University of Sussex

A University of Sussex PhD thesis

Available online via Sussex Research Online:

http://sro.sussex.ac.uk/

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details



Optical Calibration System for SNO+ and Sensitivity to Neutrinoless Double-Beta Decay

James Waterfield

Submitted for the degree of Doctor of Philosophy University of Sussex September 2016

UNIVERSITY OF SUSSEX

JAMES WATERFIELD, DOCTOR OF PHILOSOPHY

<u>Optical Calibration System for SNO+</u> AND SENSITIVITY TO NEUTRINOLESS DOUBLE-BETA DECAY

Abstract

The SNO+ experiment is primarily looking for neutrinoless double-beta decay, an unobserved, lepton number violating radioactive decay. This is achieved by loading liquid scintillator with tellurium whose isotope ¹³⁰Te decays via double beta decay with a Q-value of 2527 keV. An optical calibration system, located outside the scintillator, has been developed to help meet the radiopurity requirements of the experiment. This thesis describes the hardware component of the optical calibration system which calibrates the timing and charge response of the photomultiplier tube array of SNO+. A set of quality assurance tests showed that the system was at the required standard for installation. Data taken with SNO+ and the optical calibration system showed that the system was stable enough for photomultiplier tube calibration, identified resolvable issues with the SNO+ data acquisition system and allowed measurement of single photoelectron spectra. Data quality checks have been developed to ensure data is of calibration standard. The sensitivity of SNO+ to neutrinoless double-beta decay with nearly 800 kg of 130 Te and five years data taking is investigated with a comprehensive evaluation of systematic uncertainties. Two new methods for acquiring a greater sensitivity to neutrinoless double-beta decay were developed; a one dimensional fit in event energy and a multidimensional fit in event energy and position. A simple event counting analysis, developed previously by the collaboration, was shown to be sensitive to systematic uncertainties. A fit in an extended energy range was shown to constrain the systematics and achieve a half-life sensitivity of 9.30×10^{25} yr corresponding to a 5.6% improvement over the counting analysis which neglected systematic uncertainties. The multidimensional analysis with systematics included achieved a 20% improvement over the counting analysis with a half-life sensitivity of 1.06×10^{26} yr, corresponding to an effective Majorana mass between 52 to 125 meV.

Acknowledgements

There have been many people who have helped and encouraged me throughout my academic career and I can only mention a small number of them here.

Firstly, I would like to thank my primary supervisor Jeff Hartnell whose guidance, support and enthusiasm have led me through this process and allowed me to achieve my goals. I would particularly likely to thank Jeff for all the time and patience he has dedicated to me over the four years of my research. Without his help this thesis would not have been possible.

I would also like to thank Simon Peeters who has been a constant source of support over the years. Simon was always present for any questions that I might have and was a great help for my work on the TELLIE system. Richard White also deserves a lot of praise as he was influential on the development of the TELLIE system and without his hard work TELLIE would not have been built.

Within SNO+ I would like to thank Jose Maneira for all the help he gave me during the quality assurance tests. I would also like to thank Frejia Descamps who has proved to be a valuable source of information about the PMT calibration. Thanks also to Josh Klein who gave me essential feedback during the development of the sensitivity analyses. A big thank you goes to Matt Mottram and Phil Jones who helped me with all things software related as well as introducing me to some very fine beers. Thanks also to Ashley Back and Evelina Arushanova who proved to be great friends and housemates during my stay in Canada.

There were many people who I would like to thank for my time at SNOLab. It was a very enjoyable experience and I learned a great deal about the experiment. A special thank you goes to Mike Schwendener who I not only learnt a great deal from but who, along with his brilliant partner Jen Davidson, welcomed me into the Sudbury community and made me feel like I was at home.

Life in Brighton has been amazing and it surely would not have been as much fun without the rest of my peers in the EPP group at Sussex: James Sinclair, Ed Leming, Mark Stringer, Yusufu Shehu, Zara Grout, Stewart Martin-Haugh, Nicky Santoyo, Luke Vinton, Fabrizio Miano, Nicola Abraham, Tristan Blackburn, Oliver Winston, Fabio Tresoldi, Emma Winkles, Diana Méndez, Sam Jones and not forgetting the eternal student Steve Churchwell. They have all become great friends and an invaluable source of advice whether my questions related to physics or otherwise.

A big thank you to all the "Cov Heads". You are like a family to me and I will always be there for you like you were there for me when I needed it the most. There is a special place in my heart for you all.

A special mention goes to Grace Meheux whose love and support helped me through very difficult times. It is also important to thank Mahalia Carroll who I spent a great few years with and a much needed, fantastic trip to California after 6 months of freezing Sudbury temperatures.

Lastly and most importantly I would like to thank my family, particularly my parents and brother who I love very dearly. This thesis would have been impossible without you.

Preface

The general goals of this thesis was suggested by my supervisor, Jeff Hartnell.

Chapter 1 introduces the thesis and its motivation.

Chapter 2 describes the physics of neutrinos and the experimental status of the field. It has been written from a combination of various sources.

Chapter 3 describes the SNO+ experiment. It has been written from a combination of various sources as well as personal experience.

Chapter 4 describes the design, testing and operation of the optical calibration system: TELLIE. TELLIE was designed primarily by the University of Sussex prior to the start of my research. The time profile of the TELLIE channel as measured by the single photon counting set up was taken and analysed by James Sinclair. The light intensity measurements of the optical fibres were taken by the Lisbon group and the high resolution measurement was taken by the Sussex group prior to my research. The characterisation and quality assurance of the TELLIE channels was undertaken and analysed by myself. The goals of the TELLIE commisioning period were designed by myself and Freija Descamps. The analysis of the stability of TELLIE; the test of the SNO+ DAQ system and the gated charge spectrum is my work. Freija Descamps designed and ran the PMT calibration code. The data quality checks and code was designed by me.

Chapter 5 describes the methodology of the sensitivity studies of neutrinoless doublebeta decay. The methods and framework used was designed by myself and Ashley Back.

Chapter 6 presents the results of the sensitivity studies of neutrinoless double-beta decay. All analyses presented are my own.

Chapter 7 summarises and concludes the thesis.

Contents

Li	List of Tables xvi				cvi	
Li	st of	Figure	es		xx	xv
1	Intr	oducti	on			1
2	The	Theo	retical and Experimental Status of Neutrino Physics			5
	2.1	The N	eutrino	•		5
	2.2	Neutri	no Oscillations	•		8
	2.3	Mass o	of Neutrinos	•	 •	10
		2.3.1	Mass Splittings	•	 •	11
		2.3.2	Oscillation Parameters	•	 •	12
		2.3.3	Particle Decay Constraints	•		12
		2.3.4	Cosmological Constraints	•		14
	2.4	Major	ana Neutrinos	•		14
		2.4.1	See-Saw Mechanism	•		15
	2.5	Neutri	noless Double-Beta Decay	•		16
		2.5.1	Double Beta Experiments		 •	22
3	The	SNO-	+ Experiment			27
	3.1	The S	NO+ Detector	•	 •	27
	3.2	Electro	onics	•	 •	30
	3.3	Calibr	ation of $SNO+$	•		32
		3.3.1	Electronics calibration	•		33
		3.3.2	Optical calibration sources	•		34
		3.3.3	Radioactive calibration sources	•		37
	3.4	Liquid	scintillator	•		38
		3.4.1	Linear Alkyl-Benzene	•		38

