basic principle of the trigger system is to sum the number of coincident PMT hits and see if this rises above some pre-determined threshold. Here, we have tried to summarise the key steps in the system, but for a comprehensive description, see [72]. Figure 2.6 gives a visual summary of the full system.

Returning to the pulses released in the DB (NHit100, NHit20 and ESum), the Front-end Calibration (FEC) sums the pulses from all 32 channels, and sends the signal onwards to the crate trigger card (CTC). The CTC combines the summation signals from all 16 cards (in the crate) and pushes the sum onto the Analogue Master Trigger Card (MTCA+). This makes one final sum, adding together the signals from all 19 crates. The MTCA+ looks at seven different analogue trigger sums: the three signals we mentioned before—but ESum now has both high gain (ESumHi) and low gain (ESumLo) variants—as well as matching NHit and ESum signals for the OWL PMTs (OWLN, OWLEHi and OWLELo). The MTCA+ passes the complete sums to the Digital Master Trigger Card (MTCD), which looks at the seven analogue channels, along with other digital channels. Examples of the other channels include PulseGT, which samples the electronics at a rate of 5 Hz, or an external trigger from one of the ELLIE systems. There are 26 possible trigger channels, each of which has a threshold, determined by the current run-type. The MTCD constantly monitors all the enabled channels and issues a global trigger (GT) if any of them go above threshold.

The SNO timing system (see Figure 2.7), produces three pieces of timing information. The time card contains a 10 MHz oscillator and a 50 MHz oscillator. Each time the MTCD issues a GT, the timing system records the count values on each oscillator [74]. The system synchronises the 10 MHz clock count to a Universal Coordinated Time (UTC) time-stamp, provided by the GPS receiver. This is essential for recording the local date and time, and thus the position and rotation of the Earth, required by solar and supernova neutrino analyses. The 50 MHz clock count offers a more detailed timescale, which we use to study inter-event spacing.

When the MTCD detects an enabled trigger channel going above threshold, it first creates a raw trigger, which it synchronises against the 50 MHz clock, to create a GT. It then releases a LOCKOUT pulse, which prevents any further GTs [72]. For SNO+, LOCKOUT is 420 ns—meaning events last around 400 ns, with some “dead time” while the trigger cycle resets. At the beginning of the trigger cycle, the system records all the enabled trigger channels, that are above threshold, including the channel that initiated the GT, in the trigger word [72]. Any channels that go above threshold after this cause a missed-trigger flag. At the end of LOCKOUT, if a channel remains above threshold, the MTCD issues a new GT straight away. This is as a retrigger. An automatic retrigger forces an immediate new GT, regardless of whether any enabled channels are above threshold.

The LOCKOUT period defines a SNO+ event. For each event we record all PMT hit information, looking back 220 ns and forward 180 ns, from the start of the trigger cycle. This is often referred to as the trigger window. We also save the values of all three timers—50




Figure 2.7: An overview of the SNO timing system [74].

MHz clock, 10 MHz clock and the Universal Coordinated Time—and the trigger word. There is one more counter in the trigger system, which increases by one at the start of each new trigger cycle, this is the GT counter. We associate the current value of this counter with each event, known as its global trigger ID (GTID).

2.5.3 A run

An operator interacts with the detector through ORCA, a GUI-based DAQ software package, written in objective-C. ORCA allows the operator to set a run-type (automatically enabling the appropriate trigger channels), set the thresholds for each trigger channel and start and stop a run. In SNO+, we define a run as a period of data-taking, where the detector state and conditions remain constant. During stable running, these conditions could remain constant for hours, or even days, so we expect to start a new sub-run every two hours.

The second job of the DAQ system is to manage data readout. When the MTCD issues a GT, the FEC converts the analogue hit information to a digital format, so the XL3 system can readout the data. The XL3 system also combines it with the corresponding digital trigger information and status bits, to form a PMT bundle. The XL3 sends the PMT bundles back to the DAQ machine, which passes them onto the event builder. This combines all PMT bundles, with the same GTID, in a hierarchical data format, to form a stream of events. The builder releases the event data in both a dispatch stream and as complete runs, in one or more ZDAB files.

Chapter 3

High-level checks for SNO+ data quality

Data Quality (DQ) comprises a suite of tools that run checks on experimental data to make a quantitative recommendation on its quality. Its primary goal is to act as a direct input to the run selection process allowing us to select the highest quality data for physics analyses. Subsequent aims include the ability to select this data in a fast and efficient manner and also to select it with high purity—there should be a low probability of accidentally marking poor quality data as good. Another important factor when designing the DQ software was flexibility because different analyses may differ significantly in their definition of what constitutes good data.

DQ comes in three flavours. Low-level DQ (DQLL) deals with information straight from the event builder and the trigger system. One can find a full description of these checks in [75]. Calibrations DQ (CDQ) performs checks relating to specific calibration sources, assessing the quality of data against the source requirements. For details of these checks please see [76]. The focus of this chapter will be high-level DQ (HLDQ).

HLDQ is the last FEC process to run in the main processing chain because it requires calibrated PMT information and assumes the clean data. An important distinction between HLDQ and Data Cleaning (DC) is that DC makes a recommendation at the event level—on which events to analyse—whereas HLDQ advises on a run-by-run basis. The output of HLDQ acts as primary input, along with the other DQ flavours, to Run Selection. Through Run Selection, HLDQ ensures analysers can select data, for calibrations and physics analyses, that meet the quality standards required to achieve our physics goals.

Figure 3.1 summarises the two streams of processing: the main processing chain and nearline processing and where the HLDQ checks fit into each processing stream. Presently, there is no implementation of HLDQ processing within the nearline framework. We are

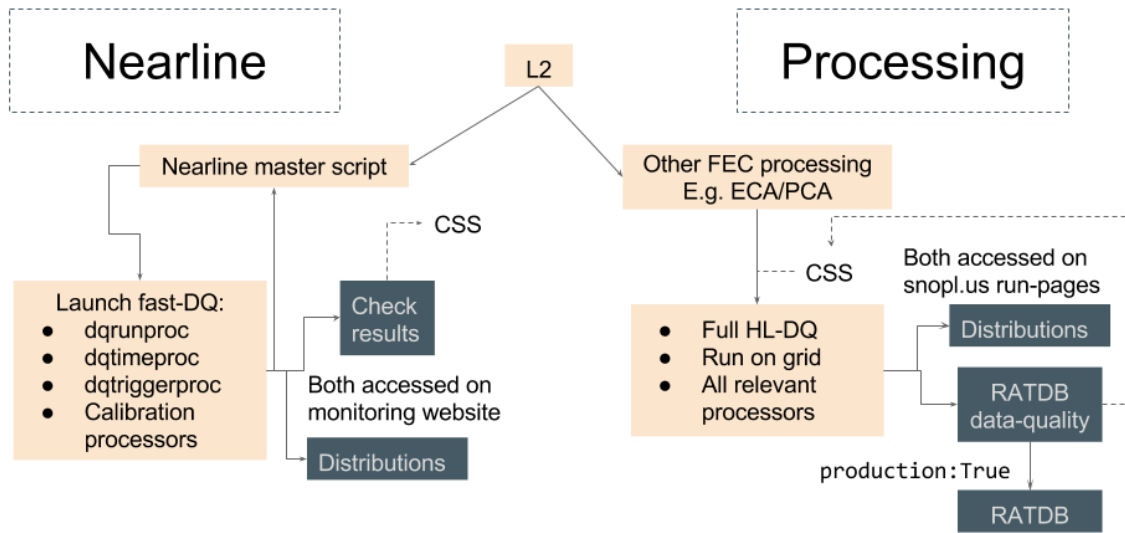


Figure 3.1: Flow diagram showing the HLDQ processors in the two processing and how they fit in relation to the other FEC processes.

working on this at present, so I have indicated where the processors fit in the nearline stream in Figure 3.1. In Sections 3.1 and 3.6, we will describe the HLDQ software from the perspective of the main processing stream. In Section 3.7, we will describe the nearline implementation in further detail.

3.1 Overview of the software

3.1.1 The parent processor

A crucial role of the parent processor is to gather information from each of the daughter processors and save it to RAT database (RATDB). The `DBTable` class-member `fRATDBOutput` takes responsibility for this. We used Reactor Analysis Tool (RAT)’s built-in mechanism for creating RATDB tables because of its uniformity in output and ease of use. Notably, the code yields some more complex data structures that it must save to RATDB, because we group the information produced by HLDQ first by processor and then by check. Working with the `DBTable` and `json::Value` classes facilitates saving these more complex structures to RATDB. We define the `DBTable` instance as a static class member and initialise it outside the parent

class definition, to collate the results from all active processors—for a given run—in a single RATDB table.

Two objects saved in the `DBTable` instance are class members `fAllChecks` and `fCriteria`. Both are JavaScript Object Notation (JSON) arrays, the former storing the result (pass or fail) for each check and the latter the criteria that the checks use. These too are static objects, initialised before the parent class constructor to accumulate information from the other processors. We store the results of each check in a `BitMask` instance to permit queries of the results of a set of checks. This enables analysers to define a series of checks that must have passed to consider a run as meeting the quality standards for a given analysis. Within the `BitMask` instance, we record check results using the convention pass = 1 and failure = 0¹. The related member `fApplied` is another `BitMask` that we use to record all the checks applied to a given run, regardless of the result (applied = 1 and not applied = 0). We also set a Boolean member variable `fCheckResult` to store the result of the current check, employing the convention pass = true and fail = false.

The HLDQ code also produces some standard output distributions, which are useful diagnostics to explain the result of a particular check or investigate a potential problem. We will discuss most of the output plots when describing the check they relate to, but there are two general plots produced by the parent processor that we will describe here. The histogram `fAppliedFlagsHist` records the occupancy for each check—essentially representing the `fApplied` bit-mask as a histogram. Similarly, `fFlagsHist` histograms the checks passed to represent `fFlags` in the same way. The Boolean variable `fOutputImages` controls whether HLDQ plots and saves the output distributions as image files. We save all images in the Portable Network Graphic (PNG) format, with the naming convention `HistogramName_RUN_pPASS.png`, where `RUN` is the run number for the current run and `PASS` is its pass number. Figure 3.2 shows a summary of checks applied and the check results when running the processors on run 8614.

Even if they are not saved as images, HLDQ collects the distributions from each processor and saves them in the `TFile` instance `fRootFile`. Again this must be a static member, initialised outside the class, to ensure we save distributions from all active processors to a single ROOT file. The methods `OpenRootFile` and `WriteToRoot` manage the ROOT file, the first opening the file and the latter writing the distributions—currently in memory—to the file. The naming convention for all output files is `DATAQUALITY_RECORDS_RUN_pPASS`, where, as before, `RUN` represents the current run number and `PASS` the pass number. For ROOT files, HLDQ appends the `.root` file extension.

¹recent discussions to standardise the convention for bit-masks across RAT had yet to reach a concrete decision, at the time of writing.



Figure 3.2: Summary of status of the HLDQ checks. Red shading identifies the applied checks, the blue histogram shows the check result (pass/fail), for run 8614.

3.1.2 General format of the processors

As its name suggests, the `BeginOfRun` method is a form of pre-processor. As such its primary role here is in loading the constants from RATDB that HLDQ requires to perform the checks. The main set of constants needed are the criteria for the checks, which the code must load from the `DQCHECKS` RATDB table. The checks realised by the HLDQ processors often stem from a count or proportion of events meeting a target level. We use the `BeginOfRun` method to initialise any count variables that these checks will use.

The `Event` method in the `Processor` base class accepts the T tree entry for the current event, as input, which I propagate when overriding this method in the HLDQ processors. This makes readily available any event-level information the checks entail. The HLDQ code does not make a direct comparison between this event-level information and the run-level criteria, but the purpose of the `Event` method is to manipulate the relevant information into a form where one can draw run-level comparisons. We achieve this by counting the number of times the event information meets an event-level condition and then setting a criterion on the total (or proportion), for example when we check that ORCA enabled the correct triggers for the run, the event level condition checks the event trigger word and then we check the run by looking at the proportion of events passing the event condition. We will give specific details of each check in Sections 3.2 to 3.5.

In the same way one can think of the `BeginOfRun` method as a pre-processor, the `EndOfRun` method is a post-processor. By this point, the processor should have analysed the event-level information ready for interpretation from the point-of-view of an entire run. The `EndOfRun` method takes as inputs the set of metrics that describe the full run and the set of criteria that define the appropriate quality standard (that the `BeginOfRun` method loaded). The method compares each metric to its corresponding criterion constant to determine a check result. The result is always a Boolean value. The second task of this method is to update the `DBTable` object, in the parent class, with the result of each check. The `Update` method in `DataQualityProc` facilitates this, taking the name of a check and Boolean `result` value as inputs.

3.2 The run processor

The first two checks ensure the run is appropriate for analysis. One way to do this is by querying the run-type bit-mask defined by ORCA. For `rat-5.3.2`, the code implements this

check to pass every time, using a hard-coded run-type value of one and defining a run-type criterion of one in the DQCHECKS RATDBtable.

To properly apply this check, HLDQ should define the criterion as a bit-mask that will query the appropriate bits of the run-type bit-mask. ORCA stores this bit-mask in the RUN table in RATDB. An improved version of this check would get both the run-type and criterion bit-masks from the database and then execute something like:

```
if( ( runType & runTypeCriteria ) == runTypeCriteria )
    // check passes
else
    // check fails.
```

We have yet to realise a check along these lines, due to recent ² instability in the run-type definition. Members of a DAQ workshop fixed the run-type bit-mask ahead of the Second Mock Data Challenge (MDCii) [77]. Using this definition (see Table 3.1), we would define the criterion as 2, so that any run defined as physics would give a pass (true) result. Note other run-types, for example with the external or deployed source bits (3 or 4), might still be suitable for other analyses, such as calibration analyses.

The second of this group of checks is the Monte Carlo (MC) flag check. This looks at the MC flag stored in the runT tree. The role of this check is to ensure that we analyse data, and not Monte Carlo, during the main physics analyses. In the BeginOfRun method, the processor loads the MC flag criterion (= 0 [78]) from DQCHECKS. The check requires no event information. The EndOfRun method obtains the MC flag value from the runT tree and compares it with the criterion. If they are equal—MC flag is 0—the check result is pass (true), but if they differ—MC flag = 1—the check results in failure (false).

The purpose of the trigger check is to check that, given the run-type, ORCA activated the correct set of triggers. The main input is a mask of triggers that should be active for a specific run-type. The current version of the code (rat-5.3.2) defines the set of trigger channels for physics runs in a bit-mask. One could extend the check with minimal development, to include different trigger channel criteria for different run-types, as we have indexed sets of criteria in DQCHECKS by run type (“neutrino” covers all physics runs). Table 3.2 summarises the trigger signals that contribute to the trig word and the bit to which each signal corresponds. We have highlighted in red the triggers that should be active in a physics run [79]. The

²At the time of writing

Table 3.1: Bit-wise description of the run-type bit-mask [77].

Bit number	Run type
0	Maintenance
1	Transition
2	Physics
3	Deployed Source
4	External Source
5	ECA
6	Diagnostic
7	Experimental
8–10	Spare
11	TELLIE
12	SMELLIE
13	AMELLIE
14	PCA
15	ECA_PDST
16	ECA_TSLP
17–20	Spare
21	DCR Activity
22	Compensation Coils OFF
23	PMTs OFF
24	Bubblers ON
25	Re-circulation
26	SLAssay
27	Unusual activity
28–31	Spare

BeginOfRun method retrieves the bit-mask from the DQCHECKS RATDB table and stores it in fTrigCheckCriteria.

Table 3.2: Bit-wise description of the trigType bit-mask.

Bit	Trigger signal
0	NHit 100 Lo
1	NHit 100 Med
2	NHit 100 Hi
3	NHit 20
4	NHit 20 LB
5	ESum Lo
6	ESum Hi
7	OWLN: 100ns NHit trigger on OWL tubes
8	OWLE Lo: ESum trigger on OWL tubes: low gain
9	OWLE Hi: ESum trigger on OWL tubes: high gain
10	PulseGT: pulsed global trigger
11	Prescale: prescale of NHit100 Lo
12	Pedestal
13	Pong
14	Sync
15	EXTASY: external async
16	Ext2
17	Ext3
18	Ext4
19	Ext5
20	Ext6
21	Ext7
22	Ext8 / pulse async
23	Special raw
24	NCD
25	Soft GT
26	Missed trigger: can't mask in/out, but appears in the trig word

A limitation of HLDQ, as of rat-5.3.2, is that it does not use a bit-mask of trigger signals that ORCA activated. I'm unclear how to access this information at present, so the check cannot compare the criteria bit-mask to a bit-mask representing the triggers activated for the run. The code does have access to the event-level trigger word via the `trigType` member of the `ev` branch, which records all the trigger channels that were above threshold when the MTCA+ issues a global trigger. We decided to apply an initial version of the check using the `trigType` bit-mask. An event-level check, in the `Event` method, passes if the trigger word contains all the trigger channels required by `fTrigCheckCriteria`. This means the trigger signals have gone above threshold in all the specified trigger channels, for the given event. We have implemented the event-level check as shown in Algorithm 3.1. This way, each call to the event method counts the number of events failing the check.

Algorithm 3.1 Checking the event-level trigger word

```

Input: ev, fTrigCheckCriteria
  for all events do
    for bit in 0 to 23 do
      queryMask = 2bit
      triggerMask = ev.GetTrigType#()
      if triggerMask & queryMask ≠ criteriaMask & queryMask
        and criteriaMask & queryMask = queryMask then
          Set passedTriggerCheck = false
      if passedTriggerCheck is false then
        Increment failed trigger check count
  
```

To convert to a run-level check, we have defined a second criterion, `fTrigCheckThresh`, that determines the percentage of events that must pass the event level criterion in order for the run to pass as a whole. Again the `BeginOfRun` method fetches this threshold from the `DQCHECKSRATDB` table. We have set it at 90% to allow for some leeway during testing. We think the check is unsuccessful in its current form because it assumes that in most events the trigger signals will go above threshold in all channels highlighted in Table 3.2. In reality, this is rare with the current list of trigger channels. Specifically, the trigger rate for the `PulseGT` trigger is around 5HZ and this will only trigger alongside one of the analogue triggers (`NHit` and `ESum`) if the `PulseGT` trigger fires within 10 ns of the analogue trigger going above threshold [78]. The easiest way to improve this check would be to redefine the trigger criteria so that it only takes into account the analogue signals from Table 3.2. But I believe to fully realise this check, HLDQ will need to be able to access the set of trigger channels that ORCA enabled for the run.

One of the criteria for run selection in SNO checked the length of each run [80]. We have written a similar check for HLDQ. Short runs can be a sign of unstable running or premature termination because of electronic problems, so the run processor looks for a minimum run-length of 1800 s (30 min), the same value used in the SNO First Pass Selection (FPS) checks. We do not constrain the run-length with an upper bound at this stage. Like the other check described so far, the `BeginOfRun` method fetches the minimum run-length value from the `DQCHECKS` and stores it in the variable `fMinRunLength`.

One can get the time-stamp of a SNO+ event from three different variables in the event branch: Universal Time, 10 MHz clock counts and 50 MHz clock counts. Although the 50 MHz clock ticks are generally more appropriate for inter-event timings, the HLDQ code calculates a run-length from all three sources—also following the example of FPS. For each event we define three time-stamp values:

```
eventTime = ev.GetUniversalTime().GetDays() * 86400 +
            ev.GetUniversalTime().GetSeconds()
clockCount10 = ev.GetClockCount10()
clockCount50 = ev.GetClockCount50()
```

We define `eventTime` in this way to calculate the correct run-length for runs spanning day boundaries. Then we use these three values in Algorithm 3.2 to calculate three run-lengths.

Looping through all events, in the `Event` method we set the value of the first event time-stamps—for each timing source—using the first event processed. We also set the each last event time-stamp variable to equal the corresponding first event time-stamp. Then, for each event, We only update the time-stamp if the required information is available. This ensures the code never calculates a negative run-length. One could turn the warning about potential orphan events into an extra check of their occupancy, with minimal development.

Evaluating the run-level result happens in the `EndOfRun` method. This calculates a run-length (in second) by subtracting the first event time-stamp from its corresponding last event time-stamp, converting from counts to seconds as required. Once HLDQ has calculated all three run-lengths, it saves their values to the RATDB output. The processor then uses the first run length value that is greater than 0, in order of preference by source: Universal Time, 10 MHz clock counts and 50 MHz oscillator counts, and posts the value of the chosen run length, along with its source, to RATDB. One could adapt this to select the shortest of the calculated run-lengths to be more conservative. If the run-lengths the processor calculates are ≤ 0 , it

Algorithm 3.2 Calculating run-lengths from three different sources

Input: Universal Time, 10 MHz clock counts and 50 MHz clock counts, from ev branch

Initialise first and last event times, e.g. for UTC

fFirstEventTime = 0

fLastEventTime = 0

for event **do** **if** first event **then**

Set first event time-stamps e.g. for UTC

fFirstEventTime = ev.eventTime

else *All other events* **if** any of UTC, 10 MHz clock counts or 50 MHz clock counts are nonzero **then** **if** Universal Coordinated Time \neq 0 **then**

fLastEventTime = ev.eventTime

if 10 MHz clock ticks \neq 0 **then**

fLastEventCount10 = ev.clockCount10

if 50 MHz clock ticks \neq 0 **then**

fLastEventCount50 = ev.clockCount50

else *No event timing available* **if** GTID is also 0 **then**

Warn of suspected orphan event

Output: Universal Coordinated Time run-length:

fLastEventTime – fFirstEventTime

Output: 10 MHz clock run-length: $1 \times 10^{-6}(\text{fLastEventCount10} - \text{fFirstEventCount10})$ **Output:** 50 MHz clock run-length: $0.2 \times 10^{-6}(\text{fLastEventCount10} - \text{fFirstEventCount10})$

Table 3.3: Definitions of the four statuses we use to identify usable channels

Status	Description
0	PMT has zero occupancy
1	PMT has nonzero occupancy but calibration failed for each hit
2	PMT has nonzero occupancy and calibration failed on at least one occasion
3	PMT has nonzero occupancy and calibration passed for every hit

sets run-length as zero. The processor compares the value of the selected run-length against `fMinRunLength`; a value greater than `fMinRunLength` results in a pass (true) for the check.

3.3 The PMT processor

The purpose of this group of checks is to ensure adequate coverage of the detector by usable PMTs—but how do we define a usable photomultiplier tube? For this processor, in its current form (rat-5.3.2), a usable PMT must meet the following criteria in every run:

- Non-zero occupancy—it must register at least one hit
- Processing must have calibrated each PMT hit, with success

With a concrete definition of a usable PMT, the processor must then analyse every PMT hit in every event to identify the PMTs that meet these criteria.

We start by defining a vector of integers indexed by logical channel number (LCN). Each integer value represents the status of the PMT corresponding to the LCN. The PMT may have one of the four statuses outlined in Table 3.3.

We use Algorithm 3.3 to set the status of each channel. After applying it during each call to the `Event` method, the processor will have a vector of PMT statuses for all channels, over the full run. The HLDQ PMT processor can then use this vector of statuses for the three coverage checks that we describe below.

The purpose of the first coverage check is to ensure the usable PMTs provide adequate general coverage of the detector. As seen in previous checks, we took Sudbury Neutrino Observatory FPS as a starting point, requiring a minimum of 8600 PMTs to be online (around

Algorithm 3.3 Finding usable channels. Parts marked in red are additions we will add to correct errors in the algorithm in rat-5.3.2

Input: ev branch

for PMT in all hits **do**

Try to get PMTUnCal object from ev

Will continue to next hit if this fails

lcn = pmtuncal.GetID()

Move iterator (over channel-calibration-status vector) to LCN

if PMT calibration status is 0 **then**

Set firstHit flag as true

First hit for PMT

Change status: 0 → 1

Set isCalibrated flag to false

Try to get PMTCal object from ev

if it exits **then**

if PMT calibration status is 1 **and** firstHit is true **then**

First time PMT registered a hit calibration succeeded

Change status: 1 → 3

else

PMT registered a hit before and calibration failed

Change status: 1 → 2

else

Catch no PMTCAL object

if PMT calibration status is 3 **then**

PMT registered a hit before and calibration passed before, but here calibration failed

Change status: 3 → 2

90% of all PMT). While SNO+ is still in its commissioning phase, we have opted for a more lenient coverage criterion. We have set the general detector coverage threshold (by usable PMTs) at 70% in the DQCHECKS RATDB table, which the `BeginOfRun` method stores in `fGeneralCoverageThresh`. We also store the total number of channels `fChannelCount` in the `BeginOfRun` method.

We calculate the detector coverage by iterating over the vector of channel statuses. We use the index (LCN) to query the PMTINFO database table and select normal PMTs. For each PMT that is both normal and has a calibration status value of 3, we call the `UpdateCoverage` method, that increments the counter `fGeneralCoverage` by one. After iterating over the vector of calibration statuses, the processor calls `RunCoverageChecks`—a method that defines the actual checks. For the general coverage check, this takes the ratio of the counter `fGeneralCoverage` to the total number of normal PMTs and converts to a percentage. If the detector coverage exceeds `fGeneralCoverageThresh` the check result is a pass (true).

The second coverage check approaches detector coverage from the point of view of the electronics. Its purpose is to ensure there are no large gaps in the online electronics, for example an entire crate offline or a large section of a crate offline. Detector operators should have noted electronics problems such as these at runtime but this check aims to highlight them within the data. One should use the result of the check in conjunction with the shift report notes to ultimately determine if a specific lack of coverage by the electronics will affect the quality of the data produced.

In the `BeginOfRun` method we define a vector of integers (`fCrateCoverage`) with one entry per crate (there are 19 crates in total). A crate's entry in this vector will act as a counter for the number of usable channels in the crate. The `BeginOfRun` method also defines two criteria—taken from DQCHECKS—the in-crate coverage threshold and the crate coverage threshold. The first sets a lower bound on the percentage of channels in a crate that must we must consider usable, at present 50% (or more), for the sake of leniency during commissioning. The crate coverage threshold specifies the percentage of crates that must meet the in-crate threshold, which we have set as all crates (100%) in `rat-5.3.2`.