		3.4.2	Tellurium
	3.5	Backg	rounds
		3.5.1	Tellurium Backgrounds
		3.5.2	Scintillator Backgrounds
		3.5.3	External Radioactive Backgrounds
		3.5.4	Cosmogenics
		3.5.5	Solar Neutrinos
		3.5.6	Pileup
	3.6	Monte	e-Carlo
4	РМ	T Tim	ing and Gain Calibration System (TELLIE) 49
-	4.1	Requi	rements of PMT Calibration 49
	4.2	Hardw	vare
		4.2.1	LED selection
		4.2.2	Push-Pull Driver
		4.2.3	Optical Fibre
		4.2.4	PIN Diodes
	4.3	Softwa	are
	4.4	Chara	cterisation and Quality Assurance of
		TELL	IE channels
		4.4.1	LED Quality Assurance
		4.4.2	Push-Pull Quality Assurance
	4.5	TELL	IE commissioning
		4.5.1	Goals
		4.5.2	Transition of TELLIE Light in SNO+
		4.5.3	Results
		4.5.4	Conclusion
	4.6	Data (Quality processor
		4.6.1	Pulse intensity checks
		4.6.2	Trigger check
		4.6.3	Pulse delay check
		4.6.4	Fibre check
		4.6.5	Timing checks
		4.6.6	Results of checks
	4.7	Conclu	usion

5	Me	hodology For Extracting Limits On Neutrinoless Double-Beta De-	
	cay	1	10
	5.1	Likelihood technique	10
		5.1.1 Verification of the penalty term $\ldots \ldots \ldots$	13
	5.2	Parametrised Monte Carlo	15
		5.2.1 Energy Resolution $\ldots \ldots \ldots$	15
		5.2.2 Energy Scale $\ldots \ldots \ldots$	18
		5.2.3 Energy Shift	19
		5.2.4 Energy Reconstruction Of Backgrounds Originating Outside The	
		Scintillator	.20
		5.2.5 Background rates	25
	5.3	Minimisation	25
		5.3.1 Grid Search	25
		5.3.2 MINUIT	26
	5.4	Conclusion	26
6	SN	0+ Sensitivity to Neutrinoless Double Beta Decay 1	28
	6.1	Current State of the Art Sensitivity Calculations	.29
		6.1.1 Transition to RAT 5.0.2	.32
	6.2	Energy Range and Binning Optimisation	33
		6.2.1 Optimisation With Systematic Uncertainties	34
	6.3	Application of systematic uncertainties	35
		6.3.1 Individual Systematics	.37
		6.3.2 Correlations Between Systematic Uncertainties	53
		6.3.3 Combined Systematic Uncertainties Study	62
	6.4	Multidimensional Limit Setting	67
		6.4.1 Optimisation With Fixed Backgrounds	.69
		6.4.2 Optimisation With Systematic Uncertainties	70
		6.4.3 Combined Systematic Uncertainties Study	71
	6.5	Summary of Results	73
	6.6	Potential Further Studies	76
7	Cor	clusions 1	77
	7.1	TELLIE	78
	7.2	Sensitivity to Neutrinoless Double-Beta Decay	79

X	
A The radioactive decay chains of 238 U and 232 Th	182
Bibliography	184

List of Tables

2.1	Results of a global three neutrino oscillation analysis from January 2016 [27].	
	Best fit values are given for each oscillation parameter with one, two and	
	three σ ranges for each hierarchy. The value for Δm^2 is positive for normal	
	hierarchy (NH) and negative for inverted hierarchy (IH)	13
2.2	The phase space factor $G^{0\nu}$ for the standard mechanism of neutrinoless	
	double-beta decay; Q -value and natural abundance for isotopes where double	
	beta decay is observable and have a Q -value >2 MeV	18
3.1	The levels of isotopes within the liquid scintillator as measured by the	
	Borexino experiment.	41
3.2	The maximum energy E_{max} of the emitted beta and its branching ratio	
	(BR) of the predominant (BR $\geq 1\%$) decay modes of the isotopes 214 Bi and	
	²⁰⁸ Tl	43
3.3	The leaching rates (LR) and leaching time constants (τ) of the isotopes	
	$^{210}\mathrm{Pb}$ and $^{210}\mathrm{Po}$ in various medium [90]. The measurements were taken at	
	two temperatures 25 and 12 °C. \ldots	45
3.4	The half life $(T_{1/2})$ decay, branching ratio (BR) and maximum energy of the	
	decay (E_{max}) of muon-induced backgrounds produced through interactions	
	with ${}^{12}C$ within the scintillator [91]	46

4.1 $\,$ The number of simulated photons per pulse in a TELLIE light pulse; prob-

	ability that a PMT will acquire a hit; per injection point and total TELLIE	
	running times for the different phases of SNO+. The hit probabilities were	
	calculated from Monte Carlo simulation of a TELLIE channel consisting of	
	a 504 nm LED at a given intensity measured in photons per pulse [97]. The	
	run times are calculated by assuming a pulse width of 6 ns and a desired	
	accuracy of 1 ns running at 10 Hz. The centre and fringe areas of PMTs	
	are defined as being between 0° and 14.5° and 14.5° and 29° respectively	
	where the angle is defined as being between the incident fibre direction and	
	the PMT position.	54
	-	
4.2	The measured full width half maximum (FWHM), rise and fall time of light	
	pulses at various intensities using the single photon counting method. All	
	measurements made with the driver running at $1\mathrm{kHz}$. The measurement	
	of the number of photons is performed after 45.5 m of optical fibre and	
	without a 2 m patch cable.	64
4.3	The results and criteria of the optical fibre quality assurance tests. To pass	
	the fibres must meet the denoted criteria for maximum intensity, I_{max} ; the	
	half width angle at 20% of the peak intensity and rise time of the time	
	profile of the fibre. The mean and RMS are the results of all 220 fibres	
	constructed for the TELLIE system [95]	71
4.4	The range of different parameters that can be set for a TELLIE channel or	
	channels for a sequence of pulses. Most parameters accept two 8-bit integer	
	inputs, denoted in the Hi/Lo Range column, except Fibre and Trigger Delay $% \mathcal{A}$	
	which accept one. These inputs are used to calculated either an integer	
	value or time which is set for the sequence. Although frequency can be	
	technically set to be infinity, a hard coded limit of 10 kHz has been set so	
	that the TELLIE channel operates safely	74
15	The cute used in the timewalk PMT collibration and the criteria required	
4.0	for a PMT channel to fail the cut	03
		90
4.6	The cuts used in the gain fit PMT calibration and the criteria required for	
	a PMT channel to fail the cut.	93

4.7	The types and number of triggers recorded when a TELLIE event has been	
	missed. Note that Orphan is not a recorded trigger type but denotes when	
	the builder has failed to store an event in the data file. PULSE_GT is a	
	global trigger generated by a pulser within the SNO+ trigger system. SYNC	
	is a trigger signal which is used to synchronise the DAQ system with the	
	GPS system on surface. Other trigger types are defined in Section 3.2	94
4.8	The expected TAC counts and the lower and upper bounds of TAC counts	
	for a given medium of the position of the reflection peaks with respect to	
	the direct light peak. The lower and upper bounds correspond to a 10 ns	
	buffer	104
4.9	The results of the data quality checks on February 2014 data. \ldots	107
5.1	The values of $\Delta \chi^2$ which correspond to a median confidence limit for sim-	
	ultaneous estimation of m parameters [108]	112
5.2	A table of the different external backgrounds arranged by their original	
	isotope location within the detector. The number of events which have a	
	reconstructed (Reco) or true quenched energy deposited (True) between 2.2	
	and 2.8 MeV is listed. The misreconstruction rate (Misreco) as calculated by	
	5.18 is also listed for each background. All external backgrounds in the rat	
	4.5 production were used and they were not weighted by the expected rate	
	of each background component. A fiducial volume cut at a reconstructed	
	radius of 3.5 m was made.	121
6.1	A list of experimental parameters used in the previous SNO+ sensitivity	
	analysis [112, 113]. The same experimental parameters were used for the	
	verification of echidna with this analysis. The BiPo rejection factors are	
	applied to both 214 Bi- 214 Po and 212 Bi- 212 Po decays	130
6.2	A table showing the different background types ranked in ascending order by	
	their extracted neutrinoless double-beta decay half-life at 90% confidence	
	level, $T_{1/2}$ 90% CL using the single bin analysis. Each background type	
	was reduced to zero in the fit and was fixed at its expected value. The	
	experimental parameters listed in 6.1 were used	131
6.3	A list of all backgrounds in descending order of their expected count rate	
	for three energy ranges with the assumptions in Table 6.1. \ldots	136

- 6.12 The correlation coefficient matrix for the multidimensional analysis in the optimistic systematics scenario. Global is the overall correlation of the parameter with all other parameters. LY is light yield which is directly related to energy resolution. Scale and shift are the energy scale and shift parameters. $2\nu 2\beta$, ⁸B ν , ²¹²BiPo, U, Th and Ext correspond to the ¹³⁰Te two neutrino double-beta decay, ⁸B neutrino elastic scattering, ²¹²Bi²¹²Po, uranium, thorium-chain and external background rates respectively. 172
- 6.13 The correlation coefficient matrix for the multidimensional analysis in the conservative systematics scenario. Global is the overall correlation of the parameter with all other parameters. LY is light yield which is directly related to energy resolution. Scale and shift are the energy scale and shift parameters. $2\nu 2\beta$, ⁸B ν , ²¹²BiPo, U, Th and Ext correspond to the ¹³⁰Te two neutrino double-beta decay, ⁸B neutrino elastic scattering, ²¹²Bi²¹²Po, uranium, thorium-chain and external background rates respectively. 172

6.15 The reductions in half-life sensitivity for each systematics when the systematics are investigated individually and when correlations are accounted for. The wide 1.8 to 3.2 MeV, 40 keV bins energy fit analysis was used with the optimistic systematics scenario. For comparison, the results where systematics are investigated individually for the narrow 2.2 to 2.8 MeV, 20 keV bins energy fit analysis are included. For the individual systematics studies a coarse binning was used so the error on each reduction is $\pm 0.18\%$ 175

List of Figures

2.1	A pictorial representation of the normal hierarchy, left, and inverted hier-	
	archy, right in terms of their mass states and their flavour states $[26]$ 1	2

- 2.4 Values of the nuclear matrix element, M^{0ν}, for several isotopes using different nuclear physics models. These are the Nuclear Shell Model (NSM);
 Quasi particle Random Phase Approximation (QRPA); Interacting Boson Model (IBM); Energy Density Functional (EDF) and the Projected Hartree-Fock-Bogoliubov model (PHFB). The error bars are due to theoretical uncertainties arising from varying model parameters. [35]

- 2.8 A stacked plot showing the energy spectrum of the expected backgrounds and observation of neutrinoless double-beta decay with an effective Majorana mass of $\langle m_{\beta\beta} \rangle = 200$ meV in SNO+. Inset is a residual histogram showing the signal contribution above the expected background rate. The events shown are for a 3.5 m fiducial volume cut, 0.3% Te loading, and a resolution of 200 NHit/MeV. $T_{\beta\beta}$ is the observed kinetic energy [67]. . . . 26
- 3.1 A schematic of the SNO+ detector in the SNO cavern (black). The acrylic vessel is shown in blue and is suspended from the deck above from the hold-up ropes (pink) and held in position by the hold-down ropes (red). The stainless steel struts of the PMT support structure are shown in grey [69]. 28
- 3.2 Diagram of the SNO+ trigger system. The PMT Interface Cards (PMTICs) supply the PMTs with high voltage and read out the signal passing it to the Daughter Board (DB) which calculates the charge and timing of the pulse. The Front End Card (FEC) then reads out the DBs and passes the trigger waveform to the Crate Trigger Card (CTC) which sums the triggers of an entire crate. The Analogue Master Trigger Card then sums up the triggers on all crates and this is passed to the Digital Master Trigger Card (MTC/D) which decides whether the trigger conditions gave been met to issue Global Trigger (GT) which will force all triggered channels to have their charge and time read out. The time taken for a signal to propagate through the electronics is denoted in red with the default values used to simulate the DAQ electronics in the Monte-Carlo in brackets.

- 3.3 An example of a time slope calibration taken suring SNO. The Time-to-Amplitude Converters (TAC) readout is plotted against the TAC time in nanoseconds. TAC times right of the green line are denoted as 'curl' as they are not monotonically increasing with TAC ADC counts. TAC ADC counts above the red dotted line are denoted as 'pre-curl' and can have two values of TAC time. These are flagged by the calibration. Adapted from [77]. . . 34
- 3.4 The activity of ²²²Rn as a function of time within the KamLAND detector [78]. Rates were calculated using ²¹⁴Bi-²¹⁴Po coincidences within a 5.5 m volume. The red and blue dashed lines indicate the start of calibration runs.
- 3.6 Sketch of a deployed source within the AV. The position of the source can be manipulated in three dimensions through the ropes and umbilical [72]. 37
- 3.7 A flow diagram of the SNO+ liquid scintillator processing and purification facility. The liquid scintillator is first stored on surface before being shipped underground for purification and eventually filling the inner AV [86]. . . . 41
- 4.1 A sketch of integrated charge (QHS) vs. time. The charge deposited in a channel is measured in counts above pedestal (cap), where a pedestal is the average number of counts a channel readouts when no charge is deposited. With respect to the centre of the peak, larger pulses cross the discriminator threshold (dashed horizontal line) at an earlier time than smaller pulses. . . 51

- 4.2 A photo of the components that host the electronics of the timing component of Embedded LED Light Injection Entity (ELLIE) housed in the ELLIE rack on the deck of the SNO+ detector. The electronics handles the control and light pulse generation of the TELLIE system. They consist of a power supply, a control box and twelve boxes with each containing eight light emitting channels. The boxes are daisy chained via a ribbon cable to the control box which handles communication between the DAQ system and the boxes. The control box is connected to the DAQ system via a USB port. The control box also hosts two BNC connections. The left BNC connection accepts trigger signals from the SNO+ DAQ system and the right outputs trigger signals from TELLIE to the SNO+ DAQ system. The power supply is located at the top of the ELLIE rack.
- 4.3 Left: A photo of an LED that has been drilled and has had an optical fibre attached. The ruler is in centimetres. Right: A photo of two TELLIE channels. Each channel consists of a driver board, an optical fibre, one of the four electronic channels housed on a PIN diode readout board and a coupler made from common brass plumbing fitments with custom end caps to make the coupler light tight. The coupler houses an LED, fibre connection and a photodiode which monitors the light output of the LED. 56
- 4.4 A picture of the internal electronics of a TELLIE box consisting of eight channels, housed on one motherboard, each containing one LED driver, coupler, fibre, terminating at the front of the box at an ST connector and a channel on one of two PIN diode boards. Two identical power sockets and a serial port for communication are visible on the front of the box. 57

- 4.11 A simplified diagram of two different stable light transmission paths, or modes, within an optical fibre. The different path lengths of the modes causes different propagation times through the optical fibre and hence a broadening of the time profile of the initial light pulse. The effect of this broadening is known as modal dispersion.
- 4.12 A projection of the PMT Support Structure (PSUP) and the location of installed fibres at the time of the commissioning runs. The black lines indicate the struts which make up the geodesic sphere. The different dashing and number of lines identify the different types of strut which vary in length and thickness. Between the struts are panels which host the PMTs. The struts are connected via three different types of hubs at 91 nodal positions indicated by numbers which provide the positions that the TELLIE fibres are mounted to. Each fibre is labelled with a fibre number which is the same as the nodal position it is connected to. The red circles indicate fibre positions which were installed via ladder in March 2012 while the cavity was not filled with water. The blue circles indicate fibre positions which were installed in October and November 2014. The upper blue circles had fibres installed by climbing on top of the PSUP where as the lower blue circles were installed via boat. The remaining fibres will be installed via boat as the water level increases in the cavity.
- 4.14 A schematic of the bench top set up at Lisbon used for intensity measurements of light emitted from optical fibre at different angles. Three duplex fibres which each host two fibres can be mounted in the dark box at a time. The PMT, powered by a high voltage power supply, has its X position fixed at 51 mm away from the termination point of the optical fibre. A measurement of the current of the PMT was recorded by a DAQ computer connected to a picoammeter at 1 mm steps in the Y direction. The position of the fibres was changed through a motorised Table top which was controlled by the DAQ computer. A drilled LED with a fibre inserted was connected to each of the fibres individually via ST connection. The LED was powered by a power supply at a constant voltage.