To perform the checks, we use the same method as the general coverage check. First the `EndOfRun` method iterates over the vector of PMT statuses and selects all instances where a normal PMT has a status of 3. For these PMTs, the processor calls `UpdateCoverage`, which determines the crate responding to the given LCN and increments the counter for that crate—in `fCrateCoverage`—by one. Back in the `EndOfRun` method, we call `RunCoverageChecks`. For the crate coverage check, this method initialises the counter `failedCrates` and then




Figure 3.3: Crate coverage maps showing the occupancy of (*left*) all hit PMTs and (*right*) the occupancy of channels where calibration always succeeded (status 3 in Table 3.3)—for run 8614.

iterates over `fCrateCoverage`. For each crate occupancy value, we normalise by the number of channels in a crate (512) and convert to a percentage. If this percentage is less than `fInCrateCoverageThresh` we increase the counter `failedCrates` by 1. After iterating over all crates, we convert the failed crates count to a percentage of crates passing (`crateCoverage`). If `crateCoverage` \geq `fCrateCoverageThresh` the check returns a pass (true) result.

Associated with this check, the processor produces two coverage maps to help interpret the results of the checks. Both coverage maps take the form of a two-dimensional histogram binned by crate and card number on the x-axis and channel number on the y-axis. This makes it easy to see if a whole crate or card has zero occupancy. The `BeginOfRun` method books the histograms as `fCrateOccupancyMap` and `fUnCalCrateOccupancyMap`. The `Event` method fills them via the `FillCrateCoverageMap` method, to ensure filling of the correct cell based on the crate, card and channel number. As its name suggests, the map `fUnCalCrateOccupancyMap`, records the occupancy of all PMT hits, regardless of whether calibration succeeded or not. we fill this histogram each time Algorithm 3.3 finds a `PMTUnCal` object. The other coverage map `fCrateOccupancyMap` complements this by displaying the occupancy of calibrated (status 3) photomultiplier tubes. We fill it each time the `Event` method detects a `PMTCal` object. Figure 3.3 shows an example of the coverage maps produced by running `DQPMTProc` on run 8614.

SNO FPS also had a criterion that there should be no gaps in coverage, by adjacent PMTs, of comparable size to the detector neck opening [80]. But this check was not automated so they had to rely on information from shift reports. In rat-5.3.2, HLDQ also lacks an automated check of this nature, but we have made provisions to accommodate this in the future. A coverage check from the point of view of the PMT geometry, the third coverage check aims to identify panels that have low occupancy. In the `BeginOfRun` method we set `fPanelCovThresh`—the minimum percentage of panels that must have passed the in-panel occupancy check. The criteria for panel coverage has a value of 80% in `DQCHECKS`.

We apply a similar method as the crate-coverage analysis, by creating a container of panel occupancy counters. Whilst crates have a fixed number of total channels, panels come in five different types, each having a different total number of PMTs. To solve this problem, we created a simple data structure (`panelProp`) to record three pieces of information about a panel: its type, coverage threshold and a counter for its actual occupancy. In rat-5.3.2, the coverage threshold is hard-coded for each panel type as half the maximum PMT occupancy of the panel, rounding down if this is odd. Then, in the `BeginOfRun` method, we form a Standard Template Library (STL) map, indexed by panel number, with a `panelProp` structure for each panel.

Continuing to follow the format of the previous coverage checks, the panel coverage check iterates over the vector of PMTs statuses in the `EndOfRun` method. Each time a photomultiplier tube has a status of three and is also a normal PMT, we call `UpdateCoverage`. For this check, the `UpdateCoverage` method converts from LCN to panel number and then selects the panel properties data structure at the index equal to the panel number. We increment the `coverageCount` member of the data structure by one. Outside the loop over PMT statuses, the code flows to the `RunCoverageChecks` method. Here we initialise count variable `failedPanels` and iterate over the panel map. Wherever the `coverageCount` member of the `panelProp` is less than the `coverageThresh` value, the processor increments the failed panels count. On completing iteration, we define `panelCoverage` as $1 - \text{failedPanels}/\text{fPanelCount}$, where `fPanelCount` is the total panel count, and convert to a percentage. The check results in `pass` (`true`) if this is greater than or equal to `fPanelCovThresh`.

A third output distribution accompanies the panel coverage check to help interpret its results and manually look for large gaps in coverage by adjacent PMTs. The coverage map `fGeoCoverageMap`, that we define in the `BeginOfRun` method, is a scatter plot of θ vs ϕ coordinates, for usable PMTs. We store the θ and ϕ values, for every PMT for which we have called `UpdateCoverage`, in two separate vectors. Then we fill a `TGraph` using the entries from




Figure 3.4: Coverage map plotting the θ and ϕ coordinates of usable PMTs—for run 8614.

these two vectors as co-ordinate pairs. Figure 3.4 illustrates the coverage map this produces by running DQPMTProc on run 8614.

3.4 The timing processor

The checks contained in this processor look at event timings, notably the time between two consecutive events, which we refer to as Δt . As such, the central purpose of the Event method in this processor, is to calculate all the inter-event Δt s. We store the Δt information in a vector containing a simple data structure for each entry. The TimingData data structure contains six data members. The first three are: fGTID, which we use to store the event's GTID; eventTime—containing the event time-stamp (ns) calculated from the 50 MHz clock count; and we store the primary inter-event Δt in fCount50DeltaT. The remaining three members act as placeholders for information needed for some proposed checks that we are yet to write.

The first check that uses the Δt information calculated is the run header check. The aim here is to make sure the gap between starting a run and detecting the first event is not too large. In DQCHECKS I have set the maximum length of the run header as 1.0×10^9 ns (1 s). The BeginOfRun method acquires this value from RATDB and then stores it as fRunHeaderThresh. By saving the first event time separately, in Algorithm 3.4, the processor already has the required information for this check. In the EndOfRun method, we compare the first event time to

Algorithm 3.4 Calculation and analysis of Δt s

Input: ev branch**for all events do**

Increment event count

Create timingData data structure

Save time-stamp and ev.GTID to data structure

if first event then

Save time-stamp of first event

else**if previous event time $\neq 0$ then**Calculate Δt s from: 50 MHz clock, 10 MHz clock and Universal TimeSave 50 MHz clock Δt to data structure**if negative 50 MHz clock Δt then**Fill negative Δt distributionCalculate: delta-GTID and Δ 10 MHz clock count**if delta-GTID is negative then***Events were probably built out of sync*

Fill negative delta-GTID distribution

Increment negative delta-GTID count

if Δ 10 MHz clock count is not negative then*50 MHz clock may have rolled over*

Increment 50 MHz clock resets count

Fill 50 MHz clock resets distribution

else if $0 \geq \Delta t \geq 420$ ns for 50 MHz clock then *Δt is less than trigger window plus lockout*Increment trigger window Δt countFill trigger window Δt distributionFill positive Δt distribution**else if $420 \geq \Delta t \geq 600$ ns for 50 MHz clock then***Possible retrigger*Increment retrigger Δt countFill retrigger Δt distributionFill positive Δt distribution**else***Normal Δt* Fill positive Δt distribution**else***No previous event data*Unable to calculate $\Delta t \rightarrow$ set 50 MHz clock Δt as 0Save current event parameters as previous event values, for use in next call to Event method

`fRunHeaderThresh` and if the time to the first event is less than or equal to `fRunHeaderThresh` the result of the check is a pass (true).

For regular analysis, we generally expect the inter-event Δt s to be positive. In this check, we consider two instances where a negative Δt is acceptable. The first is where the 50 MHz clock counts give a negative Δt , but subtracting the previous GTID from the current GTID also gives a negative value. This is indicative of the event builder not building these events in chronological order. The checks allow negative Δt s that coincide with a negative ΔGTID —`fNegativeDeltaGTIDCount` counts the number of instances where this occurs (see Algorithm 3.4).

The second permitted case for a negative Δt is where the 50 MHz clock counter rolls over. The counter is 43 bits in length, so it should roll over every couple of days [72]. The 10 MHz clock, on the other hand, has a lower frequency and is 53 bits long, so should roll over once every 28 y. We do not expect both counters to roll over simultaneously. A negative Δt from 50 MHz clock counts, coupled to a positive Δt from the 10 MHz clock, suggests a 50 MHz clock roll-over, so this is also allowed. We tally these instances in the count variable `fCount50Resets`. We expect at most 1 of these per run, as maximum run-lengths have been about two hours during recent commissioning data-taking. The count variable `fNegativeDeltaTCount` counts the total number of negative Δt s.

In the `EndOfRun` method, we subtract `fNegativeDeltaGTIDCount` and `fCount50Resets`³ from `fNegativeDeltaTCount` to give the number of disallowed negative Δt s. We define the criterion that at least 99% of all events must have a regular positive Δt (maximum 1% negative Δt s), saved as `fClockForwardThresh`. This is an extra-lenient criterion, while we are still commissioning, that we expect to increase closer to 100% in due course. We normalise `fNegativeDeltaTCount` by the total event count, convert to a percentage and subtract from 100 to get the metric `clockForwardValue`. The result of the check is a pass (true), if `clockForwardValue` \geq `fClockForwardThresh`.

The event separation check also looks at the calculated Δt s, but this check is more concerned with the inter-event spacing. A 400 ns trigger window sets the duration of events in SNO+. A short lockout period follows before the MTCD can issue the next GT. The exact duration of the lockout period can vary, but the shortest duration anticipated is 20 ns [72], giving a minimum duration between two consecutive global triggers of 420 ns. Algorithm 3.4 counts the number of events where Δt lies in the range 0 to 420 ns, in the count variable

³In rat-5.3.2 we also subtract the number of 10 MHz clock roll-overs but we do not expect a 10 MHz clock roll-over to happen within the lifetime of SNO+. We will remove this.

`fTriggerWindowDeltaTCount`. In the `EndOfRun` method we normalise this by `fEventCount` and convert to a percentage to give `eventSeparationValue`. We expect to have no events where the event separation is less than 420 ns but again, whilst we are still commissioning the detector, we have set the criterion (`fEventSeparationThresh`) at a forgiving 1%. If `eventSeparationValue` \leq `fEventSeparationThresh`, the check result will be a pass (true).

Moving to the next event separation bracket, we refer to events that trigger straight after lockout for the preceding event as retrigger events. The detector can retrigger automatically—regardless of whether any trigger signals are above threshold—or it can issue a retrigger only if any of the trigger signals remain above threshold at the end of the previous trigger window. Whatever the mechanism, retriggered events will have an event separation slightly greater than 420 ns. We will adjust the exact timings for retriggers during commissioning, so for now we count any events that fall within the $420 \geq \Delta t \geq 600$ ns range and also fill a histogram with these Δt s so we can study the event separations for retriggers. As with previous checks, we convert the count to a percentage (`retriggersValue`). We have nominally set the threshold for retriggered events at 10% of the total number of events, but we should tweak this during commissioning. We save the value as `fRetriggersThresh` in the `BeginOfRun` method. The `EndOfRun` method will yield a pass (true) result if `retriggersValue` is less than or equal to `fRetriggersThresh`.

We consider any other event separations (≥ 600 ns) as normal. The processor histograms all positive events separations in `fPositiveDeltaTHist` but we also fill distributions with the retriggers and trigger window Δt values, which act as enlarged areas of the full distribution. In total the processor produces six Δt distributions :

- All negative Δt s⁴
- Negative Δt s coupled to a negative delta-GTID
- Negative Δt s that are potential 50 MHz clock resets
- All positive Δt s (Figure 3.5)
- Positive Δt s that fall inside the trigger window (Figure 3.5)
- Positive Δt s that are potential retriggers (Figure 3.5)

We return again to SNO FPS [80] as a basis for the event rate check. In SNO, they applied a maximum event rate of 60 Hz to all events—regardless of trigger channel. For PulseGT events the nominal rate was 5 Hz and they rejected runs that deviated from the nominal rate

⁴We have not shown example plots for the negative Δt s as they were empty for run 8614




Figure 3.5: Distributions showing (*left*) all positive Δt and (*right*) positive Δt s that are within or slightly outside of the trigger window—for run 8614.

by more than 2%. We have used the 5 Hz rate as a lower bound on the total event rate. The upgraded electronics for SNO+ result in the capability to withstand much higher event rates. During commissioning so far, we have tested the detector with event rates up to four kHz, so we have extended the upper bound on event rate to 1.0 kHz during stable physics running.

The processor retrieves these criteria from DQCHECKS in the `BeginOfRun` method and saves them as `fMinEventRate` and `fMaxEventRate`. We calculate the event rate by dividing the total event count by the run-length calculated in `DQRunProc` (saved as the transient variable `run.runLength` in the `runT` tree). Note that to use this requires running both processors in the same macro. The result of the `EndOfRun` method will be a pass (`true`) if the calculated event rate lies in the range specified by `fMinEventRate` and `fMaxEventRate`. We also save the calculated event rate to the RATDB output, for reference.

Observant readers may have noticed that we missed out the second check implemented in this processor. The aim of the Δt comparison check (in `rat-5.3.2`) was an attempt to explain some problems noticed in the 10 MHz clock, by comparing the Δt s calculated from both this oscillator and UTC with the 50 MHz clock derived Δt . A recent check implemented by Mark Stringer, provides a more robust method of comparing 10 MHz clock counts with UTC.

In the new check, Mark converts Universal Coordinated Time to 10 MHz clock counts, using the method `DateTo10MHzClockTicks`. This method rounds to the nearest clock tick and returns an integer. In the `Event` method, it subtracts the result from the 10 MHz clock count

for the current event, which gives the difference (`clockCountDifference`) between the two. If this difference is 0—or 1 in the case where converting it to clock ticks has rounded up—this is a positive result at event level. Otherwise, the process sets the `f10MHzTimeCheck` flag to false. The run-level check, in the `EndOfRun` method, checks the value of this flag. If it remains true for all events, the run-level result is a pass, whilst a single event-level failure yields a fail (false) result for the run. The processor also produces a distribution of the calculated differences.

3.5 The trigger processor

The final processor covered in this document is the trigger processor. The goal of this processor is to analyse the trigger word in greater detail than the current run processor and calculate rates for each of the 27 trigger signals that we saw in Table 3.2. To calculate these rates, we first define a vector of integers (`fTrigTypeCounts`) that has a length of 27—one entry per bit in the trigger word. In the `Event` method, we iterate over this vector and query the corresponding bit of the trigger word. If a trigger channel fired in the trigger word, we increment its counter in the vector by one. In the `EndOfRun` method, we create a second vector (of floats this time), also with 27 entries. We iterate over both vectors simultaneously, calculating the event rate by dividing the trigger occupancy by the transient run-length (`run.runLength`; again one should run this processor in conjunction with `DQRunProc`), and setting the calculated rate in the second vector. We also set the bin contents of a histogram with each calculated rate to produce a visual summary of the rates calculated for each channel in the trigger word; Figure 3.6 illustrates this visual summary for run 8614. Once we have calculated the vector of rates, checking the rate of any trigger signal against the criteria is trivial.

For `rat-5.3.2`, the processor checks the rates, in two trigger channels, against criteria from the `DQCHECKS RATDB` table. We specify the minimum rate for `NHit100Lo` triggers as 5 Hz and save it to `fMinN100LRate` in the `BeginOfRun` method. The check results in a pass (true) if `fTrigTypeRates[0]` (the entry corresponding to `NHit100Lo` in the vector of rates) is greater than, or equal to, `fMinN100LRate`. Similarly, if `fTrigTypeRate[6] ≥ fMinESumHRate`, the `ESumHi` rate check results in a pass (true). The minimum `ESumHi` rate is also set at 5 Hz.




Figure 3.6: Histogram showing the calculated rate for each trigger word in the trigger word—for run 8614.

3.6 Usage of the processors

To use the HLDQ analysers should launch the required processor using standard GEANT4 commands `/rat/proc` followed by the name of the processor. Throughout RAT we use the convention of all lower case, so to run all four HLDQ processors, one would use:

```
/rat/proc dqruncproc  
/rat/proc dqpmtproc  
/rat/proc dqtriggerproc  
/rat/proc dqtimeproc
```

In rat-5.3.2 there is no command to run all four processors but we hope to add this feature in a future release, for ease of use. As with any RAT processor, analysers can also run the HLDQ processors in the event loop of a RAT macro. We have written a run-level processing macro, for use in the main processing chain, which acts as an example of how to run all four HLDQ processors in a macro (see: `rat/mac/processing/partial_water/runLevelProcessing.mac`). By default, HLDQ will save all output distributions to a single ROOT file, but an analyser may prefer to output PNG images directly. One can activate this behaviour by setting the value of the `images` parameter as 1 (true), for the appropriate processor. For example, the following snippet outputs the three coverage distributions as PNGs:

```
/rat/proc dqpmtproc  
/rat/procset images 1
```

3.7 HLDQ nearline

My description of HLDQ so far has centred on running the checks as part of the main processing chain, which relies on Worldwide LHC Computing Grid (WLCG) resources to process the data. A separate processing stream is available in the form of the nearline framework. Written by Kevin Labe, the framework runs lightweight processing on data to provide close to realtime feedback. It may take a couple of days before processing on the WLCG but the nearline framework aims to process most runs before level II trigger (L2) releases the next run. It consists of a master script that acts primarily as a job scheduler, launching registered jobs if they meet their specified job conditions. For example, jobs—including HLDQ—require data for the full run before they can start. Users can register jobs in text files and jobs may launch other processes—rat macros are common.

MDCii saw tests of the first version of the nearline framework but we have yet to incorporate HLDQ. Mark has tested running parts of CDQ nearline. It should be straightforward to expand to the rest of the processors. For HLDQ, we plan to run: the run processor (Section 3.2), the timing processor (Section 3.4) and the trigger processor (Section 3.5), but not the PMT processor, given its weight and dependence on calibrated PMT information, which may not be available.

3.8 Conclusion

In this chapter we have presented the four processors that perform HLDQ checks and the framework we have created for the HLDQ software within RAT. These checks work in conjunction with the other FEC processes, in both nearline and the main process streams, and the results act as a key input to run selection. In Sections 3.2 to 3.5, we have explained key algorithms such as how we calculate the run-length, identify usable PMTs and categorise event separations. We have also provided examples of the distributions we produce to help interpret the results of the checks.

In Table 3.4, we have summarised all the criteria I use during the checks, along with short comments or references. Whilst there are some limitations and caveats with the HLDQ code as of rat-5.3.2, HLDQ accomplishes its core goals. We will address these limitations during the rest of the detector commissioning phases.

Table 3.4: Summary of the criteria HLDQ applies for a standard physics run

Name	Value	Notes/reference
MC flag	0	0 = Data, 1 = MC
Trigger check	$0 \times 66d6$	Recommended active triggers [79]
Trigger check threshold	90.0%	Min % events with correct active triggers
Min run-length	1800 s	[80]
General coverage threshold	70.0 %	Min general % coverage detector-wide
In-crate coverage threshold	50.0%	Min % coverage within each crate
Crate coverage threshold	100.0%	Min % coverage by crates that passed in-crate check
Panel coverage threshold	80.00%	Min % coverage by panels that passed in-panel check
Run-header threshold	1.0×10^9 ns	Time to first event
Clock forward threshold	99.00%	Min % event separation (Δt) where clock count decreases i.e. neagative Δt s
Event separation threshold	1.00%	Max % events with separation (Δt) < 420 ns
Retriggers threshold	10.00	Max % event separation (Δt) that suggest potential retrigger
Min event rate	5.00 Hz	[80]
Max event rate	1.00 kHz	From commissioning
Min NHit100Lo rate	5.00 Hz	[80]
Min ESumHi rate	5.00 Hz	initially same as total event rate

Chapter 4

Setting limits and the echidna software package

Imagine a problem where we have some data points and we wish to fit a model to the data points. The model consists of a set of background contributions, from Poisson processes, plus a theoretical signal contribution. In this problem, one wishes to determine, does the data show a significant contribution from the proposed signal, enough to claim a discovery of the signal? If not, what upper-limit can one place on the signal rate, given the data? This chapter aims to present a method for answering these questions.

To help illustrate the methods the chapter describes, we have constructed a toy model scenario, that is a simple model, bearing some resemblance to the types of visible energy spectra one would see in the main SNO+ neutrinoless double-beta decay ($0\nu 2\beta$) analysis. We have created a Poisson background, like the two-neutrino double-beta decay ($2\nu 2\beta$) contribution in SNO+, and a flat background, resembling the ^8B solar neutrino background in SNO+. Figure 4.1 shows the spectral shapes of these backgrounds, along with an example signal (mimicking the expected $0\nu 2\beta$ visible energy spectrum, from either SNO+ or KamLAND-Zen).

4.1 Formalism

Consider a dataset which comprises an array of measurements in one or more of the observables x , y and z . The dataset contains N events in total, so $\mathbf{x} = (x_i, y_i, z_i)$ denotes the full dataset, where $0 < i \leq N$. Using this convention, the subscripts 0 and u denote the minima and maxima in each dimension. The data \mathbf{x} is now binned in an array, whose dimensions correspond to each of the observables x , y and z . The array has the following structure:

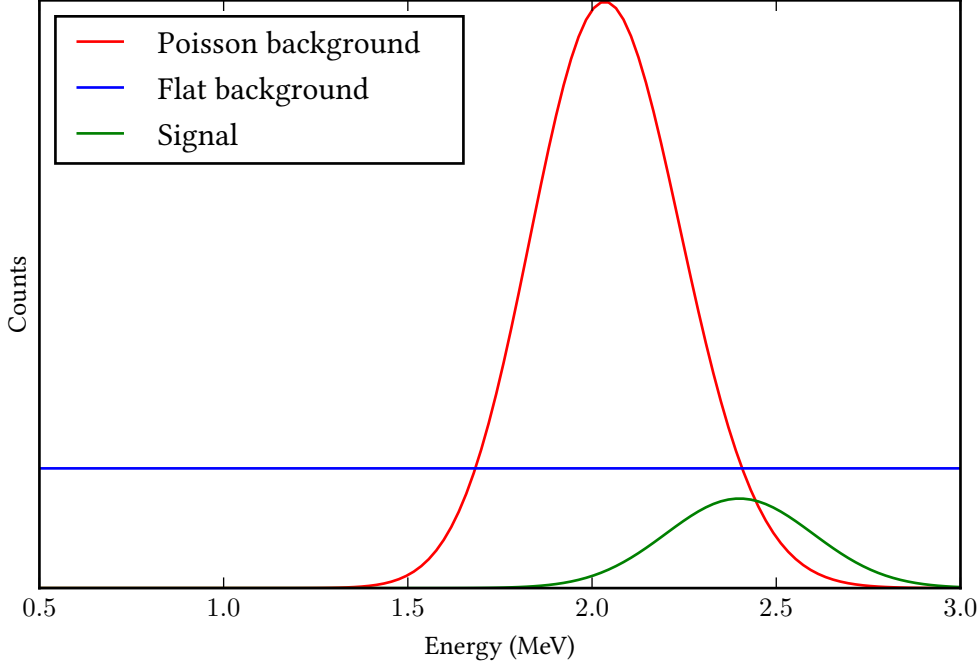


Figure 4.1: Spectral plot, showing the spectral shapes of the Poisson background, flat background and signal, that we have simulated to use in the toy model.

- K bins in observable x , width $w_k = \frac{x_u - x_0}{K}$
- L bins in observable l , width $w_l = \frac{y_u - y_0}{L}$
- M bins in observable m , width $w_m = \frac{z_u - z_0}{M}$

Taking a one-dimensional slice along a given dimension, or the full array if considering a single observable, yields an array of bin-contents:

$$\mathbf{n}_k = (n_{0lm}, n_{1lm}, \dots, n_{klm}), \quad \text{for } 0 \leq k < K. \quad (4.1)$$

Slices along the dimensions of the other observables, if applicable, have similar definitions.

The contents of a single bin are:

$$n_{klm} = \int_{x_0 + w_{kk}}^{x_0 + w_{k(k+1)}} \int_{y_0 + w_{ll}}^{y_0 + w_{l(l+1)}} \int_{z_0 + w_{mm}}^{z_0 + w_{m(m+1)}} \sum_{i=0}^{N-1} \delta(x - x_i) \delta(y - y_i) \delta(z - z_i) dx dy dz \quad (4.2)$$

There are N events in total so summing over the contents from all bins gives:

$$\sum_{k=0}^{K-1} \sum_{l=0}^{L-1} \sum_{m=0}^{M-1} n_{klm} = N. \quad (4.3)$$

Now imagine one could describe the data \mathbf{x} using a model, characterised by a Probability Density Function (PDF). The PDF $P(x|\boldsymbol{\alpha})$ evaluates the probability of obtaining \mathbf{x} , given the model parameters $\boldsymbol{\alpha}$ —which are in the fit parameters discussed in the following section—where:

$$\boldsymbol{\alpha} = (\alpha_0, \alpha_1, \dots, \alpha_k), \quad \text{for } 0 \leq j < J. \quad (4.4)$$

As with \mathbf{x} , the PDF comprises a set of measurements in the observables (x, y, \dots) , but here the measurements (or events¹) derive from Monte Carlo (MC) simulation.

The combination of all simulated MC for a unique set of α -values, forms the model to test, which requires a similar formalism to the data spectrum. Binning in the same way as the data, the array of expected values is² $\mathbf{y} = (x, y, z)$. It contains:

- K bins in observable x , width $w_k = \frac{x_u - x_0}{K}$
- L bins in observable l , width $w_l = \frac{y_u - y_0}{L}$
- M bins in observable m , width $w_m = \frac{z_u - z_0}{M}$

and the bin contents are:

$$\mathbf{y}_k = (y_0, y_1, \dots, y_k), \quad \text{for } 0 \leq k < K. \quad (4.5)$$

4.2 Overview of echidna

We designed the echidna package, with the aim of creating a fast, flexible and user-friendly software packages, specifically oriented towards spectrum fitting and limit-setting tasks. To this end, we chose to write the package in python, for its flexibility and relative simplicity

¹Given its particle physics context, we will use “events” in preference to “measurements” for the rest of the thesis, but bear in mind that the following derivations are valid for measurements of any observable obeying Poisson statistics.

²For simplicity, we show a single dimension—the observable x —but one can expand the formalism, and the following derivations, to three (or more) observables.

from the perspective of new users wishing to contribute to the project. Python also has a rich set of modules that provide interactivity, notably matplotlib’s pyplot module, which we have implemented—alongside ROOT—to handle standard output plots. This module also works well with NumPy and SciPy, two modules used extensively in echidna to achieve the desired “fast” attribute of the software. NumPy is a fundamental extension to python for scientific computing [81]. Centred around the `ndarray`³ object, NumPy provides a powerful collection of routines for N-dimensional array manipulation—including mathematical, statistical and shape transformation. The developers write most of these in a compiled language such as C or C++ to optimise execution speed. Compiled routines are also a key feature of SciPy, a package that provides key features of scientific computing: integration, interpolation and Fast Fourier Transforms, among others [82]. Statistics and optimisation routines, which form an integral part of spectrum fitting in echidna, are also available.