66

67

- 4.17 A diagram denoting the time delays due to light and signal propagation from when a TELLIE driver control board is triggered to when the SNO+ DAQ system registers an event [95]. The total delay is estimated to be approximately 615 ns. There is also a delay of approximately 70 ns between when the trigger signal is emitted by TELLIE and when it is detected by the SNO+ DAQ. The trigger delay should therefore by set to approximately 545 ns. A delay can be applied to when the trigger to the SNO+ DAQ is sent to register the TELLIE event. This can be set in steps of 5 ns between 0 and 1275 ns.

71

4.18	A graph showing the mean and RMS of the wavelength dependent intensity	
	spectrum of 120 LEDs. All spectra were normalised so that the integral of	
	each spectrum is equal to one.	76
4.19	A flow diagram of the test bench set up used to characterise the light output	
	of each TELLIE channel as a function of IOP and IPW values and to tune	
	each channel's PIN board used to monitor each channel's light output during	
	operation once installed in SNO+	78
4.20	Cathode radiant sensitivity of the PMT $[99]$ as a function of the wavelength	
	of the incident photon $[103]$. The solid and dashed lines are the sensitivities	
	of the -110 and -210 models of the PMT respectively. The -110 model was	
	used for the quality assurance measurements	80
4.21	A graph showing the average, standard deviation and range of number of	
	photons per pulse as a function of the pulse height parameter IOP of all 96	
	driver boards used in the TELLIE system. The IPW parameter was set to	
	0 to produce the highest number of photons per pulse. \ldots \ldots \ldots \ldots	81
4.22	The mean, standard deviation and range of the full width half maximum	
	(FWHM) of the time profiles of the photon pulses of all 96 channels used	
	in TELLIE as a function of the number of photons per pulse. The FWHM	
	was measured by the experimental set up described in Figure 4.19. The	
	IOP value was set to its maximum value of 16383	82
4.23	The mean, standard deviation and range of the rise times of the time profiles	
	of the photon pulses of all 96 channels used in TELLIE as a function of	
	the number of photons per pulse. The rise times were measured by the	
	experimental set up described in Figure 4.19. The IOP value was set to its	
	maximum value of 16383	82
4.24	The mean, standard deviation and range of the fall times of the time profiles	
	of the photon pulses of all 96 channels used in TELLIE as a function of	
	the number of photons per pulse. The fall times were measured by the	
	experimental set up described in Figure 4.19. The IOP value was set to its	
	maximum value of 16383.	84

xxiv

- 4.25 A two dimensional histogram showing the number of hits that a PMT registers during a TELLIE run. The PMTs have been binned as a function of distance from the emission point of the optical fibre within the detector and the time at which a PMT hit occurs. Data was taken during the February commissioning runs when the AV contained only air and only a fraction of PMTs below the AV were submerged in water.
- 4.26 A two dimensional histogram showing the number of hits that a PMT registers during a TELLIE run. The PMTs have been binned as a function of distance from the emission point of the optical fibre within the detector and the time at which a PMT hit occurs. Data was taken during the December commissioning runs when the cavity and the AV was partially filled with water. The data points highlighted in the red circle occur due to initial reflections off the side of the AV closest to the fibre. These data points are used for the AV location code. The data points highlighted in black is due to the direct light illuminating PMTs at the opposite side of the detector. These data points are used for PMT calibration. The data points highlighted in the dark blue ellipse are mainly due to water level reflections and scattering. The data points highlighted in the white ellipse are mainly due to reflections off the PMT concentrators. There is a peak in reflections off the concentrators when light is entering the PMT at a 35° incident angle. . . .
- 4.27 A profile histogram of R_{NHit} , defined in equation 4.10, as a function of the time of the event relative to the first triggered TELLIE event of the run. The plot shows the average intensity rise of 30 PMT calibration runs from the March 2014 commissioning period. The runs were taken at 500 Hz and are binned in 5 second bins. The average initial NHit of the runs was 52.7 ± 11.7 . The black line shows the result of a linear fit to the data. The gradient of the linear fit is $6.0 \pm 0.1 \times 10^{-5}$ s⁻¹ and its intercept is 0.0545 ± 0.0002 . The reduced chi-squared of the fit is 5.99.
- 4.28 A histogram of R_{NHit} , defined in equation 4.10, of all 30 PMT calibration runs taken at 500 Hz from the March 2014 commissioning period. The average initial NHit of the runs was 52.7 ± 11.7 . The total number of events recorded in the runs was 4.69×10^6 . The mean value of R_{NHit} of all events is 0.065 and the value of the RMS is 0.246.