This chapter gives a detailed description of the echidna package and how one can use it as a limit-setting tool. Using echidna as a limit setting tool is predominantly a two-stage process, Section 4.3 explores the first of these. It involves echidna’s core `Spectra` class and how we configure and convolve them—whether signal or background—ready for fitting. One can find explain the latter in Section 4.4 and Section 4.5, describing the underlying statistical formalism and the full fitting algorithm. Section 4.6 then develops this into a full limit-setting algorithm.

4.3 Spectra creation and manipulation

4.3.1 The `Spectra` class

The core of the echidna package lies in the `Spectra` class as this acts as a container for the spectral data, and provides useful methods for basic manipulation of the data. We constructed the class with NumPy’s `ndarray` at its foundations, stored as the `_data` member of the class. The array has a dimension for each spectral parameter required—often the visible energy and the simulated radial position of the event vertex. One can include any number of dimensions and can fit in any number of these dimensions, but naturally the more dimensions included, the longer the fit will take. Large arrays can also raise a `MemoryError` if they exceed NumPy’s memory allocation for arrays.

³We will format all names of relevant python objects (modules, classes, methods and functions) like this

The array contains the binned number of events in each unique combination of the spectral parameters. Users can manipulate the data contained in this `gls-ndarray` directly from the `Spectra` class, for example scaling the spectrum, which changes its normalisation; shrinking, which slices the array; and projecting the array onto one or more dimensions. These three are integral to the fitting and limit setting algorithms so we will explain them in more detail below. Other methods of the `Spectra` class provide useful manipulations and book-keeping functionality.

Before examining the `Spectra` class in greater detail, we will note the similarity between what we are trying to do with `echidna` and `NumPy` and CERN's `ROOT` package and motivate the choice to design our own package. An analyser could achieve most of the aims of the `echidna` package using `ROOT`, using the `TH2` class as a container to store spectral-dimension information for each event. `ROOT` also works well with `Minuit` as a minimiser. Yet a sophisticated framework for limit-setting—using `ROOT` as the primary data-handling package—already exists within the SNO+ Collaboration, so we designed `echidna` partly to cross-check the results of this framework. Moreover, we designed `echidna` with flexibility and ease of use in mind—a design goal it shares with `NumPy`—so we hope the package will also be useful in its own right. The main advantage of `echidna`'s approach, over using `ROOT`, is that `NumPy`'s `ndarray` is better suited as a data container for spectral fitting than `ROOT`'s histogram classes.

4.3.2 Creation and configuration

The location of the `Spectra` class, in `echidna`, is the `spectra` module. To create a `Spectra` instance, users should import the module and then call the `Spectra` constructor. This requires two pieces of information, as arguments. The first is the configuration of the `Spectra`, which we will explain in the rest of this sub-section, and the second is the `num_decays` parameter, which defines the number of decays the `Spectra` should represent. Section 4.3.4 explains this in more detail, along with how to scale `Spectra`.

When creating a `Spectra`, users must specify its dimensions, by providing a `SpectraConfig` instance. We use configs in `echidna`, wherever there is a need to specify and manage a set of parameters. The two main cases are: managing the spectral parameters—as is the case here—and managing the parameters in the fit, which we discuss further in Section 4.7. We define all configs in the `core.config` module, of `echidna`, and all the parameters one could include in a config, in the `parameter` module. In any config, we store the parameter instances in an

OrderedDict to preserve their ordering. We provide methods to add and access parameters, whether by name or index, as well as obtaining the shape of the config or a parameter's type.

The recommended prescription for creating a `SpectraConfig` instance, is to use the `load_from_file` class method. This allows one to create a `SpectraConfig`, from a YAML file (a human-readable computing language for data serialisation), listing the parameters and their properties, for example:

```
parameters:
  energy_mc:
    low: 0
    high: 10
    bins: 1000
  radial_mc:
    low: 0
    high: 15000
    bins: 1500
```

Providing the path to this file, as an argument to `load_from_file`, would return a `SpectraConfig` instance containing two spectral parameters: MC energy and MC radius. Then using this to create a `Spectra` instance, will result in a `Spectra`, with two dimensions—energy and radius. The config will divide the energy dimension into 1000 bins, between 0 to 10 MeV, and divide the radial dimension into 1.5×10^3 bins, ranging from 0 to 15×10^3 mm.

4.3.3 Filling spectra

Now we have covered how to create and configure a `Spectra` instance, we will present how to fill a `Spectra` with data. The relevant method here is the `fill` method, which is similar to the `Fill` method of ROOT histograms. Users pass values for observables as keyword-arguments corresponding to spectral dimensions. These values then locate the appropriate bin to fill in each spectral dimension. By default this increments the contents of the selected bin by one unit—a single event—but, mirroring ROOT, passing a weight allows the user to fill in more than one event at a time, or fractions of events.

4.3.4 Spectra scaling and integration

In some aspects of echidna we attempted to emulate the functionality seen in ROOT, so that similar commands or logic perform tasks common both packages. However scaling echidna Spectra, is one instance where we applied a different logic to ROOT. The parallel container in ROOT is (assuming two spectral dimensions) a TH2D. For a TH2D, the Scale method multiplies the contents of each bin by a constant value c_1 (a float), so to increase the integral of a TH2D $c_1 > 1$ and to decrease the integral $c_1 < 1$.

In echidna, the ethos is that each full Spectra should represent some number of decays N . Often during fitting problems, theoretical predictions or models provide the expected number of decays in a given livetime. In general, if one simulates N decays, using a full-detector MC, the number of reconstructed (observed) events N' in the simulation will be less than N , because the detector has a certain efficiency (ϵ_{det}). Then, when fitting the spectrum, we might cut the spectrum to a given fit region, introducing a further efficiency in each dimension, e.g. ϵ_x and ϵ_y , for a two-dimensional spectrum in observables x and y . We can still say that the full spectrum represents N decays, and the relation:

$$N' = \epsilon_{det}\epsilon_x\epsilon_y N, \quad (4.6)$$

holds.

Note, one has to be a little cautious with this logic when it comes to a spectrum of observed data points. We have no prior knowledge of the number of physical decays that produced this spectrum; that is the question fitting aims to answer. So, in this instance, the number of decays the spectrum should represent is the total number of observed decays in the data spectrum. In any case, as the following sub-section will discuss, one should never change the normalisation of the data spectrum in the fitting algorithm.

With this in mind, the Spectra.scale method also accepts a single constant (float), but here this is the number of decays the Spectra should now represent. If the number of decays before scaling is N_0 and N is the desired number of decays it should represent, the scale method carries out the following operation:

$$\mathbf{y}' = \mathbf{y} \frac{N}{N_0}, \quad (4.7)$$

such that the new binned array \mathbf{y}' represents N decays.

4.3.5 Convolutions

When studying detector systematics, such as energy resolution, one commonly models the detector response of the systematic, whether analytically or empirically, and then uses this function to convolve spectra. A general convolution is the product of two functions— f and g . Over a finite range, one can express the convolution of f and g as [83]:

$$[f * g] = \int_0^\theta f(x)g(\theta - x)dx, \quad (4.8)$$

where x is an observable measured in the range $0 \leq x \leq \theta$. For complete treatment of detector systematics one can form a full detector response function. KamLAND-Zen [49] took this approach in the analysis we refer to in Chapter 5. There are ongoing studies to build up a detector response function for SNO+, but for echidna we decided to simplify a little and study the key detector systematics. At the time of writing echidna supports treatment of resolution (in both energy and radial position), energy scale and shifts in energy.

Using the same notation introduced in Section 4.1 a convolved model spectrum $y(x)$ is:

$$y * f(x) = \int_a^b y(x)f(b - x)dx, \quad (4.9)$$

where $f(x)$ is the convolution function and a and b are the upper and lower bounds, for a fit in observable x . By first making the substitutions $a = x_0$ and $b = x_0 + Kw$, and then considering the integral as a sum over all bins in the spectrum, Equation (4.9) becomes:

$$y * f(x) = \sum_{k=0}^{K-1} y_k \int_{x_0+kw}^{x_0+(k+1)w} y(x)f(b - x)dx. \quad (4.10)$$

Rather than the full spectrum, one can express the contents of a single bin using:

$$y'_{k'} = \int_{x_0+k'w}^{x_0+(k'+1)w} y * f(x')dx', \quad (4.11)$$

$$\Rightarrow \sum_{k=0}^{K-1} y_k \underbrace{\int_{x_0+k'w}^{x_0+(k'+1)w} \int_{x_0+kw}^{x_0+(k+1)w} f(b - x)dx dx'}_{F_{kk'}}. \quad (4.12)$$

The double integral in Equation (4.11) is a matrix element $F_{kk'}$ corresponding to the contribution of from bin k in the convolution function, on bin k' in the spectrum. Making this

substitution condenses Equation (4.11) to:

$$y'_{k'} = \sum_{k=0}^{K-1} y_k F_{kk'}, \quad \text{for } k' = \text{const.} \quad (4.13)$$

Noting that this is the definition of a scalar product, gives:

$$y'_{k'} = y_k F_{kk'}, \quad (4.14)$$

for each bin of the spectrum or, for the full spectrum:

$$\mathbf{y}' = \mathbf{y} \cdot \mathbf{F}. \quad (4.15)$$

With this mathematical foundation, it becomes easy to implement different convolutions algorithmically. This involves interchanging the convolution function and producing the matrix F . For energy resolution, we have implemented both Gaussian and Poisson smearing functions. In the Gaussian regime:

$$f(y_k) = \frac{\exp \frac{-(y_k - \mu)^2}{2\sigma^2}}{\sigma \sqrt{2\pi}}, \quad (4.16)$$

where y_k is the bin contents of the bin the code is smearing and μ and σ depend on the value of the observable. In most cases the observable is energy and here we calculate sigma from the expected light yield ly of the detector and energy of the current bin E_k :

$$\sigma = \sqrt{\frac{E_k [\text{MeV}]}{ly [\text{NHit/MeV}]}}.$$

Figure 4.2[a.] shows an example of a Gaussian convolution. In the Poisson regime, the Poisson distribution for energy y_k is:

$$f(y_k) = \frac{\nu^{n(y_k)} \exp -\nu}{n(y_k)!}, \quad (4.17)$$

where, in most cases y_k is visible-energy E . The Poisson distribution deals in discrete numbers of events, so we convert from visible energy to events, by:

$$n(E) = E [\text{MeV}] \cdot ly [\text{NHit/MeV}],$$

and:

$$\nu = E_k [\text{MeV}] \cdot l_y [\text{NHit/MeV}],$$

is the expected number of events for the current bin. Figure 4.2[b.] illustrates echidna's Poisson convolution.

The remaining two convolutions are both centred around the uniform distribution. For a shift by a , in y_k :

$$f(y_k) = \begin{cases} \frac{1}{y_u - y_0} & (y_0 + a) \leq y_k \leq (y_u + a), \\ 0 & \text{otherwise,} \end{cases} \quad (4.18)$$

and for a scale convolution by scale factor b :

$$f(y_k) = \begin{cases} \frac{1}{b(y_u - y_0)} & b(y_0) \leq y_k \leq b(y_u), \\ 0 & \text{otherwise.} \end{cases} \quad (4.19)$$

For examples of these two convolutions, see Figure 4.2[c. and d.]. These four convolutions enable us to handle most systematic uncertainties in echidna.

4.3.6 Data storage

Rather than use ROOT's `TFile` format, we chose to use a similar hierarchical data format—HDF5. We chose to use the HDF5 format because it works well with `ndarrays` and complex python objects, where `ndarrays` are a major component. The `store` module of echidna, takes responsibility for data storage. One can use it to store `Spectra` instances, as well as results from fitting and limit-setting.

4.4 The Poisson likelihood chi-squared

Returning to the formalism Section 4.1 introduced, this section derives the Poisson Likelihood χ^2 (χ_λ^2), which is the main test-statistic we use in echidna. Section 4.4.1 defines the likelihood, Section 4.4.2 then proceeds to derive the extended likelihood. Section 4.4.3 extends the definition to a binned likelihood and, following this, is an explanation of how to get the χ_λ^2 , from a likelihood ratio (Section 4.4.4). The derivations that follow are most applicable to

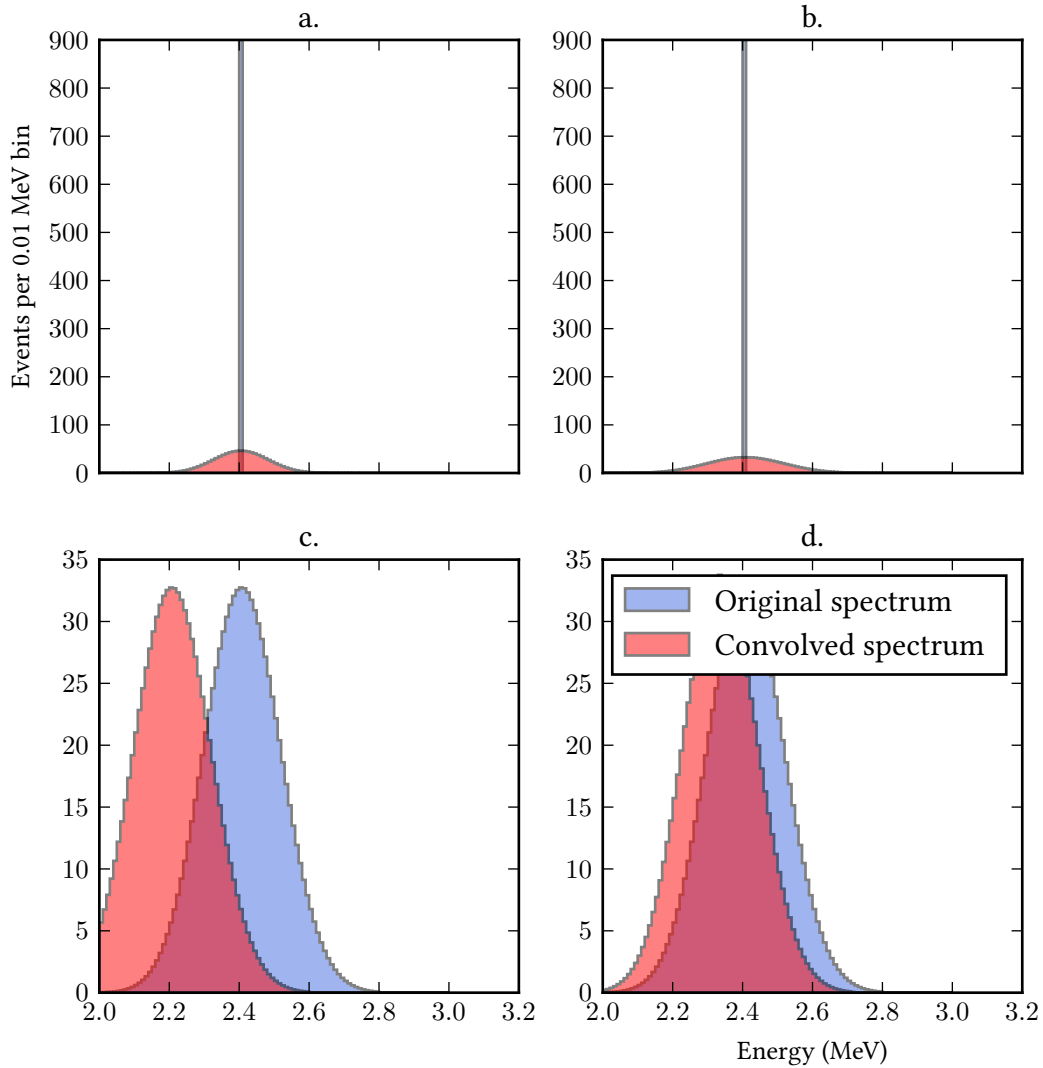


Figure 4.2: Overlaid histograms showing the effect of the four convolutions we have implemented in echidna. In each case the input is the blue curve and the convolved spectrum is the red curve. The convolutions are: a. Gaussian smearing, with mean energy 2.4 MeV and energy resolution of 5 %; b. Poisson smearing, also with mean energy 2.4 MeV and energy resolution of 5 %; c. Shift by -0.2 MeV; d. Scale by factor of 0.97.

measurements of a single variable—the visible energy. Throughout the derivation of the χ^2_{λ} , we assume a pure Poisson energy resolution for the visible energy

4.4.1 Likelihood

To construct a general expression for the likelihood we normalise the PDF such that:

$$\int_{x_0}^{x_u} P(x|\boldsymbol{\alpha})dx = 1, \quad x_0 \leq x \leq x_u. \quad (4.20)$$

Given the data (\mathbf{x}), PDF ($P(x|\boldsymbol{\alpha})$) and model parameters, the likelihood is:

$$\mathcal{L}(x|\boldsymbol{\alpha}) = \prod_{i=1}^N P(x_i|\boldsymbol{\alpha}). \quad (4.21)$$

The best-fit values for the parameters $\boldsymbol{\alpha}$, are those that maximise the likelihood. This technique is known as Maximum Likelihood Estimation and the best-fit values it produces are Maximum Likelihood Estimators (MLEs).

Algorithmically, minimising $-\log \mathcal{L}(x|\boldsymbol{\alpha})$ is more convenient than maximising the likelihood. In this case Equation (4.21) becomes:

$$-\log \mathcal{L}(x|\boldsymbol{\alpha}) = -\sum_{i=1}^N P(x_i|\boldsymbol{\alpha}). \quad (4.22)$$

4.4.2 Extended likelihood

Often, the number of events measured is not fixed but an observable of the experiment. This requires the following normalisation of the PDF:

$$\int_{x_0}^{x_u} P(x|\boldsymbol{\alpha})dx = A(\boldsymbol{\alpha}), \quad (4.23)$$

where $A(\boldsymbol{\alpha})$ ⁴ is the total number of events predicted by the model—naturally this is dependent on the model parameters. Now the normalised likelihood of each event is $P(x|\boldsymbol{\alpha})/A$ and the total likelihood is:

$$\mathcal{L}(x|\boldsymbol{\alpha}) = Q(N|A) \prod_{i=1}^N \frac{P(x_i|\boldsymbol{\alpha})}{A}. \quad (4.24)$$

The PDF $Q(N|A)$ accounts for the distribution of the total number of measured events.

⁴We will use A to represent the normalisation of the PDF from now on, where A implies the dependence on $\boldsymbol{\alpha}$ unless stated otherwise.

Since the events are independent of each other and of the observable x , one can assume fluctuations in N follow a Poisson distribution. Take the true mean of this distribution as the total number of events predicted by the model (A), such that:

$$Q(N|A) = \frac{A^N e^{-A}}{N!}, \quad (4.25)$$

On substituting into Equation (4.24), this gives:

$$\mathcal{L}(x|\boldsymbol{\alpha}) = \frac{A^N e^{-A}}{N!} \prod_{i=1}^N \frac{P(x_i|\boldsymbol{\alpha})}{A}. \quad (4.26)$$

As with the standard likelihood, working with the negative log likelihood is more convenient:

$$-\log \mathcal{L}(x|\boldsymbol{\alpha}) = -\sum_{i=1}^N \log \left(\frac{P(x_i|\boldsymbol{\alpha})}{A} \right) - N \log A + A + \log N!. \quad (4.27)$$

Given:

$$\begin{aligned} -\sum_{i=1}^N \log \left(\frac{P(x_i|\boldsymbol{\alpha})}{A} \right) &= -\sum_{i=1}^N \log P(x_i|\boldsymbol{\alpha}) + \sum_{i=1}^N \log A, \\ &\Rightarrow -\sum_{i=1}^N \log P(x_i|\boldsymbol{\alpha}) + N \log A, \end{aligned}$$

the two $N \log A$ terms cancel. In the context of minimising the negative log likelihood, the term $\log N!$ is independent of the model parameters and so remains constant during the minimisation. Ignoring this term, the function that is actually minimised is:

$$-\log \mathcal{L}(x|\boldsymbol{\alpha}) = -\sum_{i=1}^N \log P(x_i|\boldsymbol{\alpha}) + A. \quad (4.28)$$

4.4.3 Binned maximum likelihood

Up to this point, this derivation has considered event-by-event data, but often binning data from a large number of events to form a histogram, is preferable. What follows, extends the derivation of the extended likelihood from Equation (4.28), from the previous section, for

binned events. Using the formalism described in Section 4.1, one represents the likelihood as⁵:

$$\mathcal{L}(\mathbf{n}|\mathbf{y}) = \frac{A^N e^{-A}}{N!} \prod_{k=0}^{K-1} P(n_k|y_k), \quad (4.29)$$

where, $P(n_k|y_k)$ is the PDF that describes the probability of obtaining n_k events in a given bin, given a model that predicts y_k events for the same bin. This expression retains the Poisson factor multiplying the PDF, since the total number of events N , still follows a Poisson distribution.

To derive the likelihood in this case, express the bin contents of the data histogram (n_k) as a multinomial distribution:

$$P(n_k|N, p_k) = \frac{N!}{n_0! \dots n_k!} p_0^{n_0} \dots p_k^{n_k}, \quad \text{when } \sum_{k=0}^{K-1} n_k = N. \quad (4.30)$$

The best estimate for the probability p_k of a bin containing n_k events is the normalised probability from ($P(x|\alpha)$), which, for a binned PDF, is the bin contents:

$$p_k = \frac{y_k}{A}, \quad (4.31)$$

The PDF becomes:

$$P(n_k|N, p_k) = \frac{N!}{n_0! \dots n_k!} \left(\frac{y_0}{A}\right)^{n_0} \dots \left(\frac{y_k}{A}\right)^{n_k}. \quad (4.32)$$

Substituting this into Equation (4.29), the likelihood becomes:

$$\mathcal{L}(n_k|y_k) = \frac{A^N e^{-A}}{N!} \frac{N!}{n_0! \dots n_k!} \left(\frac{y_0}{A}\right)^{n_0} \dots \left(\frac{y_k}{A}\right)^{n_k}. \quad (4.33)$$

Noting that:

$$\left(\frac{1}{A}\right)^{n_0} \dots \left(\frac{y_k}{A}\right)^{n_k} = \left(\frac{1}{A}\right)^{\sum_{k=0}^{K-1} n_k} = \left(\frac{1}{A}\right)^N, \quad (4.34)$$

⁵I've continued to show a single dimension, but note that one could expand this to cover more than one dimension.

cancelling the factors of A^N and $N!$ simplifies the likelihood expression:

$$\mathcal{L}(n_k|y_k) = \frac{e^{-A}}{\underbrace{n! \dots n_k!}_{\prod_{k=0}^{K-1} n_k!}} \prod_{k=0}^{K-1} y_k^{n_k}. \quad (4.35)$$

Then, recognising that:

$$e^{-A} = e^{\sum_{k=0}^{K-1} y_k} = \prod_{k=0}^{K-1} e^{-y_k}, \quad (4.36)$$

yields, the following expression for the likelihood:

$$\mathcal{L}(n_k|y_k) = \prod_{k=0}^{K-1} \frac{e^{-y_k} y_k^{n_k}}{n_k!}. \quad (4.37)$$

As with the un-binned case, minimising negative log likelihood is preferable to maximising the likelihood as given in Equation (4.37). Taking the negative log of both sides gives:

$$-\log(\mathcal{L}(\mathbf{n}|\mathbf{y})) = -\sum_{k=0}^{K-1} (n_k \log y_k + \log n_k!) + \underbrace{\sum_{k=0}^{K-1} y_k}_A. \quad (4.38)$$

Again, one can ignore terms with no dependence on α that will remain constant during minimisation, so do not contribute. Since n_k has no dependence on α , $-\sum_{k=0}^{K-1} \log n_k!$ is one such term. This leaves:

$$-\log(\mathcal{L}(\mathbf{n}|\mathbf{y})) = -\sum_{k=0}^{K-1} (n_k \log y_k) + A. \quad (4.39)$$

Notice the similarity between Equation (4.39) and Equation (4.28); both having the $A(\alpha)$ term.

4.4.4 The Poisson likelihood χ^2

This final part of the discussion of the underlying statistical framework, continues from the binned extended likelihood to derive the χ^2_{λ} . Based on a likelihood ratio, this is a test statistic that is well suited to fitting to binned data.

Let \mathbf{m} be the true (*unknown*) values of \mathbf{n} obtained in an experiment with absolute precision. One defines the likelihood ratio (λ):

$$\lambda = \frac{\mathcal{L}(\mathbf{n}|\mathbf{y})}{\mathcal{L}(\mathbf{n}|\mathbf{m})}. \quad (4.40)$$

From this ratio, the χ_λ^2 is:

$$\begin{aligned} \chi_\lambda^2 &= -2 \log \lambda \\ &= -2 \log \mathcal{L}(\mathbf{n}|\mathbf{y}) + 2 \log \mathcal{L}(\mathbf{n}|\mathbf{m}). \end{aligned} \quad (4.41)$$

For Poisson histograms, replacing the unknown \mathbf{m} with its bin-by-bin model-independent MLE—which is \mathbf{n} —yields:

$$\mathcal{L}(\mathbf{n}|\mathbf{m}) = \mathcal{L}(\mathbf{n}|\mathbf{n}) = \prod_{k=0}^{K-1} \frac{e^{-n_k} n_k^{n_k}}{n_k!}. \quad (4.42)$$

Substituting Equation (4.37) and Equation (4.42) into Equation (4.41):

$$\begin{aligned} \chi_\lambda^2 &= -2 \log \left(\prod_{k=0}^{K-1} \frac{e^{-y_k} n_k! y_k^{n_k}}{e^{-n_k} n_k! n_k^{n_k}} \right) \\ &= -2 \log \left(\prod_{k=0}^{K-1} e^{(n_k - y_k)} \left(\frac{y_k}{n_k} \right)^{n_k} \right) \\ &= -2 \sum_{k=0}^{K-1} n_k - y_k - n_k \log \left(\frac{n_k}{y_k} \right). \end{aligned} \quad (4.43)$$

Tidying a little, gives:

$$\chi_\lambda^2 = 2 \sum_{k=0}^{K-1} y_k - n_k + n_k \log \left(\frac{n_k}{y_k} \right), \quad (4.44)$$

which is the definition of χ_λ^2 we will use throughout this thesis.

The quantity one obtains using Equation (4.44) follows a chi-square distribution and one can use it for tests of goodness-of-fit; provided y_k are not small. For low y_k , one can model the chi squared (χ^2)-like distribution using MC. From Equation (4.41), it should be clear that minimising χ_λ^2 is identical to maximising the likelihood, so one can use χ_λ^2 for parameter estimation and confidence intervals, as well as goodness-of-fit testing [84]. Section 4.6 will continue to describe how one obtains a the limit on a signal contribution, using the χ_λ^2 , but

the following section will describe how we implemented this statistical framework in echidna and the algorithm it uses for spectral-fitting.

4.5 A fixed background fit

The prerequisites for a fixed-background fit in echidna are: a data spectrum, a fixed-background spectrum that represents a model to fit to the data and a test-statistic to use during the fit. In this example fixed-background fit, we will fit the fixed-background spectrum, the summed background shown in Figure 4.3 to the data spectrum, also in Figure 4.3. The `test_statistic` module of echidna implements the calculation of some different test-statistics, including all the forms of χ^2 defined in [84] and the extended log likelihood. In this example we will use the χ^2_λ because this is the test-statistic we use in the fits in later chapters.

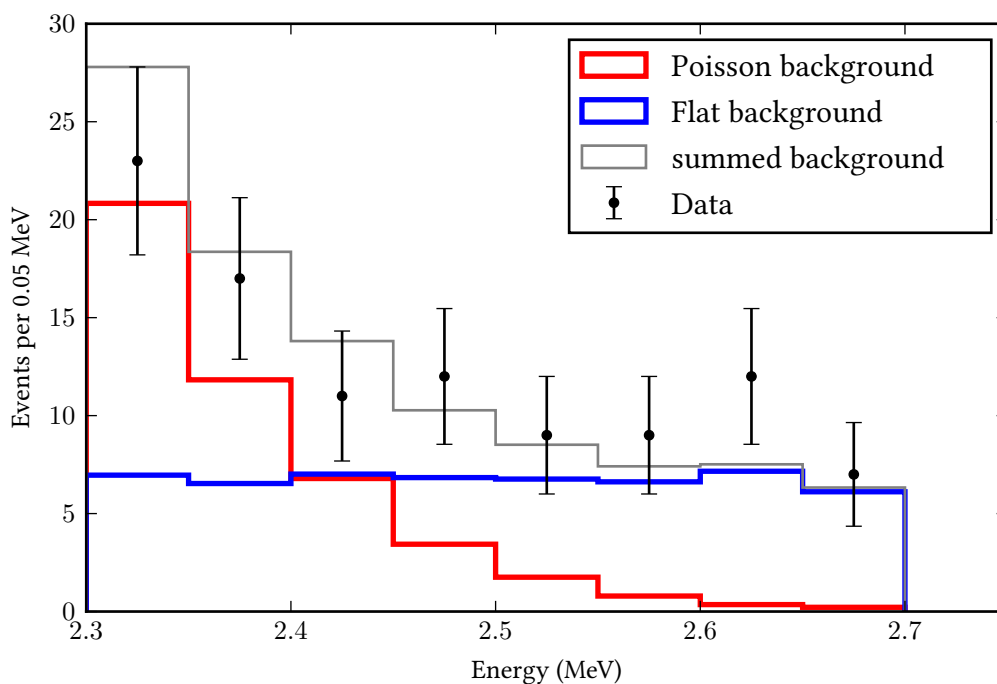


Figure 4.3: Spectral plot showing the “data” generated for the toy model, as well as the Poisson and flat backgrounds, in the 2.3 to 2.7 MeV region of interest (ROI).

The test-statistic plays an important role in quantifying the goodness-of-fit between a data spectrum and the model that describes the data. In Algorithm 4.1, we illustrate the process by which echidna determines the value of the χ^2_λ , using the definitions for n and y

that Section 4.1 establishes. Note the algorithms for the other test-statistic calculations in the `test_statistic` module follow a similar format. The algorithm assumes both `n` and `y` are one-dimensional. Within the algorithm we define the parameter ε to avoid arithmetic errors in the limits: $\lim_{y_k \rightarrow 0} \chi_\lambda^2$, $\lim_{n_k \rightarrow 0} \chi_\lambda^2$ or even $\lim_{n_k, y_k \rightarrow 0} \chi_\lambda^2$. Considering the first limit:

$$\lim_{y_k \rightarrow 0} \chi_\lambda^2(n_k, y_k) = n_k \log(n) - n_k - n_k \lim_{y_k \rightarrow 0} \log(y_k) = +\text{inf}. \quad (4.45)$$

This motivates the first use of ε . This limit requires the test statistic to diverge to $+\text{inf}$ which setting $y_k = \varepsilon$ achieves, without the need to deal computationally, with inf . In the second limit $n_k \log(n_k)$ is indeterminate. By writing this term as a quotient and using l'Hôpital's rule, the limit becomes:

$$\lim_{n_k \rightarrow 0} \chi_\lambda^2(n_k, y_k) \equiv y_k - \lim_{n_k \rightarrow 0} n_k = y_k, \quad (4.46)$$

which explains why the algorithm sets $\chi_{\lambda,k}^2 = y_k$, where $n_k < \varepsilon$. The third limit is trickier to evaluate as the only non-zero term is a function of both y_k and n_k . By temporarily re-defining n_k and y_k as functions of a common variable t — $n_k = at$ and $y_k = bt$ —and then using l'Hôpital's rule on the quotient of two functions of t :

$$\lim_{n_k \rightarrow 0} \chi_\lambda^2(n_k, y_k) \equiv - \lim_{t \rightarrow 0} \frac{bt^2}{t} = 0. \quad (4.47)$$

This is the result achieved if $n_k, y_k < \varepsilon$ in Algorithm 4.1.

Algorithm 4.1 Calculating the χ_λ^2 : single value

Input: `n, y`

Output: χ_λ^2

$\chi_\lambda^2 = 0$

$\varepsilon = 1.0 \times 10^{-34}$

for $k = 0$ **to** $k = K - 1$ **do**

if $y_k < \varepsilon$ **then**

$y_k = \varepsilon$

if $n_k < \varepsilon$ **then**

$\chi_{\lambda,k}^2 = y_k$

else

 Calculate log-likelihood contribution from this bin

$\chi_{\lambda,k}^2 = y_k - n_k + n_k \log\left(\frac{n_k}{y_k}\right)$

$\chi_\lambda^2 \rightarrow \chi_\lambda^2 + \chi_{\lambda,k}^2$

return χ_λ^2

Algorithm 4.2 Calculating the χ_λ^2 array

Input: n, y

Output: χ_λ^2

Set χ_λ^2 as empty array

for $k = 0$ **to** $k = K - 1$ **do**

if $y_k < \varepsilon$ **then**

$y_k = \varepsilon$

if $n_k < \varepsilon$ **then**

$\chi_{\lambda,k}^2 = y_k$

else

 Calculate log-likelihood contribution from this bin

$\chi_{\lambda,k}^2 = y_k - n_k + n_k \log \left(\frac{n_k}{y_k} \right)$

 Append $\chi_{\lambda,k}^2$ to χ_λ^2

return χ_λ^2

Algorithm 4.1 returns a single χ^2 value, but echidna also has the ability to work with an array of test-statistic values, corresponding to each bin. Algorithm 4.2 demonstrates how we can adapt the calculation of χ_λ^2 to return an array of values. This facilitates limit-setting or goodness-of-fit analyses by exposing the sensitivity or fit contribution from each bin. Any `TestStatistic` class in echidna will return an array of values—rather than a single test-statistic—if the user passes a `per_bin=True` flag, during initialisation of the class.

Now that we have defined each of the fit components (data spectrum, fixed-background spectrum and χ^2), we begin the fit process by initialising the `Fit` class with these components. In this instance, the fixed background consists of two spectra (as shown in Figure 4.3), but we want a single array of values to represent the model y_k in Algorithms 4.1 and 4.2. We achieve this by passing a dictionary of spectra and normalisation values to the `fixed_background` keyword argument during the `Fit` initialisation. The `make_fixed_background` method then scales each spectrum to its normalisation value and sums the spectra, to create a single fixed-background spectrum, with the correct dimensions and binning.

Suppose we initialise a `Fit` instance as the object `fitter`, running the fit is then as simple as evaluating `fitter.fit()`. This will execute the process presented in Algorithm 4.3, which, for a fixed-background fit, is the block following the first if-clause. Note that when the fit is in two or more dimensions, we flatten the arrays before calculating the test-statistic and then re-shape the return array. We flatten the arrays using NumPy's `ravel` method, which takes consecutive slices along an array dimension and then chains them together. Running the fit for the spectra in Figure 4.3 yields the result $\chi_\lambda^2 = 4.548$, with seven degrees of freedom.

Algorithm 4.3 The fit method

Input: fixed-background, floating-backgrounds and data spectra; test-statistic**Output:** value of test-statistic**if** no floating-backgrounds **then** project data onto fit dimensions \rightarrow observed project fixed-background onto fit dimensions \rightarrow expected

flatten observed and expected

calculate test-statistic, using e.g. Algorithm 4.1

re-shape array, as required

return test-statistic**else**

pass to minimiser

minimise

return minimum test-statistic

4.6 Investigating signals

There are two distinct scenarios where one could be investigating a signal process: one aims to see if a set of measurements from an experiment, evidence the discovery of the signal process and quantify its rate, whilst the other estimates the potential for an experiment to detect the signal process. In this thesis we describe two analyses that use echidna to investigate signal processes—one based on each of these two scenarios. In Chapter 5, we use echidna to analyse data from the KamLAND-Zen experiment to investigate whether there is evidence of a contribution from any of the Majoron-emitting $0\nu 2\beta$ decays, whilst in Chapter 6, we estimate the potential sensitivity of the SNO+ experiment to the same signals. In this section, we will explain how one can use the Poisson Likelihood χ^2 test statistic, that we derived in Section 4.4, to investigate signal rates in either of the two scenarios.

If we want to say something about the rate of the signal, we must include this as a parameter in the fit. This gives a J -dimensional χ^2_λ surface—whose global minimum position is the best fit to the data. Standard statistics texts [85, 65], give the prescription for working out the standard error on a best-fit parameter value, and correlations between values. In most cases one can determine these graphically from contours of constant $2\Delta \log \textit{likelihood}$ or χ^2_λ . For standard errors, one should use contours of $\Delta\chi^2_\lambda = 1$.

If we want to define a confidence region on the value of a specific parameter—e.g. signal rate—we project the path that traces the minimum values of the χ^2_λ surface onto the parameter axis. This leaves us with a χ^2_λ curve like Figure 4.4. The standard method for finding the confidence interval is to locate the points where the curve intercepts a given value of χ^2_λ . We

stick to the convention set by other $0\nu 2\beta$ experiments and quote a 90 % confidence interval. The appropriate value of $\Delta\chi_\lambda^2$ in this case is 2.71 [65]. Note subtracting the minimum χ_λ^2 to get a $\Delta\chi^2$ and projecting onto the parameter axis, leaves a single degree of freedom.

The method we have described above works well for a scenario where H_0 is significantly disfavoured by the goodness-of-fit test, but for sensitivity studies, or where the goodness-of-fit test yields a positive result, what can we conclude about the signal rate? In either of these cases, quoting a one-sided 90 % confidence limit is more appropriate—an upper limit on the rate of the signal (or an expected upper limit). Figure 4.4 shows example $\Delta\chi_\lambda^2$ curves, projected onto the signal rate axis, for a fit to data (blue) and a sensitivity study (red). The method is the same as before, we still take the intersection with the appropriate $\Delta\chi_\lambda^2$ line (2.71 for a 90 % CL), but now we have a sole intersection giving the upper bound. Figure 4.5 shows the χ_λ^2 projection onto the signal axis, for a fixed-background fit using the toy model spectra, to the toy model data. Calculating the limit with echidna, gives a 90 % CL of 20.2 counts.

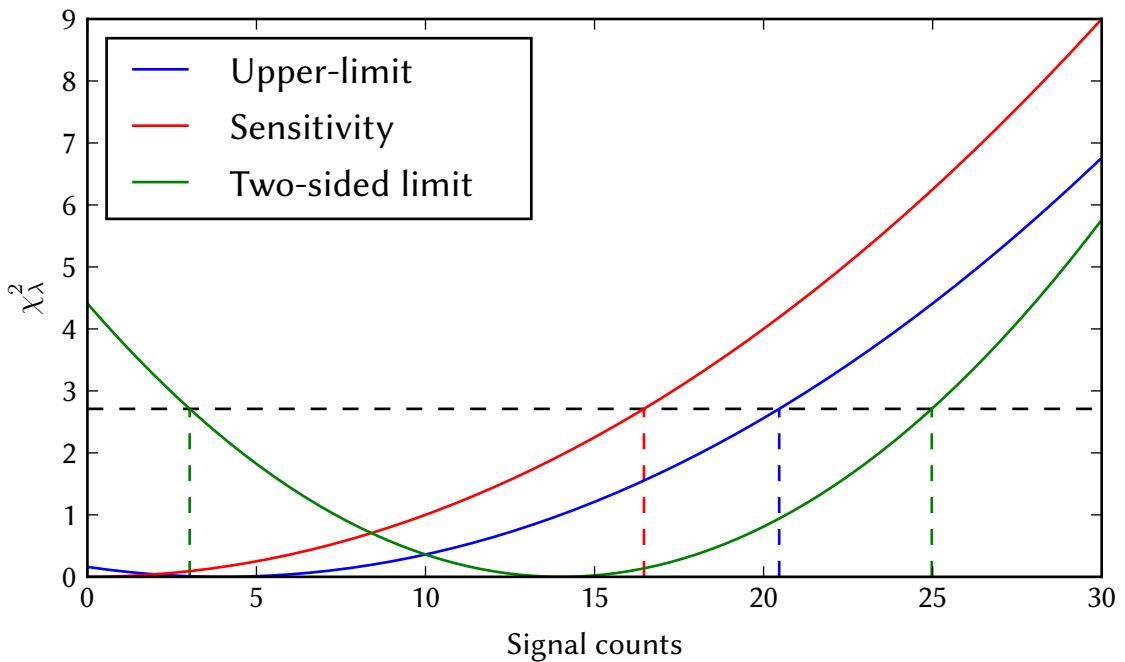


Figure 4.4: A comparison of hypothetical χ_λ^2 curves for three different scenarios. A sensitivity study (red) estimates an upper limit; a fit to data that is consistent with H_0 (blue), gives an upper limit on the signal scale; and a fit to data that significantly rejects H_0 (green), calculates a two-sided confidence interval for the signal scale.

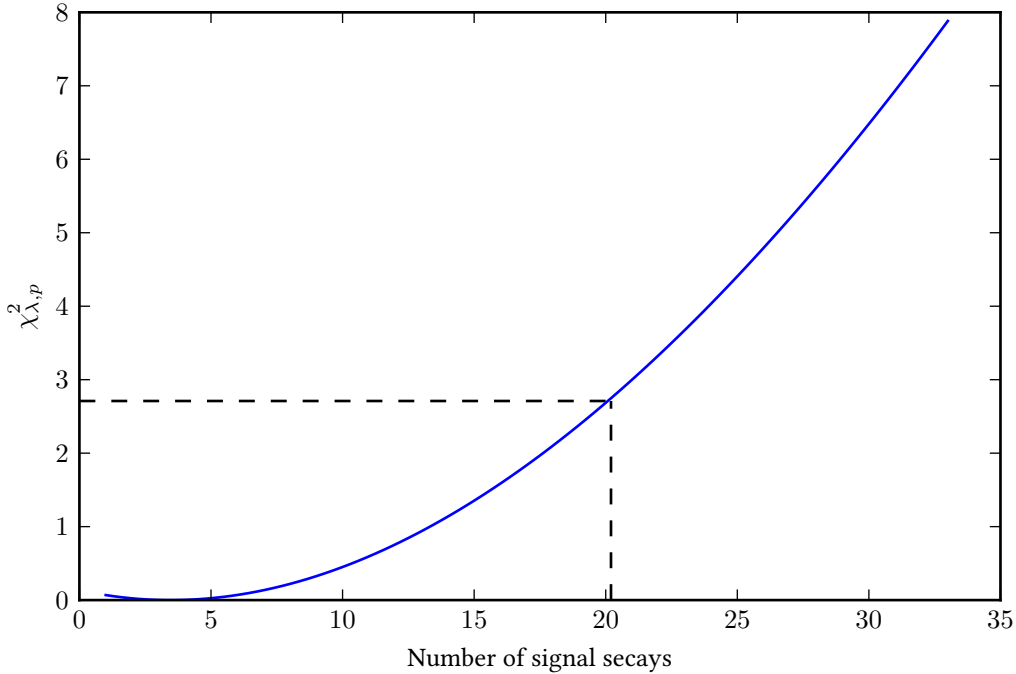


Figure 4.5: χ_{λ}^2 curve showing the limit calculated for the fixed background fit, to the toy model data.

The Bayesian method has a much more straightforward interpretation: there is a 90 % probability that the credible region contains the true value of the signal rate. Compare this to the interpretation of the Frequentist confidence interval, that says 90 % of similar experiments would produce an interval covering the true value of the signal rate. In this approach, one can convert the projection of the likelihood surface onto the appropriate parameter axis, directly to a PDF for that parameter, by normalising with a prior probability distribution. The established default choice is a “flat” or uniform prior [86], but the choice of prior can cause some contention.

A key limitation of the χ_{λ}^2 method is that it does not guarantee the interval will have the correct Frequentist coverage, this is what the Feldman-Cousins (FC) method aims to do [87]. FC is also the more widely accepted method for reporting experimental results of this nature, which means one can make a direct comparison with other FC limits. But the method is computationally intensive, requiring analysis of an ensemble of fake datasets, a technique one could also apply to the χ_{λ}^2 method to improve its Frequentist coverage. Its foundations lie in a counting experiment, so ignoring the spectral shape, whereas the χ_{λ}^2 method uses the power of individual bins to increase sensitivity.

4.7 Floating parameters in the fit

The fit in Section 4.5 and the fits in the limit-setting algorithm in Section 4.6, are fine if we accurately know the parameter values (α). Often, this is not the case. Where we do not know the exact value of a parameter, we can introduce it to the fit as a nuisance parameter. Some examples of parameters we might add as nuisance parameters are: detector systematic uncertainties, such as energy resolution; or uncertainties in background levels.

One can introduce a nuisance parameter to the fit, by allowing it to vary over a range of values. Each value for the fit parameter will result in a different test-statistic⁶ value. As we saw in section Section 4.4, the best-fit values (α_j) are those that minimise Equation (4.44). Each nuisance parameter, that one floats in the fit, introduces a further dimension to the fit. The result is a J -dimensional surface, over which χ_λ^2 can vary, where the best-fit position corresponds to the global minimum of the χ_λ^2 surface.

Some nuisance parameters should not be able to float unconstrained, for example, certain background rate, such as internal backgrounds from the ^{238}U and ^{232}Th chain decays. We can measure the rate of these decays, using assays of the water and scintillator, and use measurements taken during the *water* and *unloaded scintillator* phases, to constrain these rate parameters. In general, where we want to constrain a parameter α_j , we have some expectation or prior knowledge of what the value of the parameter should be, so we wish to penalise the fit, if it strays too far from this expectation value. We do this by adding a penalty term to the calculated χ_λ^2 value, of the form:

$$\chi^2 = \chi_\lambda^2 + \left(\frac{\alpha_j - \hat{\alpha}_j}{\sigma_j} \right)^2. \quad (4.48)$$

In Equation (4.48), $\hat{\alpha}_j$ is the expectation, or prior value for α and σ_j is an estimate of the uncertainty in the expected value. The greater the deviation from $\hat{\alpha}_j$, the larger the penalty term, and thus the total χ^2 . In a minimisation scenario, this disfavors a best-fit that deviates too far from $\hat{\alpha}_j$.

⁶We will only consider χ_λ^2 in this chapter, but note one could apply the following arguments to any valid test-statistic

4.8 Minimisation

The goal of the minimiser, as its name suggests, is to minimise the value of the test-statistic. We will write the rest of this section in the test statistic χ^2_λ , but note that one could use any valid test-statistic. The minimiser should find the lowest value of χ^2_λ , by varying the values of the parameters that describe the model (α_j). It should also find the values of α that correspond to the minimum position. Algorithm 4.4 outlines the key parts of any minimisation routine. One must supply a function to minimise and a set of parameters on which the function depends. The crucial part is the iterative section. For each iteration, the minimiser passes a set of allowed parameter values α_j to the function, which returns the value of χ^2_λ based on those values. The minimiser should use some mechanism to compare the χ^2_λ values, returned by the function, to determine when it reaches a minimum value. Note, the function to minimise, should take, as input, a list of numerical values. It receives no information on the parameter to which each value corresponds. We solve this, in echidna, by storing the parameters in an `OrderedDict`, in the `GlobalFitConfig` instance. Its method `get_pars` *always* returns the parameters in the same order, so we pass them in this order, and then deal with the parameter values in this order.

Algorithm 4.4 Format of a typical minimiser

Input: function to minimise (`_funct`), parameters (α), and any constraints or other options

Output: $\chi^2_{\lambda, min}$

Iterative process

while Minimum value not found **do**

 Generate set of parameter values α_j to test

$\chi^2_\lambda \rightarrow \text{funct}(\alpha_j)$

return $\chi^2_{\lambda, min}$

We recommend that all fits in echidna use the same callable function, which is supplied in the `Fit` class, in the form of the `_funct` method. This method accepts a list of parameter values. It first iterates over the list of values, matching each to its corresponding parameter. We use the value's index, in the list, and the `get_par_by_index` method, to select each parameter, and then the `set_current_value` method, to match the value to the parameter. We begin to build up an `ndarray` of expected values (called `expected`), by projecting the fixed background spectrum, onto all the spectral dimensions we are including in the fit. We then loop through each floating background spectrum and apply the current values of all the global parameters—such as energy shift. We can apply the global parameters either on the fly or by loading a

pre-convolved spectrum. We follow this by applying any spectrum-specific parameters, for example a rate.

Once each spectrum has the appropriate parameters applied, we shrink it to the correct ROI(s) and rebin, to match the data spectrum. We project each spectrum onto all spectral dimensions that are in the fit, sum all the projections together and add this to our existing expected array. If we have a signal spectrum, in the fit, we apply the global and spectral (rate) parameters, and project the signal spectrum onto all the fitted spectral parameters. Once we add this projection to the expected array, it becomes the y_k values in Equation (4.44). We project the data spectrum, to form an array of the same dimensions, which is our observed counts—the n_k values in Equation (4.44). We calculate the value of χ^2_λ , by passing these two arrays to the `compute_statistic` method, of the `TestStatistic` instance. Finally, we add any penalty terms to the χ^2_λ value, to get the result for this set of parameter values.

The default minimiser in echidna is `GridSearch`, a brute-force approach to minimisation problems. Suppose for each of the J parameters (α) that describe the model (see Equation (4.4)), we define a set of allowed values, based on a minimum ($\alpha_{j,min}$), maximum ($\alpha_{j,max}$) and the desired number of values. `GridSearch` constructs a J -dimensional grid, that has one point at every point for each unique combination of allowed α_j values. `GridSearch` evaluates `_func` at every point on the grid, and then determines the minimum and minimum position, using NumPy's `nanmin` and `argmin` functions. The more parameters one wishes to include, and the more allowed points to test for each parameter, the larger the grid, and the slower the minimisation. One also runs the risk of exceeding NumPy's memory allocation, for `ndarrays`, raising a `MemoryError`. To avoid these problems, one has the option to use ROOT's MINUIT or SciPy's `optimize` library to handle the minimisation. For all the analyses presented in this thesis, the `GridSearch` optimiser is good enough.

Chapter 5

Verification of echidna as a limit setting package

The goal of this chapter is to present a study that verifies echidna (described in Chapter 4) as a tool for limit-setting. The study tests echidna to see if it can reproduce the published results from another neutrinoless double-beta decay ($0\nu 2\beta$) experiment, similar to SNO+. The experiment chosen was KamLAND-Zen—a Xenon-based liquid scintillator detector—which has similarities to SNO+. As noted in Section 1.6, KamLAND-Zen has already published limits on the main Majoron-emitting modes of $0\nu 2\beta$, in ^{136}Xe [49], and in this verification study we seek to reproduce these limits. By demonstrating that echidna can reproduce these limits, we hope to add credence to its use as a limit-setting tool for SNO+, particularly in estimating its expected sensitivity to Majoron-emitting modes of $0\nu 2\beta$, which we will present in Chapter 6.

5.1 Overview of the KamLAND-Zen and its results

The KamLAND-Zen experiment, has a layered design similar to SNO+. As in SNO+, the outermost layer of the detector is water shielding—3.2 kt. KamLAND-Zen also uses photomultiplier tubes (PMTs) as to detect scintillation (and Cherenkov) light produced in the detector. It has outer PMTs in the water-shielding, which can veto events not originating in the target volume, and a spherical structure of inward facing PMTs, as is the case in SNO+. A buffer layer separates the PMT from the target volume that is fully contained within a 13-metre diameter balloon—akin to SNO+'s Acrylic Vessel (AV). Unlike SNO+'s design during the $0\nu 2\beta$ phase, this volume is filled with 1 kt pure liquid scintillator (LS). A second balloon, filled with Xe-loaded LS containing $2.44 \pm 0.01\%$ ^{136}Xe , by mass, ensuring the double-beta isotope is

more concentrated towards the centre of the detector. KamLAND-Zen also take a Fiducial Volume (FV) cut but it encompasses a larger fraction of the active volume than the SNO+ FV. After fiducialisation KamLAND-Zen contains (125 ± 7) kg ^{136}Xe .

KamLAND-Zen analysed a total of 112.3 live days of data, to determine the lifetime limits on $0\nu 2\beta$ by Majoron-emission, presented in [49]. Figure 5.1 shows the expected visible energy spectra for the Majoron-emitting modes of ^{136}Xe , with spectral indices $n = 1, 2, 3, 7$, along with the expected visible energy for the standard mode of $0\nu 2\beta$ and for two-neutrino double-beta decay ($2\nu 2\beta$). All the spectra in Figure 5.1 were convolved with a detector response function to model the energy resolution and energy-scale non-linearities of KamLAND-Zen. In the analysis presented in [49], they determined the limit on the rate of each Majoron-emitting mode, by fitting the spectrum for a single mode, along with the dominant background contribution—the $2\nu 2\beta$ spectrum. Table 5.1 summarises the results of this analysis and Figure 5.2 displays these limits in a spectral plot, along with the $0\nu 2\beta$, $2\nu 2\beta$ and combined background spectra and data points for candidate events.

Table 5.1: The KamLAND-Zen limits on the Majoron-emitting $0\nu 2\beta$ half-lives $\left(T_{1/2}^{0\nu 2\beta \chi^0(\chi^0)}\right)$ and effective Majoron-neutrino coupling constant g_{ee} , at 90 % C.L., for ^{136}Xe . The model notation follows the same form as Table 1.3. The third, fourth and fifth columns indicate whether the mode is a Goldstone boson, its leptonic charge L and its spectral index n .