87

88

 4.29 The black points show the value of ECA calibrated time versus the measured QHS at PMT 1851. The red points show the interpolation points used in the PMT calibration code. The first six interpolation points with lower values of QHS show timewalk at the PMT whereas the four interpolation points higher in QHS show the offset due to the PMT cable 90 4.30 The results of the gain fit of the PMT calibration code for PMT 2790. The peak, denoted with a blue line, is the centre of a sliding window which contains the most hits. The high half point, denoted with a black line, is the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window			
QHS at PMT 1851. The red points show the interpolation points used in the PMT calibration code. The first six interpolation points with lower values of QHS show timewalk at the PMT whereas the four interpolation points higher in QHS show the offset due to the PMT cable	4.29	The black points show the value of ECA calibrated time versus the measured	
 the PMT calibration code. The first six interpolation points with lower values of QHS show timewalk at the PMT whereas the four interpolation points higher in QHS show the offset due to the PMT cable 90 4.30 The results of the gain fit of the PMT calibration code for PMT 2790. The peak, denoted with a blue line, is the centre of a sliding window which contains the most hits. The high half point, denoted with a black line, is the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window		QHS at PMT 1851. The red points show the interpolation points used in	
 values of QHS show timewalk at the PMT whereas the four interpolation points higher in QHS show the offset due to the PMT cable 90 4.30 The results of the gain fit of the PMT calibration code for PMT 2790. The peak, denoted with a blue line, is the centre of a sliding window which contains the most hits. The high half point, denoted with a black line, is the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window		the PMT calibration code. The first six interpolation points with lower	
 points higher in QHS show the offset due to the PMT cable 90 4.30 The results of the gain fit of the PMT calibration code for PMT 2790. The peak, denoted with a blue line, is the centre of a sliding window which contains the most hits. The high half point, denoted with a black line, is the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window		values of QHS show timewalk at the PMT whereas the four interpolation	
 4.30 The results of the gain fit of the PMT calibration code for PMT 2790. The peak, denoted with a blue line, is the centre of a sliding window which contains the most hits. The high half point, denoted with a black line, is the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window 91 4.31 A histogram showing the difference in time of a non TELLIE event occurring and the previous TELLIE event during a time window where there is a suspected missing event during a total of 50 PMT calibration runs. A cut has been made such that only windows which contain one non TELLIE event are shown. The spike observed in the distribution occurs at a time when one would expect the missing TELLIE event. The true pulse delay of the runs was approximately 1.21 ms		points higher in OHS show the offset due to the PMT cable	90
 4.30 The results of the gain ht of the PMT cambration code for PMT 2790. The peak, denoted with a blue line, is the centre of a sliding window which contains the most hits. The high half point, denoted with a black line, is the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window	4.90		00
 peak, denoted with a blue line, is the centre of a sliding window which contains the most hits. The high half point, denoted with a black line, is the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window	4.30	The results of the gain ht of the PMT calibration code for PMT 2790. The	
 contains the most hits. The high half point, denoted with a black line, is the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window		peak, denoted with a blue line, is the centre of a sliding window which	
 the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window		contains the most hits. The high half point, denoted with a black line, is	
 equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window		the centre of the sliding window, moving to the right, when its integral is	
 is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window		equal to half the integral at the position of the peak value. The threshold	
 contains half the entries of the peak window		is the centre of a differently sized sliding window, moving to the left, which	
 4.31 A histogram showing the difference in time of a non TELLIE event occurring and the previous TELLIE event during a time window where there is a suspected missing event during a total of 50 PMT calibration runs. A cut has been made such that only windows which contain one non TELLIE event are shown. The spike observed in the distribution occurs at a time when one would expect the missing TELLIE event. The true pulse delay of the runs was approximately 1.21 ms. 95 4.32 Screenshots of the SNO+ event display of the gated charge spectrum run. The two displays are of the summation of TAC counts during the run. 96 4.33 A histogram of the time at which PMTs were triggered during the gated charge spectrum run. The first peak corresponds to the PMTs in crate 5 receiving pedestal triggers. The following three peaks corresponds to the time profile of TELLIE light. 97 4.34 A histogram of the charge collected on PMT 2779 during the gated (black) and ungated (red) charge spectrum runs. The large initial peak in the gated charge spectrum run is the charge collected on the PMT due to noise when no photoelectron is produced. Both histograms are normalised so there integrals are equal to 1. 97 4.35 A histogram of the charge collected on PMT 2779 during the gated (black) and ungated (red) charge spectrum runs. The histogram is integer binned from 20 to 200 cap to remove the noise peak in Figure 4.34 Both histograms are normalised so there integrals are equal to 1. 		contains half the entries of the peak window.	91
 and the previous TELLIE event during a time window where there is a suspected missing event during a total of 50 PMT calibration runs. A cut has been made such that only windows which contain one non TELLIE event are shown. The spike observed in the distribution occurs at a time when one would expect the missing TELLIE event. The true pulse delay of the runs was approximately 1.21 ms	4.31	A histogram showing the difference in time of a non TELLIE event occurring $% \left[{{\left[{{{\rm{TE}}_{{\rm{TE}}}} \right]}_{{\rm{TE}}}} \right]$	
 suspected missing event during a total of 50 PMT calibration runs. A cut has been made such that only windows which contain one non TELLIE event are shown. The spike observed in the distribution occurs at a time when one would expect the missing TELLIE event. The true pulse delay of the runs was approximately 1.21 ms		and the previous TELLIE event during a time window where there is a	
 has been made such that only windows which contain one non TELLIE event are shown. The spike observed in the distribution occurs at a time when one would expect the missing TELLIE event. The true pulse delay of the runs was approximately 1.21 ms. 4.32 Screenshots of the SNO+ event display of the gated charge spectrum run. The two displays are of the summation of TAC counts during the run. 96 4.33 A histogram of the time at which PMTs were triggered during the gated charge spectrum run. The first peak corresponds to the PMTs in crate 5 receiving pedestal triggers. The following three peaks corresponds to the time profile of TELLIE light. 97 4.34 A histogram of the charge collected on PMT 2779 during the gated (black) and ungated (red) charge spectrum runs. The large initial peak in the gated charge spectrum run is the charge collected on the PMT due to noise when no photoelectron is produced. Both histograms are normalised so there integrals are equal to 1. 97 4.35 A histogram of the charge collected on PMT 2779 during the gated (black) and ungated (red) charge spectrum runs. The histogram is integer binned from 20 to 200 cap to remove the noise peak in Figure 4.34 Both histograms are normalised so there integrals are equal to 1. 		suspected missing event during a total of 50 PMT calibration runs. A cut	
 event are shown. The spike observed in the distribution occurs at a time when one would expect the missing TELLIE event. The true pulse delay of the runs was approximately 1.21 ms		has been made such that only windows which contain one non TELLIE	
 when one would expect the missing TELLIE event. The true pulse delay of the runs was approximately 1.21 ms		event are shown. The spike observed in the distribution occurs at a time	
 the runs was approximately 1.21 ms		when one would expect the missing TELLIE event. The true pulse delay of	
 4.32 Screenshots of the SNO+ event display of the gated charge spectrum run. The two displays are of the summation of TAC counts during the run 96 4.33 A histogram of the time at which PMTs were triggered during the gated charge spectrum run. The first peak corresponds to the PMTs in crate 5 receiving pedestal triggers. The following three peaks corresponds to the time profile of TELLIE light		the runs was approximately 1.21 ms.	95
 The two displays are of the summation of TAC counts during the run 96 4.33 A histogram of the time at which PMTs were triggered during the gated charge spectrum run. The first peak corresponds to the PMTs in crate 5 receiving pedestal triggers. The following three peaks corresponds to the time profile of TELLIE light	4.32	Screenshots of the SNO+ event display of the gated charge spectrum run.	
 4.33 A histogram of the time at which PMTs were triggered during the gated charge spectrum run. The first peak corresponds to the PMTs in crate 5 receiving pedestal triggers. The following three peaks corresponds to the time profile of TELLIE light		The two displays are of the summation of TAC counts during the run	96
 charge spectrum run. The first peak corresponds to the PMTs in crate 5 receiving pedestal triggers. The following three peaks corresponds to the time profile of TELLIE light	4.33	A histogram of the time at which PMTs were triggered during the gated	
 receiving pedestal triggers. The following three peaks corresponds to the time profile of TELLIE light		charge spectrum run. The first peak corresponds to the PMTs in crate 5	
 time profile of TELLIE light		receiving pedestal triggers. The following three peaks corresponds to the	
 4.34 A histogram of the charge collected on PMT 2779 during the gated (black) and ungated (red) charge spectrum runs. The large initial peak in the gated charge spectrum run is the charge collected on the PMT due to noise when no photoelectron is produced. Both histograms are normalised so there integrals are equal to 1		time profile of TELLIE light.	97
 and ungated (red) charge spectrum runs. The large initial peak in the gated charge spectrum run is the charge collected on the PMT due to noise when no photoelectron is produced. Both histograms are normalised so there integrals are equal to 1	4.34	A histogram of the charge collected on PMT 2779 during the gated (black)	
 charge spectrum run is the charge collected on the PMT due to noise when no photoelectron is produced. Both histograms are normalised so there integrals are equal to 1		and ungated (red) charge spectrum runs. The large initial peak in the gated	
 no photoelectron is produced. Both histograms are normalised so there integrals are equal to 1		charge spectrum run is the charge collected on the PMT due to noise when	
 integrals are equal to 1		no photoelectron is produced. Both histograms are normalised so there	
4.35 A histogram of the charge collected on PMT 2779 during the gated (black) and ungated (red) charge spectrum runs. The histogram is integer binned from 20 to 200 cap to remove the noise peak in Figure 4.34 Both histograms are normalised so there integrals are equal to 1		integrals are equal to 1	97
and ungated (red) charge spectrum runs. The histogram is integer binned from 20 to 200 cap to remove the noise peak in Figure 4.34 Both histograms are normalised so there integrals are equal to 1	4.35	A histogram of the charge collected on PMT 2779 during the gated (black)	
from 20 to 200 cap to remove the noise peak in Figure 4.34 Both histograms are normalised so there integrals are equal to 1		and ungated (red) charge spectrum runs. The histogram is integer binned	
are normalised so there integrals are equal to 1		from 20 to 200 cap to remove the noise peak in Figure 4.34 Both histograms	
		are normalised so there integrals are equal to 1	98

4.36 A histogram of the difference between the thresholds, ΔT , obtained from the gain fit PMT Calibration code for gated and ungated charge spectra. 99

- 4.37 The difference between the mean pulse delay of a run and the pulse delay of an event. All TELLIE data from Decemember 2014 was used. 102
- 4.38 An example of a run that passes the fibre check. The plot shows a projection of the PSUP with the number of hits each PMT has been triggered on the z-axis. The circle denotes where the fibre check has calculated the centre of the direct light. The cross denotes where the fibre check has calculated the centre the centre of the reflected light and hence the fibre position. 103
- 4.40 An example of a run that passes the time checks described in Section 4.6.5. The blue histogram is the reflected light contained in a 15° area cut around the fibre with respect to the centre of the detector. The red histogram is the direct light contained in a 15° area cut directly opposite the fibre with respect to the centre of the detector. The triangles denote the peaks found by ROOT's TSpectrum class which are used to perform the timing checks. 106

- 5.1 A histogram showing the Gaussian distributions used to verify that the penalty term has been implemented correctly. The same non-normalised Gaussian distribution was used for the signal and the background and has a mean value of 2.5 MeV and a sigma of 0.112 MeV. The background in black was normalised to 100 counts. The histogram in blue shows the scale of the signal at the 90% Confidence Limit (CL) when no penalty term was applied. The histogram in red shows the scale of the signal at the 90% CL when a penalty term was applied where the background rate has a systematic uncertainty of $\sigma = 20$ counts and was varied in steps of one count between 50 and 150 counts.