Model	Decay mode	Goldstone boson	L	n	$T_{1/2}^{0\nu \chi^0(\chi^0)}$ (y)	g_{ee}
IB	$0\nu 2\beta \chi^0$	no	0	1	$>2.6 \times 10^{24}$	$<(0.8-1.6) \times 10^{-5}$
IC	$0\nu 2\beta \chi^0$	yes	0	1	$>2.6 \times 10^{24}$	$<(0.8-1.6) \times 10^{-5}$
ID	$0\nu 2\beta \chi^0 \chi^0$	no	0	3	$>4.5 \times 10^{23}$	<0.68
IE	$0\nu 2\beta \chi^0 \chi^0$	yes	0	3	$>4.5 \times 10^{23}$	<0.68
IIB	$0\nu 2\beta \chi^0$	no	-2	1	$>2.6 \times 10^{24}$	$<(0.8-1.6) \times 10^{-5}$
IIC	$0\nu 2\beta \chi^0$	yes	-2	3	$>4.5 \times 10^{23}$	<0.013
IID	$0\nu 2\beta \chi^0 \chi^0$	no	-1	3	$>4.5 \times 10^{23}$	<0.68
IIE	$0\nu 2\beta \chi^0 \chi^0$	yes	-1	7	$>1.1 \times 10^{22}$	<1.2
IIF	$0\nu 2\beta \chi^0$	gauge boson	-2	3	$>4.5 \times 10^{23}$	<0.013
bulk	$0\nu 2\beta \chi^0$	bulk field	0	2	$>1.0 \times 10^{24}$	




Figure 5.1: Reproduced from [49]. Visible energy spectra for the different Majoron-emitting modes of $0\nu 2\beta$ in KamLAND-Zen. The spectra have been convolved using a detector response function, including energy resolution and energy-scale non-linearities. The resolution-limited line at the Q -value represents $0\nu 2\beta$ without Majoron emission.

5.2 Digitising the KamLAND-Zen spectra

For the verification study, we chose to digitise the spectra from KamLAND-Zen’s results [49]—reproduced here in Figure 5.2 and then convert each of them to an echidna spectrum, that we can use to determine limits. The digitisation process was simple, using the Java application EasyNData [88]. The work-flow was:

1. Launch the application and load the highest-resolution version of the image containing the spectra to digitise, available in the online material for [49].
2. Calibrate the image scale by selecting two known points on the image and inputting their coordinate values. The application also has the option to specify if either axis is a logarithmic scale, which is the case for the y -axis in Figure 5.2. The choice of points here is important as it affects the calibration for the full plot. Naturally a good




Figure 5.2: Reproduced from [49]. The data points show the visible energy spectrum of candidate events. The curves show the combined best-fit background spectrum (grey dashed line), $2\nu 2\beta$ spectrum (purple) and 90 % confidence upper limits on $0\nu 2\beta$ and Majoron-emitting modes, for spectral indices $n = 1, 2, 3, 7$. The red line shows the sum of the $2\nu 2\beta$ and best-fit background spectra. The best-fit has $\chi^2/\text{d.o.f} = 100.4/87$.

choice of points are well spaced—giving good coverage of the plot—and easy to measure accurately. Given that one cannot be certain of the starting position on the x -axis (visible energy), we opted for $(1.0, 10^4)$ and $(4.0, 10^{-1})$ as the two calibration points.

3. Once we have calibrated the image, select data points or points along a curve by placing the cursor (cross-hair) over them and clicking to log the coordinates.
4. Once we have logged a full curve or dataset, dump the coordinate pairs to a data-file.

For the “data” points, in Figure 5.2 placing the cross-hair on the centre of the point is trivial, but the smooth background and signal curves required a little more thought. Using a rule on the screen, we made sure that each point selected along the curve, intersected with the energy-coordinate line of the data point—the central value of each 50 keV bin. Thus, when fitting the digitised spectra, using the same binning, we could be sure to compare corresponding bin values.

Figure 5.3 serves as an initial check of the digitisation process. In the figure, we have plotted each of the digitised spectra from Figure 5.2, to demonstrate their agreement with the original spectrum. Note that the peaks of the $n = 7$ mode, $2\nu 2\beta$ and $0\nu 2\beta$ spectra (respectively, $(0.6, 6 \times 10^3)$; $(0.8, 2.5 \times 10^4)$; and $(2.55, 3.5)$) coincide with their counterpart peaks in Figure 5.2. Or, since the remaining Majoron-modes have broader peaks, we can check features such as the configuration of the end-points. At this level, for $n = 3, n = 2$ and $n = 1$ the endpoints are spaced evenly over ~ 100 keV and the $2\nu 2\beta$ endpoint almost aligns with the $n = 2$ spectrum. These checks suggest, at least to first-order, that the digitised spectra are an accurate reproduction of the spectra in Figure 5.2.




Figure 5.3: Spectral plot using the digitised spectra to reproduce Figure 5.2.

5.3 Methodology for verification study

Now that we have digitised the spectra from Figure 5.2, we can use these spectra in echidna for setting limits. We will use a process similar to the one we describe in Section 4.6, using the digitised spectra as simulated Monte Carlo (MC), and fitting to the digitised data spectrum. Since KamLAND-Zen already convolved the spectra with a detector response function, we treat all the spectra except the data points as fully reconstructed visible energy information, and so do not need to apply any detector convolutions to the digitised spectra. We will not float these as nuisance parameters in the fit but we will float the rates of the background spectra in the fit. Floating the rate of $2\nu 2\beta$ is an essential part of the fit because the correlation

(or anti-correlation) between the $2\nu 2\beta$ and signal rates, requires that we fit their contributions simultaneously. Note that we also cannot constrain the rate of the $2\nu 2\beta$, because of the correlation with the signal spectra, and because previous measurements of the $2\nu 2\beta$ for ^{136}Xe , will have assumed zero contribution from Majoron-emitting modes. We also float the combined spectrum of other radioactive backgrounds. Since we do not have information on the individual components, we float the rate of the full spectrum.

All of the processes can be approximated as Poissonian, so we take the uncertainty on the rate of the full spectrum as \sqrt{N} , where N is the total number of decays, the spectrum represents. We assume KamLAND-Zen simulated ten times the number of MC events they expect to observe. It is common, particularly among liquid scintillator experiments like KamLAND-Zen and SNO+, to strive to simulate ten times the exposure of their data. However we consider here the consequence of this assumption. Figure 5.4 shows the visible energy for the KamLAND-Zen dataset of candidate events. The red curve is the sum of the dominant and combined other backgrounds and represents the best fit to the data points, assuming no additional signal contributions. It includes a light red shaded band, which represents the $\pm 1\sigma$ Poisson uncertainty range on the summed spectrum, assuming an MC sample with ten times the exposure of the data spectrum. In the lower panel we show the deviation of the data points from the best fit curve and compare the size of these fluctuations with the $\pm 1\sigma$ error bands, for different MC sample exposure scenarios. The outer light blue band represents the error band for an MC sample with five times the exposure of the data spectrum. It covers 71 of the 84 data-points or $\sim 84\%$. The medium blue band represents 10-times the exposure in the MC sample and covers 58 data-points (69%). The final scenario is a factor of 20 more exposure in the MC—the dark blue band—enclosing 43 points (51%). Since, by definition the 1σ error band on a best-fit curve, should cover approximately 68% of the data-points we fit to, the assumption of a factor of 10 increase in exposure in the MC sample appears to be the most appropriate interpretation for this fit. Nevertheless, if this assumption is incorrect it would affect size of our error margins on the lifetime limits we compare to KamLAND-Zen's results.

5.4 Estimating digitisation uncertainty

Before fitting the digitised spectra, we need to understand the potential additional uncertainty introduced, due to human error when locating the cross-hair on the appropriate point in the spectrum. To estimate this uncertainty, we reload the original plot Figure 5.2, but this time




Figure 5.4: Spectral plot, showing the visible energy for the KamLAND-Zen dataset of candidate events, along with the dominant $2\nu 2\beta$ background contribution (magenta) and the combined contribution from all other radioactive backgrounds (grey). The red curve is the best fit summed spectrum, assuming no additional signal contributions. The light red shaded region shows a $\pm 1\sigma$ Poisson error band. The lower panel plots the ratio of the deviation of the data points from the best fit curve compared to $\pm 1\sigma$ error bands, MC samples with a factors of 5, 10 and 20 times the exposure of the data.

calibrate it assuming a linear scale on the y -axis; giving a y -scale in the range -1.0 to 4.0. From this, we estimate the height of each pixel h_p as 0.005 and that the radii of the data points and half-widths of the curves, is approximately four pixels. Noting that each point read actually requires three measurements—the two calibration points and the data-point itself—and that, conservatively, the two calibration points could equally be out by up to four pixels, the error on each spectral value y_k is:

$$\delta y_k = \pm 4\sqrt{3h_p^2}, \quad \text{where } h_p = 0.005. \quad (5.1)$$

To write this more completely, we are applying a perturbation (P_{lin}) to the spectral values y_k and then subtracting the original y_k values to get the offset:

$$\delta y_k = y_k \pm 4\sqrt{3h_p^2} - y_k, \quad (5.2)$$

but of course the perturbation is uniform across the spectrum, so the y_k values cancel.

Although, the plot in Figure 5.2 does not have a linear y -scale, we can still apply the linear-scale perturbations calculated above, in log-space:

$$\log_{10} P_{lin} = \log_{10} y_k \pm 4\sqrt{3h_p^2}. \quad (5.3)$$

So, taking the exponent of each side, this becomes:

$$P_{lin} = 10^{\log_{10} y_k \pm 4\sqrt{3h_p^2}}, \quad (5.4)$$

$$\Rightarrow y_k 10^{\pm 4\sqrt{3h_p^2}}. \quad (5.5)$$

When we convert this back to an offset, this gives:

$$\delta y_k = y_k 10^{\pm 4\sqrt{3h_p^2}} - y_k, \quad (5.6)$$

as an estimation of the uncertainty introduced through digitising the spectra. Note also, that we only consider the affect of uncertainty along the y -scale, since any perturbation introduced along the x -scale, should be negligible compared to the width of the bins.

To better understand the size of the additional uncertainty introduced during the digitisation process described above, we turn again to a spectral plot showing the dominant and total combined background contributions, and the best-fit to candidate events. In Figure 5.5, we this time show two error bands in the lower ratio plot. We again take the ratio of the

deviation in the data points compared to a 1σ Poisson uncertainty on the MC best-fit point, which the blue band between ± 1 portrays. We then overlay the size of the digitisation error (red band) for comparison. As one would expect from the logarithmic scale of the original plot, the digitisation error dominates over the Poisson error band at lower energies, where the background rate is considerably larger. Whilst at higher energies > 2 MeV, the primary uncertainty is due to the Poissonian nature of the processes. We expect the low energy-region to provide little sensitivity to most of the Majoron-emitting modes of $0\nu 2\beta$, because of the high backgrounds in this region. It is only for the $n = 7$ mode, which (as Figure 5.1 shows) peaks at a lower energy than the $2\nu 2\beta$ contribution, that we are likely to get a significant χ^2 contribution in the energy region where the digitisation error dominates. This should be reflected by a larger error margin on our calculated lifetime for this mode. We will include the error contributions from both sources in the fits in this chapter, summing the digitisation and Poisson uncertainties in quadrature.

5.5 Measurement of the double-beta rate

As an initial step in verifying echidna as a fitting and limit-setting tool, using the χ^2_λ test statistic we calculate the goodness-of-fit for the digitised spectra, to the data spectrum. Figure 5.6 shows the fitted spectra, and residuals from MC, in units of σ underneath. The fit has a value of $\chi^2_\lambda = 112.8$ with 85 degrees of freedom. Note the “dip” in the data spectrum, close to the end-point of the $2\nu 2\beta$ spectrum. KamLAND-Zen observed less events than expected in the energy region where they should get the most sensitivity from the lower spectral index Majoron modes, so we would expect these modes to have a stronger limit than the sensitivity would suggest. We also float the $2\nu 2\beta$ spectrum, without any signal contribution, and compare to the measurement of the $2\nu 2\beta$ lifetime that KamLAND-Zen publish. We measure $T_{1/2}^{2\nu 2\beta} = 2.5 \times 10^{21}$ y with $\chi^2_\lambda = 94.8491$ and 84 degrees of freedom. This is in good agreement with the value KamLAND-Zen published: $T_{1/2}^{2\nu 2\beta} = 2.30 \pm 0.02(stat) \pm 0.12(syst)10^{21}$ y.

5.6 Setting limits on Majoron-emitting $0\nu 2\beta$ modes

For a full fit to the data, to set limits on the signal contributions, we must define the range over which we will float the signals and background spectra. For the $2\nu 2\beta$ rate, we use the measured KamLAND-Zen lifetime $T_{1/2}^{2\nu 2\beta} = 2.30 \pm 0.02(stat) \pm 0.12(syst) \times 10^{21}$ y [49] to calculate an expected rate for this contribution. We then determine an expected 1σ error on

this rate, combining uncertainty due to the Poisson nature of the process and the digitisation uncertainty, as discussed above. We calculate a similar 1σ error on the rate of the combined background contribution.. We then define the range of values over which the $2\nu 2\beta$ rate can float, as $N \pm 10\sigma$. We define the range of values for the total radioactive backgrounds in a similar way, using the range $N \pm 4\sigma$. We also constrain this rate with a penalty term using the expression from Equation (4.48) (Equation (4.48)), where σ is our combined Poisson and digitisation uncertainty and the our prior is the combined background best-fit contribution, when we fit to the data, along with the $2\nu 2\beta$ contribution, but with no additional signal. Note that we cannot apply a penalty term to constrain the $2\nu 2\beta$ rate because our the prior estimate of this rate assumes no additional contribution from Majoron mode signals, but we now fit simultaneously with such a signal contribution.

We float each signal spectrum based on the KamLAND-Zen limit we are trying to reproduce. We define a range from zero counts to twice the counts corresponding to the KamLAND-Zen limit divided over 101 evenly-spaced bins, for all four signal spectra. For the $n = 1$ mode this corresponds to floating a total signal contribution from 0.0 to 90.8 counts and for the $n = 2$ mode, we vary the contribution from 0.0 to 236 counts. The $n = 3$ mode sees its contribution vary from 0.0 to 1.05×10^3 counts and the $n = 7$ spectrum is allowed to contribute from 0.0 to 1.07×10^4 counts. To set limits on the Majoron-emitting modes of $0\nu 2\beta$ we create a grid with 251 values on the $2\nu 2\beta$ axis and 101 values on the radioactive backgrounds rate axis. Then we use echidna's `GridSearch` minimiser to simultaneously fit each signal contribution with the dominant $2\nu 2\beta$ contribution and sub-dominant combined background spectrum to determine a 90 % confidence lower limit on $T_{1/2}^{0\nu 2\beta \chi^0(\chi^0)}$ for the Majoron-emitting mode.

To understand if the fit behaved as expected we consult some diagnostic plots. The Poisson Likelihood χ^2 (χ_λ^2) curves in Figure 5.7 show the value of the χ_λ^2 as a function of number of signal decays. From these plots that we extract the final 90 % confidence limit on the signal rate—marked by the dashed line on the plot. As described in Section 4.6, we take this limit as the first signal contribution that pushes the χ_λ^2 over a value of 2.71 units of χ^2 . In the plots in Figure 5.7, we are also looking to see nice smooth curves, with the x -axis centred around the limit. This is the mark of a successful fit with sufficient granularity that we can be sure it did not settle on a local minimum.

Since we ran the fit using a grid search minimiser, we can also inspect the full χ_λ^2 surface, to examine the fit for correlation (or anti-correlations) between parameters. Figure 5.8 shows the number of $2\nu 2\beta$ decays verses the number of signal decays, χ_λ^2 surfaces, for all four spectral index modes. We clearly see that as we increase the signal contribution, the fit compensates with fewer $2\nu 2\beta$ decays. This anti-correlation is most significant for the signal

mode with spectral index $n = 7$ because the peak in this signal spectrum occurs in a similar energy region to the peak in the $2\nu 2\beta$ spectrum (see Figure 5.1). An important point to note with the contours in Figure 5.8, is their smoothness and the fact they are nicely contained within the parameter space. If a contour reaches the bounding axes, this suggests we are not sampling sufficient parameter space in the fit, and we run the risk of artificially constraining a parameter value.

Figure 5.9 displays a further set of contour plots. This time we see the projection of the number of decays in the combined radioactive backgrounds axis and the signal decays axis. Again, as a check on the quality of the fit, we observe smooth contours that are well contained by the parameter space. However, we see no correlation or anti-correlation between in projection of the χ^2_λ surface. This suggests that floating the number of decays in the combined background spectrum does little to affect the 90 % confidence limit on the number of signal decays.

5.7 Comparison of limits

To compare the 90 % confidence lower limits on $T_{1/2}^{0\nu 2\beta \chi^0(\chi^0)}$, that we calculate here, with their counterpart limits published by KamLAND-Zen (see Table 5.1, or [49]), we take the ratio of $T_{1/2}/T_{1/2, KLZ}$, where $T_{1/2}$ are the limits we calculate here and $T_{1/2, KLZ}$ are the KamLAND-Zen limits. We plot this ratio in Figure 5.10. The purple lines in Figure 5.10 denote the value of this ratio, compared to the dashed line at one, which are the KamLAND-Zen limits. The purple shaded band represents our best estimate of the uncertainty on the limits we calculate here. To construct this error band we run the full limit-setting process twice more. For the lower bound, we assume the two background spectra always fluctuate to their lowest 1σ value, by subtracting the 1σ errors from the spectra on each iteration of the fit, before calculating the χ^2_λ . Then we assume the signal contribution always fluctuates to its maximum value permitted by the 1σ statistical uncertainty. This combination guarantees the maximum signal contribution possible in the fit and so the shortest lifetime limit. For the upper bound we do the reverse, allowing all background contributions to fluctuate to the largest value permitted by the statistical uncertainty on the spectra and the smallest signal contribution possible. Finally, on Figure 5.10, the red lines mark the corresponding expected sensitivity, if we were to fit, using echidna, to the fixed, summed $2\nu 2\beta$ and combined other backgrounds spectra, instead of fitting to the data points.

In Figure 5.10, as per our expectation, the expected sensitivities for the lower spectral index modes of $0\nu 2\beta$ by Majoron-emission show a weaker limit than both our calculated limits and the KamLAND-Zen limits. We remarked previously that the spectrum of candidate data points fluctuates low close to the $2\nu 2\beta$ end-point, where we get the largest contributions to the χ^2 , which causes this disparity between sensitivity estimates and the calculated values. The $n = 1$ mode for $0\nu 2\beta$ by Majoron-emission shows reasonable agreement between our calculated value and the KamLAND-Zen limit and the $n = 2$ mode shows excellent agreement. However, as the spectral index increases, we see more tension between our calculated value and KamLAND-Zen's published result; particularly for the $n = 7$ mode. One potential explanation for this tension is the result of the lack of correlation or anti-correlation we saw in any of the contours in Figure 5.9. Floating the combined background spectrum should have some impact on the limit we set, but the long tail where there are no signal contributions (at energy >2.6 MeV in Figure 5.3), helps to constrain this spectrum during the fit. If we were able to decouple the individual components of the combined background spectrum, we may see correlations or anti-correlations develop, particularly in the higher spectral index modes. But, this would require a more detailed background model than we have available from the digitised spectra.

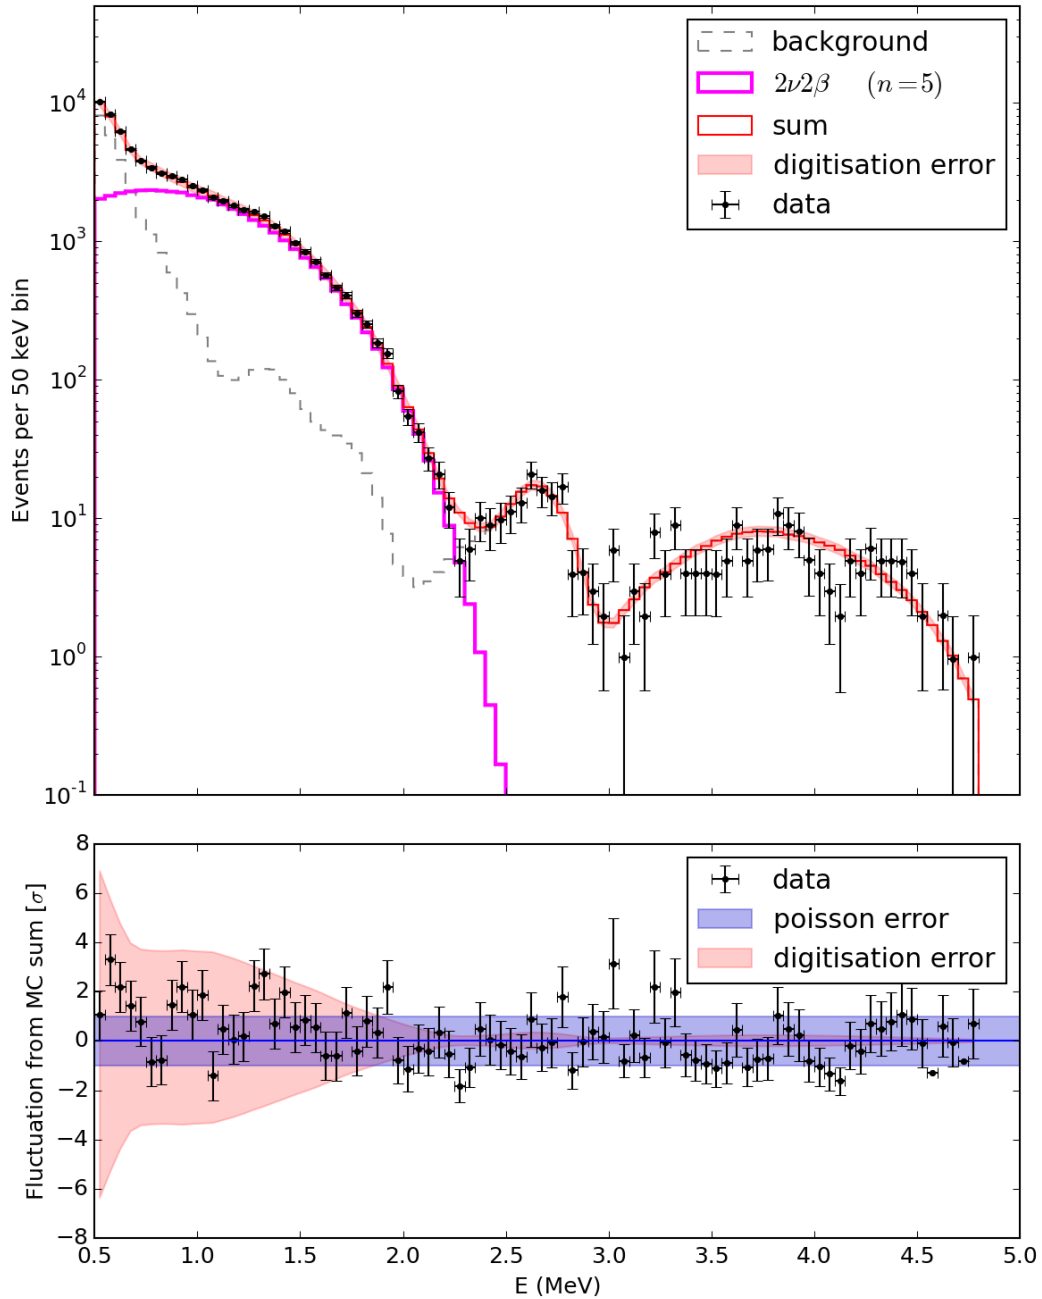


Figure 5.5: Spectral plot, showing the visible energy for the KamLAND-Zen dataset of candidate events, along with the dominant $2\nu 2\beta$ background contribution (magenta) and the combined contribution from all other radioactive backgrounds (grey). The red curve is the best fit summed spectrum, assuming no additional signal contributions. The light red shaded region shows the estimated size of the additional uncertainty in this spectrum, introduced through digitisation. The lower panel plots the ratio of the deviation of the data points from the best fit curve compared to $\pm 1\sigma$ Poisson error band (blue). The red band on the lower panel represents the size of the digitisation uncertainty in relation to the Poisson error band.




Figure 5.6: Spectral plot showing the fixed best-fit of the digitised spectra to the digitised data array. Underneath we include the residuals from MC, in units of σ .




Figure 5.7: Quartet of chi squared (χ^2) curves for the KamLAND-Zen fit, with floating $2\nu 2\beta$ and floating combined background, showing χ^2_{λ} as a function of signal decays. With (*top-left* to *bottom-right*) spectral indices $n = 1, 2, 3, 7$. The dashed lines mark the position of the 90% confidence upper limit on the rate.