- 5.7 The number of expected external background counts per year as a function of energy in 0.3% Te-loaded scintinllator with a light yield of 200 NHit/MeV. A fiducial volume cut was made on reconstructed radius R normalised to av radius R_{AV} at $(R/R_{AV})^3 = 0.2$. All events were scaled to the background rates in [110] and an external background event rejection of 50% was assumed. The red points are reconstructed energies, blue and yellow points are true quenched energy deposited in the scintillator. The yellow points have a further energy cut so their reconstructed energy is between 2.2 and 2.8 MeV. The green points are smeared from the blue points by echidnas energy resolution function as decribed in Section 5.2.1. A further weighting was applied to the blue and green points to account for misreconstruction rates described above. The misreconstruction rates for ²¹⁴Bi and ²⁰⁸Tl background types were combined due to the large uncertainty associated with ²¹⁴Bi. The error bars represent statistical and systematic uncertainty. 124

- 6.2 The ¹³⁰Te $0\nu 2\beta$ decay energy spectrum in blue fitted with a Gaussian denoted by the dashed red line. The results of the fit was a mean of 2.476 ± 0.018 MeV and sigma of 0.111 ± 0.012 MeV. The vertical solid line corresponds to the fitted mean and the small dashed vertical lines correspond to \pm one sigma with respect to the mean. The larger dashed vertical lines from left to right correspond to -1/2 sigma and 3/2 sigma with respect to the mean. 130
- 6.4 The half-life at 90% confidence limit as a function of energy range of the region of interest centred at 2.5 MeV. The assumptions used are listed in Table 6.1. Seven different bin sizes were investigated listed in the legend. 133

- 6.7 A graph showing the best fits at the 90% confidence limit as a function of uncertainty in energy resolution. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. . . . 139
- 6.9 A graph showing the best fits at the 90% confidence limit as a function of uncertainty in energy scale. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. 142

- 6.11 A graph showing the best fits at the 90% confidence limit as a function of uncertainty in energy shift. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. 144
- 6.12 Top: the half-life at 90% confidence limit as a function of uncertainty in the 130 Te double beta decay rate. Bottom: the reduction in half-life sensitivity as a function of uncertainty in the 130 Te double beta decay rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line in the bottom plot corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis. . . 146
- 6.13 A graph showing the best fits at the 90% confidence limit as a function of uncertainty in the ¹³⁰Te double beta decay rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins.

- 6.15 A graph showing the best fits at the 90% confidence limit as a function of uncertainty in the ⁸B neutrino elastic scattering rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins.
- 6.16 A graph showing the reduction in half-life sensitivity as a function of uncertainty in the ²¹²Bi²¹²Po pileup rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.

- 6.23 A two dimensional histogram showing the χ^2 space of energy shift on the y-axis and energy scale on the x-axis at the 90% confidence limit. There is a 10% uncertainty in energy scale and a 0.1 MeV uncertainty in energy shift. The best fit of the two parameters is marked with a white cross. . . . 158

6.24	A two dimensional histogram showing the χ^2 space of ¹³⁰ Te double beta
	decay rate normalised to its prior on the y-axis and energy scale on the
	x-axis at the 90% confidence limit. There is a 10% uncertainty in energy
	scale and a 20.3% uncertainty in $^{130}\mathrm{Te}$ double beta decay rate. The best fit
	of the two parameters is marked with a white cross
6.25	A two dimensional histogram showing the χ^2 space of ¹³⁰ Te double beta
	decay rate normalised to its prior on the y-axis and energy shift on the x-
	axis at the 90% confidence limit. There is a 0.1 MeV uncertainty in energy
	shift and a 20.3% uncertainty in $^{130}\mathrm{Te}$ double beta decay rate. The best fit
	of the two parameters is marked with a white cross
6.26	The half-life sensitivity at 90% CL as a function of range of volume, $(R/R_{AV})^3$
	where R is the radius from the centre of the detector and R_{AV} is the radius
	of the Acrylic Vessel. Several bin sizes where investigated, 0.01 in the black
	spots, 0.05 in the vertical blue crosses, 0.1 in the diagonal red crosses and
	0.2 in the magenta diamonds. The analysis with an energy range of $2.2~{\rm to}$
	2.8 MeV was used to encapsulate the entire signal region
6.27	The χ^2 space at the 90% CL sensitivity in the parameters of volume element,
	$(R/R_{AV})^3$, and energy
6.28	The reduction in half life sensitivity as a function of volume element $(R/R_{AV})^3$,
	where R is the radius from the centre of the detector and R_{AV} is the radius
	of the Acrylic Vessel, with a 50% uncertainty on external background rate.
	Two bin sizes where investigated: 0.05 in the vertical blue crosses and 0.1
	in the diagonal red crosses
A.1	The decay chain of 238 U. The energy of the Q values of alpha and beta
	decays are in MeV and gammas in keV. The half life of each isotope is also
	shown [127].
A.2	The decay chain of 232 Th. The energy of the Q values of alpha and beta
	decays are in MeV and gammas in keV. The half life of each isotope is also
	shown [127].

Chapter 1

Introduction

Neutrinos are fundamental particles that may hold the answers to some of the biggest questions in physics today. They can be used to investigate physics beyond the Standard Model, the scale of fermion masses and may explain the matter-antimatter asymmetry of the Universe. Beyond particle physics, they also play a significant role in a large range of physics topics such as geophysics, cosmology, astrophysics and nuclear physics.

In 1998 Super-Kamiokande provided evidence that atmospheric neutrinos oscillate [1] and in 2002 the Sudbury Neutrino Observatory (SNO) proved that neutrinos change flavour [2]. This implies that neutrinos have mass which was not predicted by the Standard Model. The upper limit of the sum of neutrino masses was made by a set of cosmological experiments. The value they found is six orders of magnitude smaller than the next lightest fermion, the electron.

One theory known as the see-saw mechanism could explain the lightness of neutrino mass. This hypothesises the existence of extremely heavy neutrinos which suppress the masses of the three familiar light neutrinos. For this to be true then neutrinos have to be Majorana fermions and hence their own antiparticles. The only fundamental fermions that can be Majorana are neutrinos because they are the only fundamental fermions with neutral charge.

As Majorana fermions are their own antiparticles, interactions involving Majorana fermions may violate lepton number, which is conserved in the Standard Model. Some cosmological models require lepton number violation to explain the matter-antimatter asymmetry in the universe.

The only known and potentially feasible experimental test for determining whether neutrinos are Majorana or not is the detection of neutrinoless double-beta decay. Double beta decay is observable in nuclides where it is energetically forbidden to decay via single
beta decay. In neutrinoless double-beta decay only two electrons are emitted in comparison to standard double beta decay where two neutrinos and two electrons are emitted. Neutrinoless double-beta decay therefore violates lepton number by $\Delta L = 2$ plus other implicit consequences.

The main goal of SNO+ is to search for neutrinoless double-beta decay. It also has a range of other physics goals including examining solar, reactor, supernovae and geo neutrinos. SNO+ is using liquid scintillator as its active mass instead of the heavy water used in SNO. The liquid scintillator will be doped with natural tellurium whose isotope ¹³⁰Te decays via double beta decay with a Q-value of 2527 keV [3]. Liquid scintillator provides a high light yield per unit energy deposited. This lowers the energy threshold and increases the energy resolution of the detector, which is important in distinguishing the monochromatic signal of neutrinoless double-beta decay from the standard double beta decay energy spectrum.

The method SNO+ utilises in its search for neutrinoless double-beta decay is to maximise the signal to background ratio around the Q-value of ¹³⁰Te. Therefore, it is imperative to minimise the risk of contamination from radioactive backgrounds which may reduce this signal to background ratio. As a result, SNO+ needs to limit the number of times calibration sources are deployed in the detector. To calibrate the timing and gain response of the photomultiplier tube (PMT) array regularly throughout the lifetime of the experiment, an external LED based calibration system has been developed known as the Timing component of the Embedded LED Light Injection Entity (TELLIE).

This thesis has two major parts: firstly the results of the testing and commissioning of the external calibration system, TELLIE, that I performed; and secondly the development of a new analysis is described that improves the sensitivity of the experiment to neutrinoless double-beta decay. My analysis does a comprehensive evaluation of systematic uncertainties and shows how their effect on SNO+ sensitivity can be reduced.

A more in depth description of the theory of neutrinos and motivation for and current results of searches in neutrinoless double-beta decay is given in Chapter 2. The chapter begins with a brief history of the key theoretical and experimental developments in the understanding of the neutrino. This is followed by a description of the physics of neutrino oscilliation. The theory of neutrino mass is then given along with the current experimental results of neutrino mass and the oscillation parameters. The chapter then goes on to describe the theory of Majorana neutrinos and the See-Saw Mechanism. Finally, a description of the physics of neutrinoless double-beta decay is given along with the experimental status of the field.

The SNO+ experiment is detailed in Chapter 3. The SNO+ detector is first described followed by the electronics and trigger system of SNO+. A detailed description of the calibration of SNO+ in terms of the types of calibration systems: the electronics calibration, optical calibration sources and radioactive calibration sources is given. This is followed by a review of all the radioactive backgrounds that SNO+ expects to detect. The chapter concludes with a description of the Monte Carlo simulation used by the SNO+ experiment.

Chapter 4 describes the LED based external calibration system and its commissioning. The requirements of PMT calibration are first detailed followed by a description of the hardware and software of the TELLIE system. There was an extensive period of quality assurance and characterisation testing before TELLIE was installed at SNO+ which is detailed. A review is given of the commissioning periods after TELLIE was installed where data was taken with the SNO+ detector. Finally a description of data quality checks is given which were designed to ensure that any data that is taken with TELLIE is of a high enough standard to be used for PMT calibration.

The methodology used in the analysis for sensitivity to neutrinoless double-beta decay is described in Chapter 5. The chapter begins with an outline of the likelihood technique used in the analysis and followed by a description of the parametrised Monte Carlo used to mimic the reconstruction of events in the SNO+ detector. Included in this section is a description of the motivation and how backgrounds which originate outside the liquid scintillator are handled by the analysis. The different minimisation techniques used in the analysis are also delineated.

Chapter 6 presents the results of the sensitivity studies on neutrinoless double-beta decay with systematic uncertainties of the experiment incorporated. The chapter begins with a description of the current state of the art sensitivity calculations. This is followed by details on how the energy range and binning was optimised for an analysis with and without systematic uncertainties included. The chapter then goes on to describe the results of analyses where: firstly individual systematic uncertainties are included; secondly pairs of systematic uncertainties are investigated. Finally two cases are compared: one where all systematic uncertainties considered are included; and then one where one systematic uncertainty is removed; in order to determine the reduction in neutrinoless double-beta decay half-life sensitivity when correlations with other systematic uncertainties are included. Also described is a new multi-dimensional fitting technique in energy and fractional Acrylic Vessel (AV) volume that improves the sensitivity of SNO+ to neutrinoless double-beta decay through self shielding from backgrounds originating outside the scintillator. This analysis is again first optimised with and without systematic uncertainties included, and then the reductions in sensitivity are examined for each systematic with all correlations accounted for. The chapter concludes with proposals for further studies which may be of interest to the SNO+ experiment.

Finally, Chapter 7 summarises the findings of this thesis. Here the main results are presented and the sensitivity SNO+ could achieve on the neutrinoless double-beta decay process is compared to the current results in the field.

Chapter 2

The Theoretical and Experimental Status of Neutrino Physics

This chapter details the history and physics of the neutrino that signifies the importance of the search for neutrinoless double-beta decay. Section 2.1 details the history of the neutrino. The physics of neutrino oscillation is described in Section 2.2. Section 2.3 explains the physics of neutrino mass and gives the current bounds on its value along with the oscillation parameter values. The theory of Majorana neutrinos is described in Section 2.4 and also details the see-saw mechanism that could explain the lightness of neutrino mass. The physics and current experimental results of neutrinoless double-beta decay is described in Section 2.5.

2.1 The Neutrino

The neutrino is a fundamental particle with neutral charge and a spin of one-half. They exist in the Standard Model as massless particles. However, due to the discovery of neutrino oscilliation, it was confirmed that neutrinos possess a non-zero mass which are several orders of magnitude below the other fundamental particles in the Standard Model. They exist in doublets with their corresponding charged leptons in three distinct neutrino flavours: electron ν_e , muon ν_{μ} and tau ν_{τ} . Neutrinos only interact via gravity due to their non-zero mass and the weak force. The weak force gauge bosons, W^{\pm} and Z^0 , only couple to negative chirality (left-handed) fermions and positive chirality (right-handed) anti-fermions. This implies that if right-handed neutrinos or left-handed anti-neutrinos exist then they will not interact via the weak force in the Standard Model and are hence known as sterile. The neutrino was first hypothesised by Pauli in 1930 in a letter to the "Radioactive Ladies and Gentlemen" of the Tübingen conference to explain the non-discrete energy spectrum of the electron observed in beta decay [4]. Pauli proposed an electrical neutral particle which he called the neutron that was emitted in addition to the electron in beta decay. This secondary particle would explain the continuous beta-decay energy spectrum as the neutral particle was carrying away energy unobserved in the experiment and hence would maintain the conservation of energy in the interaction. Following Chadwick's discovery in 1932 of the hadron that we know today as the neutron [5] and Perrin's observation in 1933 that the particle hypothesised by Pauli must have significantly less mass than the electron [6], Fermi coined the term neutrino, little neutral one, in his 1934 theory of beta decay [7].

The anti-neutrino was discovered in 1956 by Cowan and Reines [8]. Their experiment was designed to detect anti-neutrinos produced in a nuclear reactor located at the Savannah River in South Carolina. The anti-neutrinos interacted with protons within a tank of water via inverse beta decay, $\bar{\nu}_e + p \rightarrow n + e^+$. The positron produced in the interaction promptly annihilates with an electron producing gammas which were detected in tanks of liquid scintillator placed between the water tanks. The neutron produced interaction is then captured by cadmium, which were dissolved in the water, producing gammas 5 μ s later than the annihilation gammas. The delayed coincidence between the two sets of gammas allowed the inverse beta decay events to be distinguished from other background events and hence confirmed the existence of the anti-neutrino.