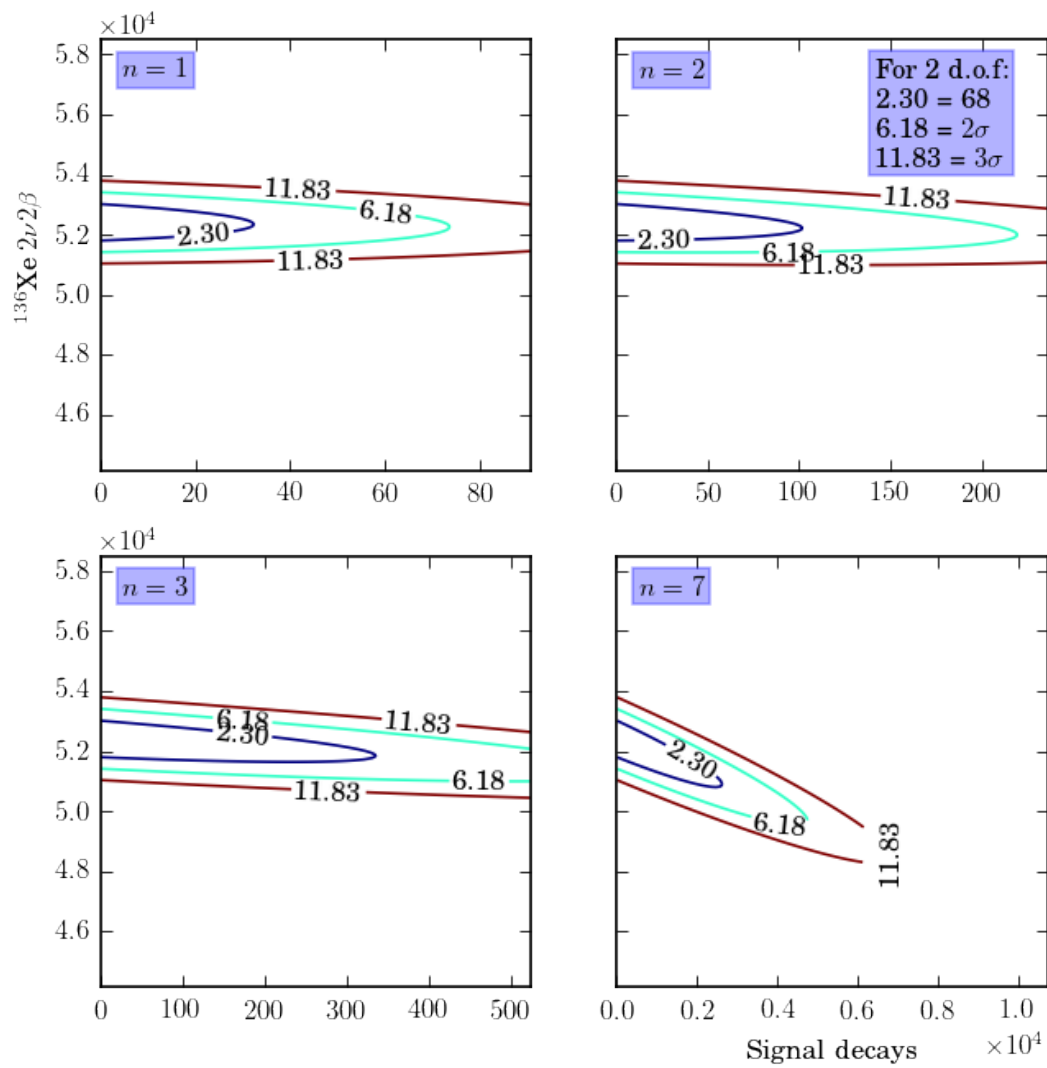


Figure 5.8: Quartet of χ^2 contour plots, showing the projection of the χ^2_λ onto the $2\nu 2\beta$ rate and signal scale axes, with (*top-left to bottom-right*) spectral indices $n = 1, 2, 3, 7$.

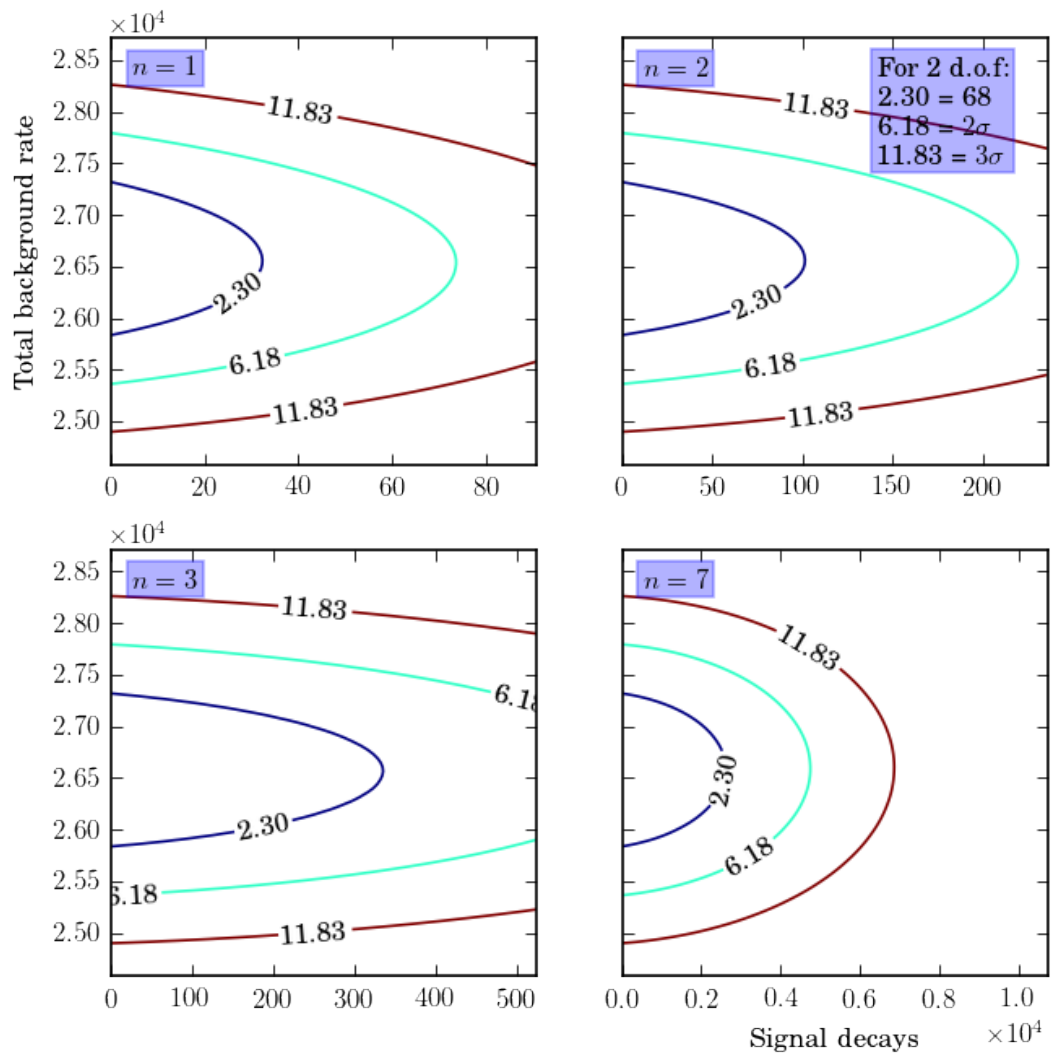


Figure 5.9: Quartet of χ^2 contour plots, showing the projection of the χ^2_λ onto the combined background rate and signal scale axes, with (*top-left to bottom-right*) spectral indices $n = 1, 2, 3, 7$.




Figure 5.10: Comparison of echidna's limits for the digitised spectra, with KamLAND-Zen's limits. The dashed line at 1 signifies the KamLAND-Zen, limit. The purple lines show the ratio of echidna's limit divided by the corresponding KamLAND-Zen limit, with and shaded purple band representing the statistical uncertainty on the calculated limit. The red lines represent the corresponding sensitivity estimate for each mode, calculated using echidna

Chapter 6

Estimating SNO+ sensitivity to Majoron-emitting modes

6.1 General assumptions

In the limit setting analyses this chapter presents, we make the following general assumptions. All Monte Carlo (MC) used is from the rat-5.0.2 production, which nominally assumes a light-yield of 200 NHit/MeV; a trigger sum of 200 signals from individual photomultiplier tubes (PMTs) corresponds to ~ 1 MeV visible energy. For all modes of ^{130}Te double-beta decay, we assume 0.3% loading natural tellurium, by mass, when calculating the rate. As discussed in Section 2.3, the default loading option for the main neutrinoless double-beta decay ($0\nu 2\beta$) sensitivity analysis is 0.5%. The change in the default loading choice was coupled to the choice of the composition of the scintillator cocktail, as we described in Section 2.3 when discussing [60]. The most significant impact the loading and scintillator cocktail choice has on this analysis, is the expected light yield, which then affects the energy resolution we model using echidna's Poisson convolution method. Where the previous default loading choice and scintillator cocktail, gave a light yield of 200 NHit/MeV, the new default choice of 0.5% loading $^{\text{Nat}}\text{Te}$, using the Te-diol method almost doubles the expected light-yield—350 NHit/MeV is the conservative estimate from [60]. This means a much improved energy resolution, but we did not adopt this loading choice in the analysis presented in this chapter, because the decision to switch to 0.5% loading using the Te-diol method, occurred after we had pre-convolved the rat-5.0.2 MC, assuming a light-yield of 200 NHit/MeV.

Following the convention of the main SNO+ $0\nu 2\beta$ sensitivity analysis [54], we assume 5 y livetime for all sensitivity studies and a self-shielding Fiducial Volume (FV) cut at 3.5 m

radius. This has the effect of reducing the active volume, and therefore the number of ^{130}Te nuclei (assuming they are homogeneously distributed), to around 20 % the volume enclosed by the Acrylic Vessel (AV). In all fits, we use MC true energy convolved using echidna, not reconstructed energy, to generate the spectra, so we assume 100 % detection efficiency. The actual detection efficiency of SNO+ will be less, meaning that the analyses presented here overestimate the number of decays in all spectra. Yet, we have propagated this assumption, throughout all the signal and background spectra and in the total background spectra we fit to as “fake data”, so the effects of this assumption on individual spectra, should cancel when we run the fits.

6.2 Backgrounds

As we saw in Section 1.4, a critical consideration for any experiment probing neutrino mass, is the level of background. To make an observation of $0\nu 2\beta$, SNO+ must reduce contributions to the energy spectrum, or at least develop a full understanding of the energy spectrum we expect in the absence of a signal. In a limit setting analysis, this background model is the array \mathbf{y} , defined in the formalism in Section 4.1, and consists of an expected count in each bin of the visible energy spectrum, or in a given region of interest (ROI). This section aims to describe how we construct this array of counts, to develop a background model we can use in a fit.

6.2.1 ^{130}Te double-beta decay

The primary background for a $0\nu 2\beta$ search with SNO+, is the irreducible background from the two-neutrino mode. In a classic $0\nu 2\beta$ experiment, the finite energy resolution of the detector broadens the delta function at the Q -value that is the signal, and extends the tail of the continuous two-neutrino double-beta decay ($2\nu 2\beta$) spectrum into the region where one searches for a signal. We have discussed the assumptions relating to global systematic uncertainties, such as the detector energy resolution, but the number of $2\nu 2\beta$ decays that fall in the ROI is also a key uncertainty in the fit.

We calculate the activity A of $2\nu 2\beta$ (y^{-1}), using the expression:

$$A = \frac{\log 2}{T_{1/2}^{0\nu 2\beta}} \frac{N_A}{M_r} f(^{130}\text{Te}) L m_{LAB}, \quad (6.1)$$

where we use the lifetime limit of $T_{1/2}^{2\nu2\beta} = 7.0 \times 10^{20}$ y [59]. From [89], the mass of linear alkylbenzene (LAB) (m_{LAB}) is 780 t. We multiply by the loading fraction (by mass) of natural tellurium ($L = 0.003$), to get the mass of ^{Nat}Te , which evaluates to 2.3 t. The factor $f(^{130}\text{Te})$, is the mass fraction of ^{130}Te in natural tellurium. Using a relative natural abundance of 34.08 %, we calculate its value as 0.3470. Multiplying by this gives us ~ 806 kg as the mass of ^{130}Te in the AV. Next we need the number of ^{130}Te nuclei. We take the value of Avogadro's number N_A from [65] and values for the atomic mass of tellurium (M_r). Multiplying the mass of ^{130}Te by the factor N_A/M_r , gives us $\sim 3.8 \times 10^{27}$ nuclei in the AV. The final part of Equation (6.1) comes from the standard decay rate relation:

$$A = -\frac{dN}{dt} = \frac{\log 2}{T_{1/2}}N, \quad (6.2)$$

where N is the number of nuclei and $T_{1/2}$ is a generic decay half-life. Adding this factor into Equation (6.1) and the value for $T_{1/2}^{2\nu2\beta}$, yields an activity of 3.7×10^6 y^{-1} , for the full AV and across the full energy region. If we make an FV cut, this adds a further factor of r_{FV}^3/r_{AV}^3 , which is ~ 0.2 for a 3.5 m radius FV cut. In Figure 6.1, the red curve shows the expected $2\nu2\beta$ spectrum, after the FV cut, when scaled to the expected number of decays we calculate here. We also note, in Table 6.2, that $2\nu2\beta$ ranks as the second background, by counts in the ROI, for Majoron studies.

The continuous nature of the Majoron signal spectra, means that the level of $2\nu2\beta$ decays, is an even more important parameter in fits that probe these signals. Whilst, in standard $0\nu2\beta$ searches, one may use the precision of ex-situ measurements of the $2\nu2\beta$ rate, for the same isotope, to constrain this rate parameter in the fit, we must fit simultaneously, the $2\nu2\beta$ and the potential contribution from a Majoron-emitting mode. As shown by the fits to KamLAND-Zen data, in Chapter 5, the $2\nu2\beta$ rate is highly anti-correlated with a possible Majoron contribution—most notably for modes with a higher spectral index. Previous measurements of the lifetime for $2\nu2\beta$ in ^{130}Te , such as the result we use to calculate the expected number of decays above, do not assume a possible contribution from Majoron-emitting modes, so we can use this number to estimate the expected total number of decays in a fixed fit and use it to guide the range over which we allow the $2\nu2\beta$ rate to float, in a full fit, but we cannot use it to constrain the rate of $2\nu2\beta$ in the full fit.

6.2.2 Solar neutrinos

The solar neutrino spectrum from decay of ${}^8\text{B}$, in the Standard Solar Model (SSM), is another irreducible background for any $0\nu 2\beta$ search, in SNO+. We take the yearly number of solar neutrino interactions, within the scintillator volume and across the range 0 to 15 MeV, as 1270 y^{-1} [90]. This assumes three-flavour mixing and that we only detect electron neutrinos. This value is from June 2013. Whilst, [90] (version 27) presents another value 1021 y^{-1} and later versions of the document suggest 1338 y^{-1} , due to an updated flux. The most recent number relies on more recent production MC. As a compromise, we chose to use the initial calculation, which is the median of the three numbers and retains the assumptions of rat-5.0.2 MC, for consistency.. The flat light-green curve in Figure 6.1, illustrates the expected solar neutrino spectrum, after the FV cut and Table 6.2 provides the number of counts with the range 1.0 to 3.5 MeV. In contrast to the main $0\nu 2\beta$ analysis for SNO+, ${}^8\text{B}$ solar neutrinos are not the dominant background contribution for this analysis, due to the fact we have a much wider energy ROI.

6.2.3 Internal backgrounds

${}^{238}\text{U}$

Decays from ${}^{238}\text{U}$ and other isotopes in this decay chain, occur naturally in most materials. Emissions from this radioactive decay chain will originate from all components in the scintillator cocktail, including the telluric acid used for tellurium loading. We aim to minimise this contribution through a rigorous underground purification of the scintillator cocktail—as Section 2.3 described. Since the chain is in equilibrium, we scale all its constituent parts to the same expected rate (except for the AV leaching contributions, due to ${}^{210}\text{Pb}$, ${}^{210}\text{Bi}$ and ${}^{210}\text{Po}$, which we discuss separately). We base the rate on a target level, equal to the radiopurity of Borexino, during Phase I [91]. One determines the contributions from individual components of the scintillator cocktail and then sums these to get a total rate.

A key concern in this chain is the ${}^{214}\text{Bi} \rightarrow {}^{214}\text{Po}$ transition, with a branching ratio of 99 %, via either a 3.27 MeV β or a 2.2 MeV γ [90]. The ${}^{214}\text{Po}$, is short-lived, with a lifetime of 164 μs , so one can tag these events by the coincident delayed 7.83 MeV α . There are two types of coincidence, one out of window, where the initial decay and delayed α must be in separate triggered events, whilst the in-window tag considers events that pile-up in the same trigger window. The delayed coincidence tag rejects 99 % ${}^{214}\text{Bi}$ decays in the ROI for $0\nu 2\beta$, via the

standard mechanism. At present it is not clear how this rejection cut applies to an extended ROI, so we took the conservative decision not to apply this cut. The in-window cut rejects 98 % events in the standard $0\nu 2\beta$ ROI. We chose to apply this cut.

²³²Th

The ²³²Th decay chain, like ²³⁸U, also occurs naturally in most materials. We base the rate on the same target radiopurity levels that we use for the uranium chain and, since the chain is also in equilibrium, scale each individual component to the same rate. This chain also has a delayed coincidence transition. A 2.25 MeV β emission, is the product of the ²¹²Bi \rightarrow ²¹²Po transition, that occurs in 64 % of decays. The delayed coincidence tag, rejects 99.999 % of the bismuth decays, whilst the in-window tag also rejects about 98 % of the single trigger events. As with the uranium chain, we chose to apply the in-window rejection factor, but not the delayed coincidence suppression.

Other internal backgrounds

We see some contributions from other internal backgrounds. Most of these are at low energies. We consider contributions from ³⁹Ar, ¹⁴C, ⁸⁵Kr and ⁴⁰K. We scale each internal background to the rate given by [90]. For ⁴⁰K, we include contributions from both the hold-up and hold-down ropes and both surfaces of the AV, as well as the acrylic itself. The activity is $2.00 \times 10^4 \text{ y}^{-1}$. The activity of ³⁹Ar is $9.00 \times 10^4 \text{ y}^{-1}$ and for ¹⁴C it is $3.90 \times 10^5 \text{ y}^{-1}$. From [90], we get the ⁸⁵Kr activity as $6.00 \times 10^4 \text{ y}^{-1}$.

6.2.4 AV and leaching backgrounds

In the final section of the uranium chain, ²¹⁰Pb is a long-lived isotope with a lifetime of 22.6 y. The remaining decays to ²¹⁰Bi and ²¹⁰Po become a problem, if the levels of lead build up. This can happen because of leaching from the surface of the AV. One can calculate the expected rates of these backgrounds on a year-by-year basis, from modelling the expected leaching rate of ²¹⁰Pb into the AV. We use the rates from [90] and sum over the expected rates for years one through to five. Because we set everything as a per year rate, and then multiply by the lifetime within the limit setting code, We divide the summation by five, to get a mean yearly rate (which is then multiplied by five in the limit setting script). This means the rates

we calculate are strictly only valid for a five-year livetime, but that is all we consider in this analysis.

6.2.5 α, n backgrounds

One of the main α, n background sources stems from energetic α emission, interacting with oxygen and carbon nuclei. The primary source of these energetic α particles is the ^{210}Po from above, and is therefore intrinsically linked to the leaching rate. One calculates the rates for oxygen and carbon interactions separately. Table 6.1 summarises the calculation of the rates for these backgrounds. The α, n interactions also produce a delayed 2223 keV $-\gamma$, which one can use in a coincidence tag of these events. Tagging like this reduces the rate of the prompt decay by a factor of ~ 250 . Each α, n background produces a delayed γ so we calculate the rate of these separately, by summing all α, n contributions from both oxygen and carbon. The rejection factor for the delayed γ is about a factor of 10.

Table 6.1: Summary of the expected background contributions, over 5 y, from the α, n backgrounds and corresponding delayed 2223 keV $-\gamma$ [90].

Source	Scaling	Decays	Decays (with cuts)
AV surface internal ^{13}C	0.004	855	3.44
AV surface internal ^{18}O	0.004	318	1.28
AV surface scintillator ^{13}C	0.004	1349	5.42
AV surface scintillator ^{18}O	0.004	12	0.048
AV surface external ^{13}C	0.004	2190	8.80
AV surface external ^{18}O	0.004	675	2.71
Scintillator ^{13}C	0.004	6080	24.5
Scintillator ^{18}O	0.004	56	0.225
2223 keV γ	0.097	11500	1120

6.2.6 External backgrounds

We include external background contributions from the internal calibration ropes, external ropes—both hold-up and hold-down ropes—the water shielding and the PMTs and PMT

support structure (PSUP). We also consider contributions from three sources on the AV, the acrylic and both internal and external dust. We scale all external backgrounds to the expected rates from [90] and apply a factor of 0.5 analysis cut, based on PMT hit times, to all the rates. Note, for externals we must use reconstructed MC, to get a true representation of the rates, as the FV cut yields no MC truth external events. This is not a problem because we do not wish to float any contributions from external sources, so we may use the reconstructed events and keep the rates fixed.

Table 6.2: Summary of the number of counts in the ROI 1.0 to 3.5 MeV, for the main (grouped) background contributions.

Background	ROI counts
^{210}Bi	2.395e+07
$2\nu 2\beta$	1.439e+06
^{238}U chain	2.352e+05
^{210}Po	1.117e+05
^{232}Th chain	2.432e+04
Scintillator	1453
^8B ν ES	507.8
2223 keV $-\gamma$	220.5
External	58.91
(α, n)	2.895
Total	2.577e+07

6.3 Energy systematics

In addition to setting the correct rate from Section 6.2, whether we fix the rate or float it in the fit, we also need to convolve the spectra correctly, to account for the detector response. We use all three of echidna's convolution functions, described in Section 4.3.5. For fixed backgrounds we convolve to the nominal resolution of 200 NHit/MeV and maintain default values for the energy scale factor and energy shift, of 1.0 and 0.0. For backgrounds, whose rate we are floating in the fit, there is the potential to float all three of these systematics in the fit as well.




Figure 6.1: Spectral plot, showing contributions to the fixed background spectrum, in the energy range 1.0 to 3.5 MeV.

6.4 A fixed background fit

In this section we describe a limit setting analysis using a fit to the combined fixed background spectrum. This analysis follows the procedure outlined by Sections 4.5 to 4.6. The fixed background spectrum from Figure 6.1 acts as our observed data spectrum. We then fit the fixed total background spectrum to the data spectrum, with an additional contribution from each signal in turn, using `echidna`'s `Limit` class, to set a 90% confidence lower limit on the expected lifetime of that signal. We use the default χ^2_λ (Baker-Cousins chi squared (χ^2)) [84] as the test statistic in the fit. In this first fit, we test an extended energy range 0 to 3.5 MeV, so that we can determine an appropriate ROI for further fits, based on these results.

To determine an appropriate range for the signal scalings, we defined an initial range from zero decays, up to the number of $2\nu 2\beta$ decays (3.7243×10^6), but then refined this number to improve the granularity of the fit and determine an accurate limit. For each limit

setting analysis, we produce a quartet of χ^2 curves—see Figure 6.2¹. These are essential for diagnosing any problems in the fit. The curve should be smooth and parabolic to indicate no local minima, any jumps could suggest that the granularity of the search is too coarse. The limit should be close to the centre of the range in signal decays. If it is too low, one again risks a search that is too coarsely binned to determine an accurate limit and if it is too high, one cannot guarantee that the smooth χ^2 curve continues well after reaching the limit. The χ^2 curves in Figure 6.2, however, all look smooth and nicely centred, suggesting an accurate limit for each of these fits. We summarise the 90 % confidence limits obtained, for this fit, in Table 6.3.

Table 6.3: Summary of the estimated 90 % confidence lower limits on the half life ($T_{1/2}^{0\nu 2\beta\chi^0(x^0)}$), for the fixed background fit.

Spectral index	$T_{1/2}^{0\nu 2\beta\chi^0(x^0)}$ (y)
1	5.431×10^{24}
2	1.795×10^{24}
3	8.604×10^{23}
7	1.516×10^{23}

We note the improvement, when comparing the estimated $T_{1/2}^{0\nu 2\beta\chi^0(x^0)}$ 90 % confidence lower limits, from the fixed background fit, with the corresponding limits set by KamLAND-Zen (see Table 5.1). Although not strictly a fair comparison, since we are comparing limits across two different isotopes, it is nevertheless interesting to note an improvement over KamLAND-Zen by at least a factor of two, in the $n = 1$ mode, increasing to almost an order of magnitude improvement in the $n = 7$ mode. For the fixed-background fit, we expect an improvement on this scale because we have not floated any parameters in the fit, yielding an idealised result. As we introduce more complex fits in this chapter will see our estimates of the half life limits, come more in line with those from KamLAND-Zen. We can draw further comparisons with the SNO+ sensitivity estimate in the main $0\nu 2\beta$ analysis, where we estimate $T_{1/2}^{0\nu 2\beta} \geq 1.96 \times 10^{26}$ y. Our studies of the Majoron-emitting modes of $0\nu 2\beta$, clearly set a much lower half life limit than we do for the standard mode of $0\nu 2\beta$. This is predominantly a consequence of the fact the Majoron mode signals span a much wider energy range than the ROI the main analysis uses. As such, we see much more of the low-energy backgrounds and the inherent $2\nu 2\beta$ background, in the fit, leading to a poorer half life estimate.

¹Note, we will maintain the order of the signals, on all quartets of plot, unless stated otherwise.




Figure 6.2: Quartet of χ^2 curves for the fixed background fit, showing Poisson Likelihood χ^2 (χ^2_{λ}) as a function of signal decays, with (top-left to bottom-right) $n = 1, 2, 3, 7$. The dashed lines mark the position of the 90% confidence limit.

In Figures 6.3 to 6.6, we have produced a spectral plot for each of the fits detailed above. They show the total fixed background in grey and in blue, each Majoron-emitting mode of $0\nu 2\beta$, scaled to the estimated $T_{1/2}^{0\nu 2\beta \chi^0}(\chi^0)$ limit. Underneath the main spectral plot, we have displayed the χ^2 contribution from each bin in the fit. An interesting observation is the fact that for modes $n = 1, 2$ we continue to get the majority of our sensitivity, as in the main $0\nu 2\beta$ analysis, from the energy region close to the end point of the $2\nu 2\beta$ spectrum. The reason for this is due to the distinct low background region you can see in Figure 6.1 as the $2\nu 2\beta$ spectrum falls away. But where the spectral index shifts the peak of the Majoron-mode energy distribution further towards, and past in the case of $n = 7$, the peak in the $2\nu 2\beta$




Figure 6.3: Spectral plot showing the fixed background spectrum and $n = 1$ signal spectrum, at the 90% confidence limit. Underneath the histogram shows the contribution from each bin to the full sensitivity.

spectrum, we get progressively more sensitivity at lower energies. For the $n = 7$ mode, the largest contributions to our sensitivity come from an energy region dominated by completely different backgrounds than at the $2\nu 2\beta$ end-point. Here the dominant backgrounds are the two leaching backgrounds ^{210}Bi and ^{210}Po , as well as the $2\nu 2\beta$ spectrum.

Figures 6.3 to 6.6 are also useful in determining an appropriate ROI for the remainder of the fits. We note that no spectra have much sensitivity below 1.0 MeV, so we will use this as the lower bound of our energy ROI. For the upper bound, we wish to keep some of the background spectrum beyond the endpoint of the signals, so that we can use it to self-constrain background rates. For this reason, we chose to keep an upper bound of 3.5 MeV. We will use this ROI for all the remaining fits in this chapter.

6.5 Floating the double-beta rate

The first step to improve on the limits we obtained in the fixed background fit, is to float the rate of $2\nu 2\beta$ in ^{130}Te . As we saw for ^{136}Xe , in Chapter 5, there is high anti-correlation between the $2\nu 2\beta$ background and the Majoron-emitting modes, particularly for higher




Figure 6.4: Spectral plot showing the fixed background spectrum and $n = 2$ signal spectrum, at the 90 % confidence limit. Underneath the histogram shows the contribution from each bin to the full sensitivity.




Figure 6.5: Spectral plot showing the fixed background spectrum and $n = 3$ signal spectrum, at the 90 % confidence limit. Underneath the histogram shows the contribution from each bin to the full sensitivity.




Figure 6.6: Spectral plot showing the fixed background spectrum and $n = 7$ signal spectrum, at the 90 % confidence limit. Underneath the histogram shows the contribution from each bin to the full sensitivity.

spectral indices. We expect to get a less stringent limit, compared to those in Table 6.3, by allowing the $2\nu 2\beta$ rate to float.

Now that we are floating parameters in the fit, we will have to use a minimiser. We opt for echidna’s default GridSearch minimiser, to ensure we can accurately determine the position of the best-fit and not locate a local minimum. Also, using GridSearch permits us to scrutinise the χ^2 surface after the fit and look at parameter correlations (or anti-correlations). To float the $2\nu 2\beta$ rate in the fit, we must create a configuration for the fit, which defines the boundaries of the grid on this axis and the number of cells in the grid. Although we cannot constrain the rate using the previous measurement of $T_{1/2}^{2\nu 2\beta} = 7 \pm 0.9(\text{stat}) \pm 1.1(\text{syst}) \times 10^{20}$ y from NEMO-3 [59] that we used to set its rate in the fixed background fit, we can use this value to guide our choice of grid. The lifetime is equivalent to 18.62×10^6 decays. Since the process is Poissonian, we take the \sqrt{N} , where N is the total number of decays in the spectrum, as an estimate for the uncertainty in the rate. Thus $\sigma = 4315$ decays. Then we define the range as $N \pm n\sigma$, where n is the number of sigma we wish to deviate from our prior rate (the NEMO-3 measurement), in the fit. For this analysis, we define a grid with the range $N \pm 12\sigma$, along the $2\nu 2\beta$ rate axis.

For the signals, we set a high of 5000 counts for the $n = 1$ mode. We use 251 bins for all modes to achieve a good granularity. For $n = 2$, we float the rate in the range 0 to 9310 counts, for $n = 3$ we define the range 0 to 5.4×10^4 counts and for $n = 7$ we float from 0 to 6.0×10^5 counts. Table 6.4 summarises the limits of the fit. We have also included a spectral plot for the fit, in Figure 6.7, that portrays the signal spectra, scaled to the 90% confidence lower limits, from Table 6.4.




Figure 6.7: Spectral plot, showing the four signals scaled to their limits, and the floated $2\nu 2\beta$ spectrum at its best fit value.

As we did for the fixed background fit, the first step in reviewing these limits is to confirm the status of the fit. The χ^2_λ curves in Figure 6.8, from which we extract the 90% confidence limits in Table 6.4, form nice smooth curves, with x -axis centred on the limit. These suggest successful fits with sufficient granularity that did not settle in a local minimum. Since we ran the fit as a grid search, we can also inspect the full χ^2_λ surface, to examine the fit for correlation (or anti-correlation) between parameters. Figure 6.9 shows the number of $2\nu 2\beta$ decays versus the number of signal decays, χ^2_λ surfaces, for all four spectral index modes. The best-fit positions in each case take the number of $2\nu 2\beta$ decays as 18.62×10^6 , the number we used to guide the parameter range when setting up the grid. But, we see clearly that as we increase the signal contribution, the fit prefers fewer $2\nu 2\beta$ decays, to compensate. This

anti-correlation exactly mirrors the behaviour we saw when reproducing the results from KamLAND-Zen in Chapter 5. An important point to note with these contours, is that they are smooth and also nicely contained within the parameter space. If we see a contour that reaches the boundary of the parameter space, this suggests that the fit is artificially constraining a parameter by not sampling enough of the parameter space.

Note that the 90 % confidence lower limits, on $T_{1/2}^{0\nu 2\beta\chi^0}(\chi^0)$, that we present in Table 6.4 are weaker than the limits we got from the fixed background fit. In this fit, the limit for the $n = 1$ mode is slightly less stringent than its counterpart from Table 6.3 but as we increase the spectral index, we see a greater reduction in the predicted lifetime limit, in comparison to Table 6.3. For the $n = 7$ mode, the limit is now less than a third of the value we determined in the fixed background fit. This is expected, and is a direct consequence of the interplay between the size of the $2\nu 2\beta$ and signal contributions, that the anti-correlation in Figure 6.9 reveals.

Table 6.4: Summary of 90 % confidence limits for the fit floating the $2\nu 2\beta$ rate.

Spectral index	$T_{1/2}^{0\nu 2\beta\chi^0}(\chi^0)$ (y)
1	5.299×10^{24}
2	1.522×10^{24}
3	4.752×10^{23}
7	4.642×10^{22}

6.6 Floating the detector energy resolution

Our next step in improving the limits, is to float the detector energy resolution, as a nuisance parameter in the fit. Detector energy resolution is a global fit parameter that one should apply to both the signal spectra and any backgrounds we are floating. We will continue to float the $2\nu 2\beta$ rate and apply the energy resolution parameter, on top of this—adding an additional dimension to the fit.

We specify the values for the global fit parameter by creating a `GlobalFitConfig` in `echidna`. This ensures that `echidna` applies the global parameters uniformly, across all the spectra we are floating in the fit. The `GlobalFitConfig` also specifies the values over which we wish to




Figure 6.8: Quartet of χ^2 curves for the fixed background fit, with floating $2\nu 2\beta$ rate, showing χ^2_{λ} as a function of signal decays. With (top-left to bottom-right) $n = 1, 2, 3, 7$. The dashed lines mark the position of the 90% confidence limit.

float the global fit parameter. Accurate convolution of spectra, to simulate detector energy resolution, remains a computer intensive process, in echidna, so we opted to pre-convolve the spectra and let echidna load them, when needed, to use in the fit. This makes the fit considerably faster to run, but has the disadvantage that the values used to pre-convolve the spectra fix the size of this dimension in the grid. As with defining the $2\nu 2\beta$ grid, we used the prior energy resolution—the nominal value of 200 NHit/MeV that accompanies the rat-5.0.2 MC—to determine an appropriate range of values to test for the energy resolution. We create a dimension in the fit, with 41 bins, in the range 180 to 220 NHit/MeV. Table 6.5 summarises the parameter values for the signal contribution. For the number of $2\nu 2\beta$ decays, we used

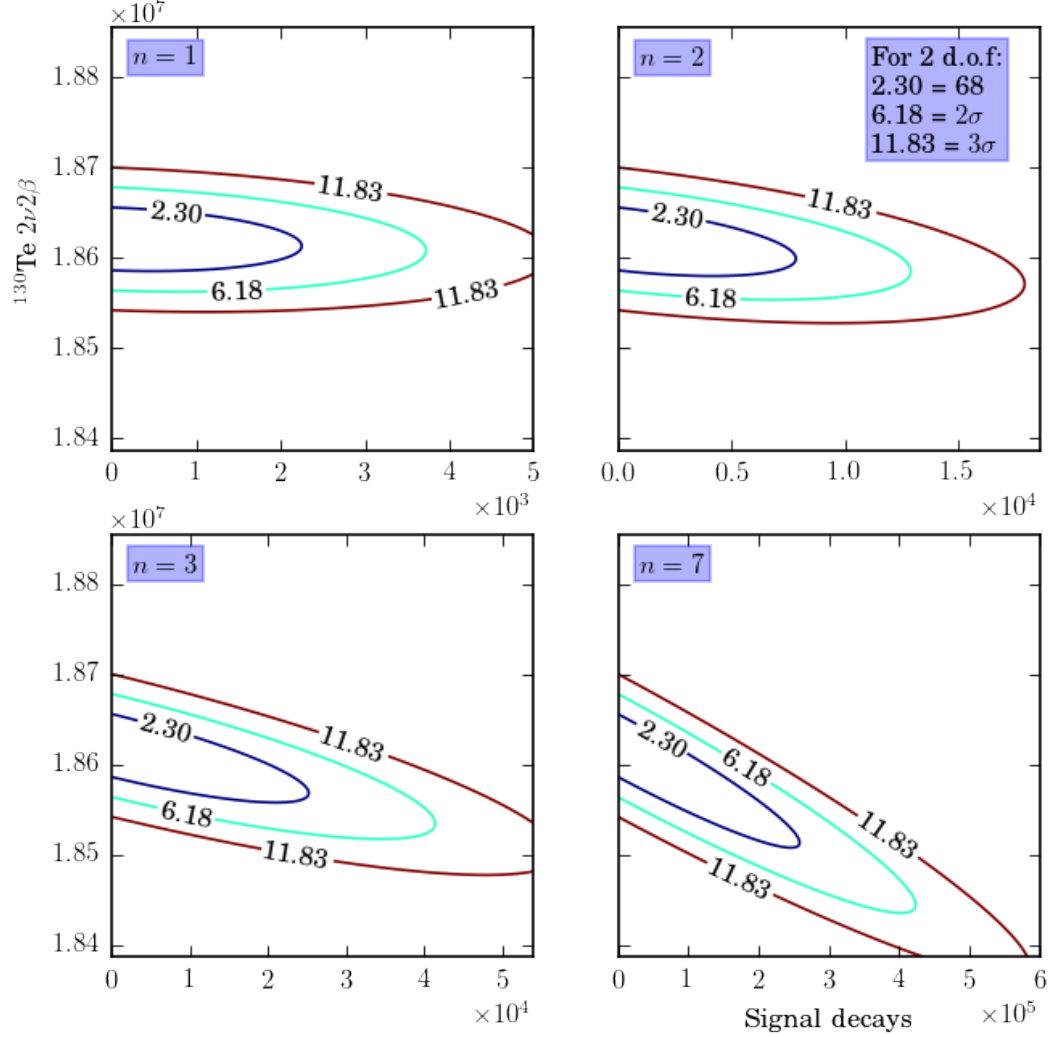


Figure 6.9: Quartet of χ^2 contour plots, showing the projection of the χ^2_λ onto the $2\nu 2\beta$ rate and signal scale axes, with (top-left to bottom-right) $n = 1, 2, 3, 7$.

different ranges for different signals in this fit. For $n = 1, 2$ we defined the rate for $2\nu 2\beta$ as $N \pm 8\sigma$, whilst for $n = 3$ we used $N \pm 12\sigma$ and we specified $N \pm 15\sigma$ for $n = 7$.

As with the previous two fits, we produce an array of different χ^2 plots, to help diagnose if the fit ran successfully. Figure 6.11 shows the χ^2_λ surface projected onto the $2\nu 2\beta$ decays and signal decays, parameter axes. As we saw from the previous fit, the χ^2 contours suggest an anti-correlation between the rate of $2\nu 2\beta$ and the signal scale. This is most noticeable for the higher spectral index modes.

Table 6.5: Summary of the signal parameter values used in the fit.

Parameter	Minimum	Maximum	Bins
$n = 1$ rate	0	9388	101
$n = 2$ rate	0	5.396×10^4	101
$n = 3$ rate	0	1.297×10^4	101
$n = 4$ rate	0	3.902×10^4	101

We also consult the χ_λ^2 surface, projected onto the light yield vs signal decays, parameter axes, see Figure 6.12. For this parameter combination, we see a correlation between the signal scale and light yield value for the spectral index modes $n = 1, 2, 3$ but we see a slight anti-correlation between these parameters in the $n = 7$ mode. This means that for modes $n = 1, 2, 3$, the fit prefers a better energy resolution as the signal contribution increases, whilst the $n = 7$ mode prefers a less resolved signal spectrum, as the number of decays increases. A possible explanation of this behaviour is that, as we saw in Figures 6.3 to 6.5, most of our sensitivity to the lower spectral index modes comes from the bins around the $2\nu 2\beta$ end-point. The poorer the energy resolution, the more events will spill out, past the $2\nu 2\beta$ end-point, into the low background region, which would be highly disfavoured by the fit and yield a higher χ_λ^2 value. In contrast the spectrum for the $n = 7$ mode falls off before the $2\nu 2\beta$ end-point, and so a more poorly resolved spectrum here, would tend to more closely resemble the $2\nu 2\beta$ spectrum. Since we still have the interplay between the signal and $2\nu 2\beta$ rates, the fit can compensate by lowering the number of $2\nu 2\beta$ decays and get a good fit.

Another key function of this array of plots, is to confirm that the parameter space we have created is large enough, such that it does not artificially constrain any of the parameters. A clear indication that the parameter space is too small, would be if the half-ellipse shape of the contours were significantly cut short by the axes of the plot. This would suggest we were not sampling enough of the parameter space. Figure 6.11, confirms that we have allowed sufficient parameter space for the rate of $2\nu 2\beta$ to float unconstrained, however this is not the case for Figure 6.12. We can see that for all three of the modes $n = 1, 2, 3$, the contours “hit” the edge of the parameter space. This is reflected in the χ_λ^2 curves in Figure 6.13, which no longer show a smooth χ^2 curve for the $n = 2, 3$ modes, instead we see a sharp increase in gradient at just over 2×10^4 signal counts in $n = 2$ and at just past 0.6×10^5 signal counts in $n = 3$. These would correspond to the points where the best-fit path through the χ_λ^2 surface reaches the edge of the surface in the $2\nu 2\beta$ versus signal decays plane. The result is an artificially lower limit in the number of signal decays, as the signal rate and $2\nu 2\beta$ rate

then must compensate for unintentional constraint on the energy resolution parameter. This results in estimating a slightly stronger lower limit than we should, on $T_{1/2}^{0\nu 2\beta \chi^0(x^0)}$ for $n = 2, 3$. Since we opted to load the pre-convolved spectra, we cannot increase the parameter space without pre-convolving more spectra and there was insufficient time during the preparation of this thesis to prepare the additional spectra.

Table 6.6 summarises the 90 % confidence lower limits on the lifetime of the signals that we achieve with this fit. This time we see the greatest worsening in the limit, from the lower spectral index modes—the $n = 1$ limit decreasing by almost a factor of two and the limits decreasing by around a factor of three for $n = 2, 3$ —when compared to the fit where we only floated the $2\nu 2\beta$ rate. This is a consequence of the direct correlation between the value of the light yield parameter and the size of the signal contribution. In contrast, we see a modest decrease in the limit for the $n = 7$ mode due to the slight anti-correlation between the light yield value and signal contribution. The results are also displayed visually in the spectral plot in Figure 6.16. Each of the signal spectra are scaled to represent the 90 % confidence lower limit on $T_{1/2}^{0\nu 2\beta \chi^0(x^0)}$ that we present in Table 6.6 and the $2\nu 2\beta$ spectrum is scaled to its best fit rate. We plot the spectra convolved with the global best-fit light yield value and corresponding energy resolution.

Table 6.6: Summary of 90 % confidence limits for the fit floating the $2\nu 2\beta$ rate.

Spectral index	$T_{1/2}^{0\nu 2\beta \chi^0(x^0)}$ (y)
1	2.777×10^{24}
2	4.645×10^{23}
3	1.478×10^{23}
4	3.630×10^{22}

6.7 Floating other systematics

For the sake of completeness, we will note the other detector convolutions that echidna can perform. We investigated floating shifts in energy and energy scale factors, in the fit. Following the same prescription as for floating the light yield, we apply each of these parameters, in turn, on top of the $2\nu 2\beta$ rate. We show the arrays of contour plots in Figure 6.14 and Figure 6.15, noting the lack of correlation between these nuisance parameters and the signal rates. This is




Figure 6.10: Spectral plot showing the four fitted Majoron spectra, scaled to the 90 % confidence upper limit, that we set. Also shown is ^{130}Te $2\nu 2\beta$ background, that we float in the fit. We scale this to its best-fit rate. We also show all floating spectra convolved to their best-fit energy resolution.

in direct contrast to the clear anti-correlation we saw for the energy resolution, and suggests that floating these parameters will not have a significant effect on the limit. The resulting contours are also fairly jagged in both Figure 6.14 and Figure 6.15, which suggests that our parameter scan does not have sufficient granularity to determine the preferred values of these parameters. Like the energy resolution, we also pre-convolved the spectra to use in this fit, using echidna, so without pre-convolving additional spectra, we are limited to the granularity shown here. We also uncovered an imperfection in echidna that resulted in having to load these spectra from file, whenever the fit required, even when not pre-convolving the spectra. The input/output requirements of this fix, significantly harm the performance of the fit and made it nonviable to include these parameters in the fit, at this stage.




Figure 6.11: Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the $2\nu 2\beta$ rate and signal scale axes. With (top-left to bottom-right) $n = 1, 2, 3, 7$.

6.8 Discussion

We present the results of the fit from Section 6.6, where we float the rate of $2\nu 2\beta$ in ^{130}Te and the detector energy resolution, as our best estimate of SNO+ sensitivity to $0\nu 2\beta$ via Majoron-emitting modes. Figure 6.16 illustrates the fixed background spectrum, floated $2\nu 2\beta$ spectrum and the spectrum corresponding to each of the four Majoron-mode signals, scaled to the 90% confidence lower limit we set in the fit. Table 6.8 summarises the best-fit values of the parameters we float in the fit.

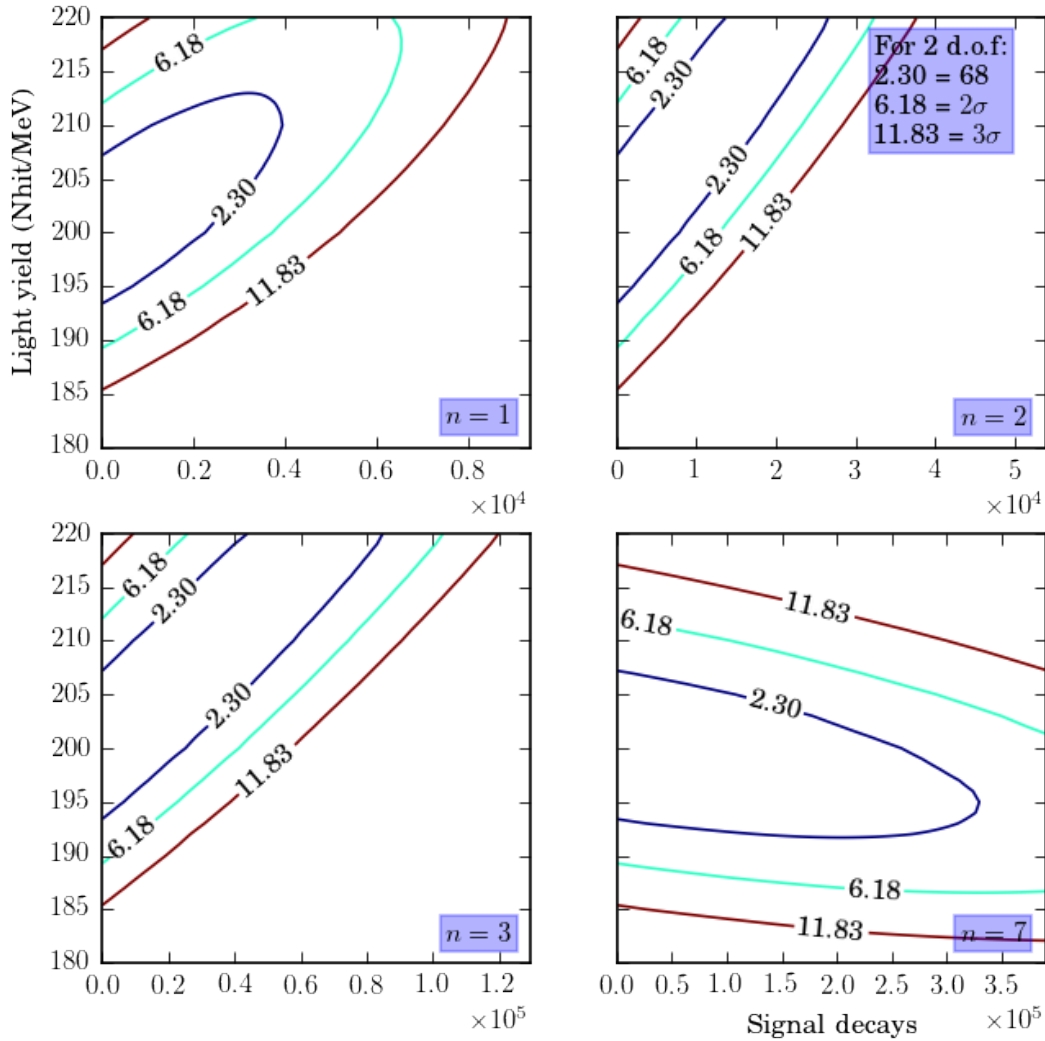


Figure 6.12: Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the energy resolution and signal scale axes. With (top-left to bottom-right) $n = 1, 2, 3, 7$.

We can use the values of the effective coupling g_{ee} , that Table 6.7 presents, as an isotope-independent tool for comparisons between experiments. We note that our sensitivity estimates for SNO+, compare favourably with the 90% confidence upper limits, set by KamLAND-Zen. KamLAND-Zen has set the most stringent limits, by probing Majoron-emitting modes in ^{136}Xe . Comparing to the limits set by KamLAND-Zen, we see that SNO+ can expect to do slightly better, in all but one of the Majoron-emitting modes. We fall short of the constraint on g_{ee} , set by KamLAND-Zen in all multi-majoron modes, with a spectral index of three. SNO+ could, nevertheless, offer an improved limit in three other modes that conserve lepton number, not including the disfavoured $n = 1$ mode, which are key channels for new physics. For ^{130}Te , the current best limit on g_{ee} , for $n = 1$ Majoron emitting modes only, comes from NEMO-3 [59].




Figure 6.13: Quartet of χ^2 curves for the fixed background fit, with floating $2\nu 2\beta$ rate and energy resolution, showing χ^2_λ as a function of signal decays. With (*top-left to bottom-right*) $n = 1, 2, 3, 7$. The dashed lines mark the position of the 90% confidence upper limit.

Our expected limit on the effective coupling, improves on the NEMO-3 limit by two orders of magnitude.

The caveat that accompanies these sensitivity estimates, whilst we have floated the two most important systematic uncertainties in the fit, floating further parameters will likely reduce the limits on $T_{1/2}^{0\nu 2\beta}$, resulting in a less constraining value for g_{ee} . One such parameter is the rate of ^{210}Bi . Figure 6.1, suggests a high anti-correlation between ^{210}Bi and the $n = 7$ Majoron-emitting mode. Equally one would expect the 2223 keV $-\gamma$ background to be correlated with the $n = 1, 2, 3$ Majoron-modes. It would be interesting to develop the fit




Figure 6.14: Quartet of χ^2 contour plots, showing the projection of the χ^2_{λ} onto the energy scale and signal scale axes. With (*top-left to bottom-right*) $n = 1, 2, 3, 7$.

presented here to incorporate these background rates as floating parameters. One could constrain the rates of these backgrounds, using the leaching rates. One could float other background rates, such as those for the uranium and thorium chain backgrounds, in the fit. As the chains are in equilibrium, one would have to couple together the rates of all backgrounds in the chain. We expect these to have a minimal effect on the limit, compared to the rate of $2\nu 2\beta$ or even the rate of ^{210}Bi , because we can use the part of the spectrum beyond the end points of all the signals, to self-constrain these background rates.

In the fit we present here, we have included the detector energy resolution as a parameter in the fit. We have also shown that the other two energy scale systematics—the scaling factor and a shift in energy—do not contribute significantly to the fit because they are not correlated

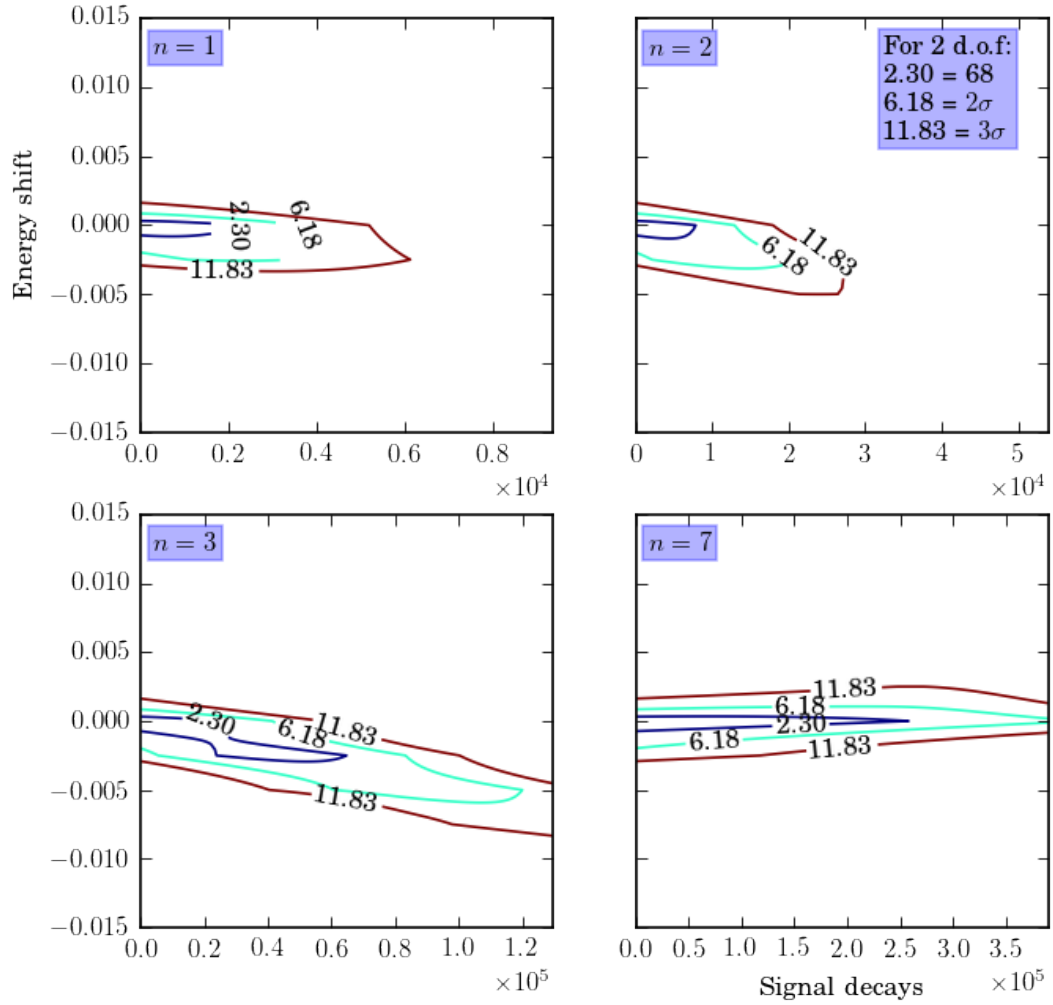


Figure 6.15: Quartet of χ^2 contour plots, showing the projection of the χ_λ^2 onto the energy shift and signal scale axes. With (top-left to bottom-right) $n = 1, 2, 3, 7$.

with the number of signal decays. The fit neglects to include any further energy systematics, including non-linear effects such as quenching.

A final point worth noting, is that all signals we consider in this analysis, assume 0.3% by mass, loading of $^{\text{Nat}}\text{Te}$, where as the default loading for SNO+, with the innovative Te-diol loading approach, is 0.5%. Unfortunately, one cannot scale the rates for 0.3% loading to the 0.5% level, because the optical properties of the new scintillator cocktail, differ noticeably. This results in, among other effects, a different detector resolution. An increased loading of $^{\text{Nat}}\text{Te}$ would undoubtedly shift limits on $T_{1/2}^{0\nu 2\beta}$ and the effective coupling in a favourable direction.

Table 6.7: Summary of limits.

	Decay mode	Goldstone boson	L	n	$T_{1/2}^{0\nu\chi^0(\chi^0)}$ (y)	$g_{ee\ min}$	$g_{ee\ max}$
Model							
IB	$0\nu 2\beta\chi^0$	no	0	1	2.8×10^{24}	5.2×10^{-6}	1.2×10^{-5}
IC	$0\nu 2\beta\chi^0$	yes	0	1	2.8×10^{24}	5.2×10^{-6}	1.2×10^{-5}
ID	$0\nu 2\beta\chi^0\chi^0$	no	0	3	1.5×10^{23}	0.89	0.89
IE	$0\nu 2\beta\chi^0\chi^0$	yes	0	3	1.5×10^{23}	0.89	0.89
IIB	$0\nu 2\beta\chi^0$	no	-2	1	2.8×10^{24}	5.2×10^{-6}	1.2×10^{-5}
IIC	$0\nu 2\beta\chi^0$	yes	-2	3	1.5×10^{23}	0.0097	0.0097
IID	$0\nu 2\beta\chi^0\chi^0$	no	-1	3	1.5×10^{23}	0.89	0.89
IIE	$0\nu 2\beta\chi^0\chi^0$	yes	-1	7	3.6×10^{22}	0.87	0.87
IIF	$0\nu 2\beta\chi^0$	gauge boson	-2	3	1.5×10^{23}	0.0097	0.0097
bulk	$0\nu 2\beta\chi^0$	bulk field	0	2	4.6×10^{23}		

Table 6.8: Table summarising the best fit values of the systematic uncertainties.

n	$^{130}\text{Te } 2\nu 2\beta$ rate (y)	Light yield (NHit/MeV)
1	$7_{-0.016}^{+0.012} \times 10^{20}$	$200_{-7}^{+12.5}$
2	$7_{-0.016}^{+0.044} \times 10^{20}$	200_{-7}^{+20}
3	$7_{-0.016}^{+0.059} \times 10^{20}$	200_{-7}^{+20}
7	$7_{-0.016}^{+0.055} \times 10^{20}$	$200_{-8}^{+7.5}$




Figure 6.16: Spectral plot showing the four fitted Majoron spectra, scaled to the 90% confidence upper limit, that we set. Also shown is ^{130}Te $2\nu 2\beta$ background, that we float in the fit. We scale this to its best-fit rate. We also show all floating spectra convolved to their best-fit energy resolution.

Chapter 7

Conclusions

In this thesis we have motivated the search for $0\nu 2\beta$, as a means of determining the mass of the neutrino. The SNO and KamLAND experiments, have proved that transitions between neutrino flavour states occur, which means that there must be at least two massive neutrino states. This remains the single experimental conclusion that the Standard Model (SM) cannot explain. The solution to this problem, requires either an extension of the SM, by introducing a right-handed term for the neutrinos, or the development of a Majorana mass term for the neutrino. We then introduce the $0\nu 2\beta$, as a probe of neutrino mass. The Schechter-Valle theory, formulates $0\nu 2\beta$ as an effective operator, that induces a transition from ν_e to $\bar{\nu}_e$. The effective operator theory means that an observation of $0\nu 2\beta$, regardless of the underlying physics mechanism, would induce a non-zero amplitude for the Majorana neutrino mass. This suggests that one could exploit any mechanism, for the underlying physics process, to probe a Majorana neutrino mass for the neutrino. $0\nu 2\beta$ via the emission of one or two additional scalars, called Majorons, is a potential model of the underlying physics mechanism, for $0\nu 2\beta$.

Models for these decays are plentiful and offering much scope for investigation by experiments. One of the key advantages of this signal, is the continuous form of the visible energy spectrum—as opposed to the resolution limited delta function, that is the signal for $0\nu 2\beta$ via the standard mass mechanism. The broad signal allows one to search for the process over a wider energy range. One can categorise the observed spectra according to their spectral index. The higher the spectral index of a model, the closer to the lower the peak in the visible energy range. The background spectra of $0\nu 2\beta$ experiments, tends to vary across the visible energy region. This means that different $0\nu 2\beta$ experiments may be more sensitive to different spectral indices. We have briefly reviewed the experimental status of both standard searches for $0\nu 2\beta$ and searches probing Majoron-emitting modes.

SNO+ is one experiment, aiming to probe $0\nu 2\beta$ via the standard mechanism, involving the exchange of a light Majorana neutrino. We have presented an overview of the detector and also reviewed the varied physics program of SNO+. Whilst $0\nu 2\beta$ is the main goal of the experiment, we hope to build on the success of Sudbury Neutrino Observatory (SNO) in studying the solar neutrino spectrum. SNO+ is aiming to make the first precision measurement of the flux from the CNO cycle as well as probing the transition region of the ^8B spectrum. With a sufficiently low ^{14}C background, SNO+ could make a measurement of the p p flux. In the initial *water phase*, SNO+ will investigate invisible nucleon decay, in ^{16}O nuclei, via the mechanism $n \rightarrow 3\nu$. In both *loaded* and *unloaded* scintillator phases, SNO+ can study reactor antineutrinos, from reactors at over three different baselines. In the main $0\nu 2\beta$ analysis, SNO+ will use a novel loading technique to load 0.5% $^{\text{Nat}}\text{Te}$, by mass. SNO+ estimates a 90% confidence lower limit on the lifetime of $0\nu 2\beta$, as $T_{1/2}^{0\nu 2\beta} = 1.96 \times 10^{26}$ y. We have presented a detailed account of the techniques SNO+ employs for detection and data acquisition (DAQ). The SNO+ trigger system is vital for ensuring we save the events we need to analyse.

Equally important for ensuring we have high quality data to analyse, is the Data Quality (DQ) code. In SNO+ high-level DQ (HLDQ) ensures we have the highest quality data available to analyse, without it, SNO+ will not be able to achieve its physics goals. We have presented a detailed account of the development and testing of a suite of processors that perform the HLDQ checks. The processors include checks of basic run properties, such as ensuring a run duration over 30 minutes, to determining the detector coverage afforded by usable PMTs. We have described key algorithms employed when performing the checks and the criteria that we use to qualify each check result. The HLDQ checks we have described, act as the last part in a group of analyses referred to as the front-end calibrations processes. These are the front line in the analysis of any data from SNO+ and we run them, in some capacity, both nearline, to deliver almost realtime feedback to the detector operator, and offline, in the main Worldwide LHC Computing Grid (WLCG) processing.

We have described an limit setting software framework (echidna), which we have developed predominantly at Queen Mary, University of London and the University of Sussex, and of which the author was a primary developer. The framework provides complete functionality for performing fitting and limit setting analyses involving the fitting of spectral shapes. The first stage in the work-flow involves the creation of a spectrum and filling it with events from SNO+ MC. One can then manipulate the spectra, by applying convolutions, to emulate detector reconstruction effects, and by scaling the spectra to a specific rate or shrinking to an ROI. We centred the fitting aspect around the creation of a χ^2_λ surface. By default we determine the global minimum of this surface using a grid searching method, but we also

include other, more optimised methods for finding the minimum. We then build on the fitting framework by introducing a spectrum that we have specified as the signal. By gradually increasing the rate of the signal, and determining the $\Delta\chi^2_{\lambda}$, at each step, we can determine the contribution that corresponds to a 90 % confidence lower limit, on the signal rate.

We tested echidna using both a toy model and by attempting to reproduce the results of Majoron-mode search, conducted by KamLAND-Zen. We were able to reach a rough agreement between the results echidna produces and the published KamLAND-Zen results, but because we do not have information on the components of the spectrum of combined radioactive backgrounds, we were unable to account for all the correlation in, the higher spectral index modes.

The final result of this thesis is study of the potential sensitivity of SNO+ to Majoron-emitting modes of $0\nu 2\beta$, in ^{130}Te . We ran a fit including detector energy resolution and the $2\nu 2\beta$ rate, as floating parameters. We determine the effective coupling for each of the main Majoron-emitting modes. With five live years of data, we expect SNO+ to set a more stringent limit than the current best limits from KamLAND-Zen, in all but the multi-Majoron modes with a spectral index of three. We expect to set a world leading limit on the effective coupling, for ^{130}Te . Most interesting about these results, is that SNO+ has high sensitivity to the higher spectral index modes that conserve lepton number. These modes hold interesting prospects for the search for new physics.

Bibliography

- [1] W. Pauli, “Dear radioactive ladies and gentlemen,” *Phys. Today*, vol. 31N9, p. 27, 1978.
- [2] H. Becquerel, “Emission de Radiations Nouvelles par l’Uranium Metallique / Emission of New Radiation by Metallic Uranium,” *Compt. Ren.*, vol. 122, p. 1086, 1896.
- [3] E. Rutherford, “Uranium Radiation and the Electrical Conduction Produced by it,” *Phil. Mag.*, vol. 47, p. 109, 1899.
- [4] J. Chadwick, “Intensitätsverteilung im magnetischen Spektrum der β -Strahlen von Radium B + C / The Intensity Distribution in Magnetic Spectrum of β -Rays of Radium B + C,” *Verhandl. Dtsch. phys. Ges.*, vol. 16, p. 383, 1914.
- [5] J. Chadwick, “Possible Existence of a Neutron,” *Nature*, vol. 129, p. 312, 1932.
- [6] E. Fermi *Ricerca Scientifica*, vol. 2, p. 12, 1933.
- [7] F. Perrin *Comptes Rendus*, vol. 197, p. 1625, 1933.
- [8] F. Reines, “The neutrino: From poltergeist to particle,” *Rev. Mod. Phys.*, vol. 68, pp. 317–327, 1996.
- [9] F. Reines and C. L. Cowan, “The neutrino,” *Nature*, vol. 178, pp. 446–449, 1956.
- [10] R. Davis, Jr., “Attempt to detect the antineutrinos from a nuclear reactor by the $^{37}\text{Cl}(\bar{\nu}, e^{-})^{37}\text{A}$ reaction,” *Phys. Rev.*, vol. 97, pp. 766–769, 1955.
- [11] E. J. Konopinski and H. M. Mahmoud, “The Universal Fermi interaction,” *Phys. Rev.*, vol. 92, pp. 1045–1049, 1953.
- [12] G. Danby, J. M. Gaillard, K. A. Goulianos, L. M. Lederman, N. B. Mistry, M. Schwartz, and J. Steinberger, “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos,” *Phys. Rev. Lett.*, vol. 9, pp. 36–44, 1962.

- [13] K. Kodama *et al.*, “Observation of tau neutrino interactions,” *Phys. Lett. B*, vol. 504, pp. 218–224, 2001, hep-ex/0012035.
- [14] S. L. Glashow, “Partial Symmetries of Weak Interactions,” *Nucl. Phys.*, vol. 22, pp. 579–588, 1961.
- [15] S. Weinberg, “A Model of Leptons,” *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 1967.
- [16] A. Salam, “Weak and Electromagnetic Interactions,” *Conf. Proc. C*, vol. 680519, pp. 367–377, 1968.
- [17] S. Bilenky, “Neutrino oscillations: From a historical perspective to the present status,” *Nucl. Phys. B*, vol. 908, pp. 2–13, 2016, 1602.00170.
- [18] J. N. Bahcall, N. A. Bahcall, and G. Shaviv, “Present status of the theoretical predictions for the Cl-36 solar neutrino experiment,” *Phys. Rev. Lett.*, vol. 20, pp. 1209–1212, 1968.
- [19] R. Davis, Jr., D. S. Harmer, and K. C. Hoffman, “Search for neutrinos from the sun,” *Phys. Rev. Lett.*, vol. 20, pp. 1205–1209, 1968.
- [20] B. T. Cleveland, T. Daily, R. Davis, Jr., J. R. Distel, K. Lande, C. K. Lee, P. S. Wildenhain, and J. Ullman, “Measurement of the solar electron neutrino flux with the Homestake chlorine detector,” *Astrophys. J.*, vol. 496, pp. 505–526, 1998.
- [21] T. Kajita, E. Kearns, and M. Shiozawa, “Establishing atmospheric neutrino oscillations with Super-Kamiokande,” *Nucl. Phys.*, vol. B908, pp. 14–29, 2016.
- [22] B. Pontecorvo, “Neutrino Experiments and the Problem of Conservation of Leptonic Charge,” *Sov. Phys. JETP*, vol. 26, pp. 984–988, 1968. [Zh. Eksp. Teor. Fiz.53,1717(1967)].
- [23] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*. 2007.
- [24] Q. R. Ahmad *et al.*, “Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory,” *Phys. Rev. Lett.*, vol. 89, p. 011301, 2002, nucl-ex/0204008.
- [25] J. N. Bahcall, “Solar models and solar neutrinos: Current status,” *Phys. Scripta*, vol. T121, pp. 46–50, 2005, hep-ph/0412068.
- [26] Y. Fukuda *et al.*, “Evidence for oscillation of atmospheric neutrinos,” *Phys. Rev. Lett.*, vol. 81, pp. 1562–1567, 1998, hep-ex/9807003.

- [27] P. A. R. Ade *et al.*, “Planck 2015 results. XIII. Cosmological parameters,” *Astron. Astrophys.*, vol. 594, p. A13, 2016, 1502.01589.
- [28] J. J. Gomez-Cadenas, J. Martin-Albo, M. Mezzetto, F. Monrabal, and M. Sorel, “The Search for neutrinoless double beta decay,” *Riv. Nuovo Cim.*, vol. 35, pp. 29–98, 2012, 1109.5515.
- [29] M. G. Inghram and J. H. Reynolds, “Double Beta-Decay of Te^{130} ,” *Phys. Rev.*, vol. 78, pp. 822–823, Jun 1950.
- [30] W. H. Furry, “On transition probabilities in double beta-disintegration,” *Phys. Rev.*, vol. 56, pp. 1184–1193, 1939.
- [31] J. Schechter and J. W. F. Valle, “Neutrinoless Double beta Decay in $\text{SU}(2) \times \text{U}(1)$ Theories,” *Phys. Rev. D*, vol. 25, p. 2951, June 1982.
- [32] S. M. Bilenky and C. Giunti, “Neutrinoless Double-Beta Decay: a Probe of Physics Beyond the Standard Model,” *Int. J. Mod. Phys.*, vol. A30, no. 04n05, p. 1530001, 2015, 1411.4791.
- [33] J. Barea, J. Kotila, and F. Iachello, “ $0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements in the interacting boson model with isospin restoration,” *Phys. Rev.*, vol. C91, no. 3, p. 034304, 2015, 1506.08530.
- [34] F. Šimkovic, V. Rodin, A. Faessler, and P. Vogel, “ $0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements, quasiparticle random-phase approximation, and isospin symmetry restoration,” *Phys. Rev.*, vol. 87, no. 4, p. 045501, 2013, 1302.1509.
- [35] J. Menendez, A. Poves, E. Caurier, and F. Nowacki, “Disassembling the Nuclear Matrix Elements of the Neutrinoless beta beta Decay,” *Nucl. Phys.*, vol. A818, pp. 139–151, 2009, 0801.3760.
- [36] J. Kotila and F. Iachello, “Phase-space factors for double- β decay,” *Phys. Rev. C*, vol. 85, pp. 034316–, Mar. 2012.
- [37] M. Agostini *et al.*, “Results on Neutrinoless Double- β Decay of ^{76}Ge from Phase I of the GERDA Experiment,” *Phys. Rev. Lett.*, vol. 111, no. 12, p. 122503, 2013, 1307.4720.
- [38] R. Arnold *et al.*, “First results of the search of neutrinoless double beta decay with the NEMO 3 detector,” *Phys. Rev. Lett.*, vol. 95, p. 182302, 2005, hep-ex/0507083.
- [39] R. Arnold *et al.*, “Search for neutrinoless double-beta decay of ^{100}Mo with the NEMO-3 detector,” *Phys. Rev.*, vol. D89, no. 11, p. 111101, 2014, 1311.5695.

- [40] C. Arnaboldi *et al.*, “A Calorimetric search on double beta decay of Te-130,” *Phys. Lett.*, vol. B557, pp. 167–175, 2003, hep-ex/0211071.
- [41] E. Andreotti *et al.*, “ ^{130}Te Neutrinoless Double-Beta Decay with CUORICINO,” *Astropart. Phys.*, vol. 34, pp. 822–831, 2011, 1012.3266.
- [42] J. B. Albert *et al.*, “Search for Majorana neutrinos with the first two years of EXO-200 data,” *Nature*, vol. 510, pp. 229–234, 2014, 1402.6956.
- [43] A. Gando *et al.*, “Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen,” *Phys. Rev. Lett.*, vol. 117, no. 8, p. 082503, 2016, 1605.02889. [Addendum: *Phys. Rev. Lett.* 117, no. 10, 109903 (2016)].
- [44] J. Argyriades *et al.*, “Measurement of the Double Beta Decay Half-life of Nd-150 and Search for Neutrinoless Decay Modes with the NEMO-3 Detector,” *Phys. Rev. C*, vol. 80, p. 032501, 2009, 0810.0248.
- [45] G. B. Gelmini and M. Roncadelli, “Left-handed neutrino mass scale and spontaneously broken lepton number,” *Physics Letters B*, vol. 99, no. 5, pp. 411–415, 1981.
- [46] S. Schael *et al.*, “Precision electroweak measurements on the Z resonance,” *Phys. Rept.*, vol. 427, pp. 257–454, 2006, hep-ex/0509008.
- [47] C. P. Burgess and J. M. Cline, “New class of Majoron-emitting double- β decays,” *Physical Review D*, vol. 49, no. 11, pp. 5925–5944, 1994.
- [48] P. Bamert, C. P. Burgess, and R. N. Mohapatra, “Multi-majoron modes for neutrinoless double-beta decay,” *Nuclear Physics B*, vol. 449, pp. 25–48, 1995.
- [49] A. Gando *et al.*, “Limits on Majoron-emitting double-beta decays of Xe-136 in the KamLAND-Zen experiment,” *Phys. Rev. C*, vol. 86, no. 2, p. 021601, 2012, 1205.6372.
- [50] M. Hirsch, H. V. Klapdor-Kleingrothaus, S. G. Kovalenko, and H. Päs, “On the observability of majoron emitting double beta decays,” vol. 372, no. 1–2, pp. 8–14, 1996.
- [51] J. Suhonen and O. Civitarese, “Weak-interaction and nuclear-structure aspects of nuclear double beta decay,” *Physics Reports*, vol. 300, pp. 123–214, Jul 1998.
- [52] M. Agostini *et al.*, “Results on $\beta\beta$ decay with emission of two neutrinos or Majorons in ^{76}Ge from GERDA Phase-1,” *Eur. Phys. J.*, vol. C75, no. 9, p. 416, 2015, 1501.02345.
- [53] J. Boger *et al.*, “The Sudbury neutrino observatory,” *Nucl. Instrum. Meth. A*, vol. 449, pp. 172–207, 2000, nucl-ex/9910016.

- [54] S. Andringa *et al.*, “Current Status and Future Prospects of the SNO+ Experiment,” *Adv. High Energy Phys.*, vol. 2016, p. 6194250, 2015, 1508.05759.
- [55] G. Alimonti *et al.*, “The Borexino detector at the Laboratori Nazionali del Gran Sasso,” *Nucl. Instrum. Meth. A*, vol. 600, pp. 568–593, 2009, 0806.2400.
- [56] A. Suzuki, “Antineutrino Science in KamLAND,” *Eur. Phys. J. C*, vol. 74, no. 10, p. 3094, 2014, 1409.4515.
- [57] R. J. Ford, “A Scintillator Purification Plant and Fluid Handling System for SNO+,” *AIP Conf. Proc.*, vol. 1672, p. 080003, 2015, 1506.08746.
- [58] M. Redshaw, B. J. Mount, E. G. Myers, and F. T. Avignone, “Masses of ^{130}Te and ^{130}Xe and Double- β -Decay Q Value of ^{130}Te ,” *Phys. Rev. Lett.*, vol. 102, p. 212502, may 2009.
- [59] R. Arnold *et al.*, “Measurement of the Double Beta Decay Half-life of ^{130}Te with the NEMO-3 Detector,” *Phys. Rev. Lett.*, vol. 107, p. 062504, Aug. 2011, 1104.3716.
- [60] M. Askins *et al.*, “Background Information for the Tellurium Loading Down Select,” Tech. Rep. SNO+-doc-3505, SNO+ collaboration, 2015.
- [61] K. Eguchi *et al.*, “First results from KamLAND: Evidence for reactor anti-neutrino disappearance,” *Phys. Rev. Lett.*, vol. 90, p. 021802, 2003, hep-ex/0212021.
- [62] G. Alimonti *et al.*, “The liquid handling systems for the Borexino solar neutrino detector,” *Nucl. Instrum. Meth. A*, vol. 609, pp. 58–78, 2009.
- [63] C. Arpesella *et al.*, “Direct Measurement of the Be-7 Solar Neutrino Flux with 192 Days of Borexino Data,” *Phys. Rev. Lett.*, vol. 101, p. 091302, 2008, 0805.3843.
- [64] T. Kaptanoglu, “Te Diol 0.5and Documentation,” Tech. Rep. SNO+-doc-3689, University of Pennsylvania, 2016.
- [65] K. A. Olive *et al.*, “Review of Particle Physics,” *Chin. Phys. C*, vol. 38, p. 090001, 2014.
- [66] H. Ejiri, “Nuclear deexcitations of nucleon holes associated with nucleon decays in nuclei,” *Phys. Rev. C*, vol. 48, pp. 1442–1444, 1993.
- [67] S. N. Ahmed *et al.*, “Constraints on nucleon decay via ‘invisible’ modes from the Sudbury Neutrino Observatory,” *Phys. Rev. Lett.*, vol. 92, p. 102004, 2004, hep-ex/0310030.
- [68] T. Araki *et al.*, “Search for the invisible decay of neutrons with KamLAND,” *Phys. Rev. Lett.*, vol. 96, p. 101802, 2006, hep-ex/0512059.

- [69] I. T. Coulter, *Modelling and reconstruction of events in SNO+ related to future searches for lepton and baryon number violation*. PhD thesis, 2013.
- [70] R. Bonventre, “Solar sensitivity studies,” Tech. Rep. SNO+-doc-3690, UC Berkeley, 2016.
- [71] A. Gando *et al.*, “Constraints on θ_{13} from A Three-Flavor Oscillation Analysis of Reactor Antineutrinos at KamLAND,” *Phys. Rev. D*, vol. 83, p. 052002, 2011, 1009.4771.
- [72] J. Klein, M. Neubauer, F. M. Newcomer, and R. Van Berg, “The sno trigger system,” Tech. Rep. SNO+-doc-827-v1, The SNO+ Collaboration, 1997.
- [73] A. Mastbaum, “Electronics overview,” Tech. Rep. SNO+-doc-2850, University of Pennsylvania, 2014.
- [74] R. Van Berg, J. Klein, and J. Wilkerson, “Design of the SNO timing system,” Tech. Rep. SNO-STR-95-008, The SNO Collaboration, 1995.
- [75] G. Prior, “Data Quality Low Level Checks,” Tech. Rep. SNO+-doc-3713-v3, LIP, 2016.
- [76] M. Stringer, “The Calibration Data Quality Processors,” Tech. Rep. SNO+-doc-3726-v1, University of Sussex, 2016.
- [77] J. Caravaca, F. Descamps, N. Gagnon, P. Harvey, M. Howe, B. Land, and A. Latorre, “Standard runs, run type word and operator GUI,” Tech. Rep. SNO+-doc-3446-v1, The SNO+ Collaboration, 2016.
- [78] The SNO+ Collaboration, “The SNO+ RAT Companion,” 2016.
- [79] J. Kaspar, “SNO+ data quality call,” May 2013. Fortnightly phone meeting.
- [80] J. Wilson, *A Measurement of the ^8B Solar Neutrino Energy Spectrum at the Sudbury Neutrino Observatory*. PhD thesis, Jesus College, Oxford, 2004.
- [81] The Scipy community, “What is numpy?,” 2008.
- [82] The Scipy community, “Introduction,” 2008.
- [83] E. W. Weisstein, “Convolution, from mathworld—a wolfram web resource,” 2016.
- [84] S. Baker and R. D. Cousins, “Clarification of the Use of Chi Square and Likelihood Functions in Fits to Histograms,” *Nucl. Instrum. Meth.*, vol. 221, pp. 437–442, Apr. 1984.
- [85] R. J. Barlow, *Statistics: A Guide to the Use of Statistical Methods in the Physical Sciences*. Wiley, 1989.

-
- [86] A. Mastbaum, “ $0\nu\beta\beta$ -Phase Signal Extraction with a Maximum Likelihood Fit,” Tech. Rep. SNO+-doc-2266-v1, University of Pennsylvania, 2014.
- [87] G. J. Feldman and R. D. Cousins, “A Unified approach to the classical statistical analysis of small signals,” *Phys. Rev. D*, vol. 57, pp. 3873–3889, 1998, physics/9711021.
- [88] P. Uwer, “EasyNData: A Simple tool to extract numerical values from published plots,” 2007, 0710.2896.
- [89] S. Ashahi *et al.*, “DRAFT Report of the ^{130}Te Verification Task Force,” Tech. Rep. SNO+-doc-1728, SNO+ Collaboration, 2013.
- [90] V. Lozza, M. C. Chen, and H. O’Keeffe, “Expected radioactive backgrounds in SNO+,” Tech. Rep. SNO+-doc-507v27, Queens University, ON, 2014.
- [91] V. Lozza, “Background table & document,” Tech. Rep. SNO+-doc-1981-v21, SNO+ collaboration, 2013.