In 1957, the helicity of the neutrino, defined as the projection of a particle's spin along its momentum, was measured by Goldhaber et al [9]. The experiment used electron capture on a europium nucleus to produce an excited samarium nucleus and a neutrino. The excited samarium nucleus then returns to its ground state emitting a gamma ray. Through conservation of momentum and angular momentum, the helicity of the gamma rays emitted in the opposite direction of the neutrino is identical to that of the neutrino and thus the helicity of the neutrino can be determined through the measurement of the gamma's circular polarisation. The helicity of the neutrino was found to be negative (lefthanded) meaning its spin is anti aligned to its momentum. In the limit of massless particles helicity and chirality are identical. However, if a particle has mass then it is possible to boost to a reference frame which is moving faster than the particle's momentum where its helicity flips. Chirality, on the other hand, is a Lorentz invariant property of the particle.

In 1962, Lederman et al. discovered muon neutrinos, ν_{μ} , which were produced at

the Alternating Gradient Synchrotron and detected in a neon-filled detector known as a spark chamber [10]. Spark trails of muons produced from the interaction of neutrinos in aluminium plates confirmed the existence of ν_{μ} and confirmed the doublet nature of leptons.

The neutrino flux from the Sun was first measured by Davis et al. at the Homestake Experiment in 1968 [11]. The electron neutrinos from the Sun interacted with chlorine nuclei in the experiment which produce argon nuclei. The argon nuclei can then be extracted and counted to estimate the solar neutrino flux. The measured flux was approximately a third less than the amount predicted by theoretical models. This became known as the solar neutrino problem.

In 1976 the τ lepton was discovered by Perl et al. using an electron-positron collider and the MARK I detector [12]. This discovery implied the existence of a third neutrino to maintain the doublet structure of leptons; the tau neutrino, ν_{τ} .

Analogous to the solar neutrino flux, the neutrinos produced in interactions within the Earth's atmosphere, known as atmospheric neutrinos, were measured by the water Cerenkov experiments IMB [13] and Kamiokande [14] in 1988. A deficit was also found for this flux which became known as the atmospheric neutrino problem.

In 1989 the electron positron collider LEP became operational. The ALEPH, DELPHI, L3 and OPAL experiments based at the collider measured the total number of non-sterile neutrinos which have a mass below half the mass of the Z^0 boson from the line width of the Z^0 mass [15]. This was found to be consistent with 3 neutrinos.

The atmospheric neutrino problem was solved in 1998 by Super-Kamiokande, which measured an energy and zenith angle dependent deficit in the muon neutrino flux [16, 17]. This provided hints for neutrino oscillation and hence that neutrinos had mass.

The final neutrino predicted by the Standard Model, the tau neutrino ν_{τ} , was discovered in 2000 by the DONUT experiment [18]. This detected a ν_{τ} beam obtained via charmed meson decays that were produced at the Tevatron from accelerating protons, which collided with a fixed target.

The solar neutrino problem was solved by the SNO experiment in 2002 by measuring a deficit in the expected ν_e flux and no deficit in the total neutrino flux when compared to the expected ν_e flux [2]. This implies that a portion of the electron neutrinos produced in the Sun have changed flavour to ν_{μ} and ν_{τ} whilst propagating to the detector. We know the mechanism is neutrino oscillation. The fact that neutrinos change flavour implies that neutrinos must have a non-zero mass. From 2011 onwards, the mixing angle θ_{13} in the neutrino oscillation matrix was measured to be non-zero and relatively large in comparison to the analogous matrix in the quark sector by the reactor experiments Daya Bay, Double Chooz and RENO [19, 20, 21, 22] as well as the accelerator experiments T2K [23] and MINOS [24]. A non-zero value of this oscillation parameter is required for being able to observe CP violation in neutrino oscillation.

2.2 Neutrino Oscillations

Neutrinos are produced in weak eigenstates of definite lepton number $(|\nu_e\rangle, |\nu_{\mu}\rangle$ or $|\nu_{\tau}\rangle)$ referred to in this thesis as flavour eigenstates. A neutrino with momentum p is in an eigenstate with an energy eigenvalue of $E_i = \sqrt{m_i + p}$. There are three eigenstates each with a value of mass m_i , $i \in \{1, 2, 3\}$ known as mass eigenstates. The flavour eigenstates are related to the mass eigenstates via a mixing matrix U_{li}

$$\nu_l = \sum_{i=1}^3 U_{li} \nu_i, \quad l \in \{e, \mu, \tau\}.$$
(2.1)

 U_{li} is a 3×3, complex, unitary matrix known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. The results of the Super-Kamiokande and SNO experiments showed that neutrinos can change flavour as they propagate from the neutrino source to the detector. This is known as neutrino oscillation. Neutrino oscillation occurs due to the flavour eigenstates not being identical to their respective mass eigenstates and instead are a linear superposition of all neutrino mass eigenstates i.e. the off-diagonal elements of the PMNS matrix are non-zero. The PMNS matrix can be represented in the form of six parameters consisting of three mixing angles θ_{12} , θ_{13} , θ_{23} and three phases δ , ϕ_1 , ϕ_2 . A standard representation is then given by

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \cdot \operatorname{diag}(e^{i\phi_1}, e^{i\phi_2}, 1), \ (2.2)$$

where c_{ij} and s_{ij} notate $\cos \theta_{ij}$ and $\sin \theta_{ij}$. This first matrix is analogous to the Cabibbo-Kobayashi-Maskawa (CKM) matrix observed in the quark sector. The second matrix which contain the Majorana phases ϕ_1 and ϕ_2 arises due to how neutrino masses can be described differently to other fermions.

To calculate the probability of a neutrino of flavour α oscillating to flavour β then we

begin by assuming that at time t=0 a neutrino is created in a pure $|\nu_{\alpha}\rangle$ state

$$|\nu_{\alpha}(x=0,t=0)\rangle = \sum_{i=1}^{3} U_{\alpha i} |\nu_{i}\rangle,$$
 (2.3)

then after time t the neutrino would have propagated to a distance x = L

$$\left|\nu_{\alpha}(L,t)\right\rangle = \sum_{i=1}^{3} U_{\alpha i} \left|\nu_{i}\right\rangle e^{-i\mathbf{p_{i}}\cdot\mathbf{x}},\tag{2.4}$$

where \mathbf{p}_i and \mathbf{x} are the momentum and spatial 4-vectors. If it is assumed that the neutrino is ultra relativistic which is reasonable given their light mass, m_i , in comparison to their energy, E, then given

$$E_i = \sqrt{p^2 + m_i^2} \approx |p| + \frac{m_i^2}{2|p|},$$
(2.5)

and since $|p| \approx E$ then

$$\mathbf{p_i} \cdot \mathbf{x} = E_i t - |p|L \approx (E_i - |p|)L \approx \frac{m_i^2}{2E}L.$$
(2.6)

Rewriting the mass eigenstates in equation 2.4 in terms of their flavour eigenstates

$$|\nu_{\alpha}(L)\rangle = \sum_{i=1}^{3} (U_{\alpha i} U_{ei}^{*} e^{-i\frac{m_{i}^{2}}{2E}L} |\nu_{e}\rangle + U_{\alpha i} U_{\mu i}^{*} e^{-i\frac{m_{i}^{2}}{2E}L} |\nu_{\mu}\rangle + U_{\alpha i} U_{\tau i}^{*} e^{-i\frac{m_{i}^{2}}{2E}L} |\nu_{\tau}\rangle).$$
(2.7)

Using the orthonormality condition $\langle \nu_{\alpha} | \nu_{\beta} \rangle = \delta_{\alpha\beta}$, the probability of a neutrino ν_{α} oscillating to a flavour of ν_{β} in a vacuum is obtained to be

$$P(\alpha \to \beta) = |\langle \nu_{\beta} | \nu_{\alpha}(L) \rangle|^{2} = \delta_{\alpha\beta} - 4 \sum_{i < j}^{1,3} \operatorname{Re} \left[U_{\alpha i} U_{\beta i}^{*} U_{\alpha j}^{*} U_{\beta j} \right] \sin^{2} \left(\Delta m_{ij}^{2} \frac{L}{4E} \right) + 2 \sum_{i < j}^{1,3} \operatorname{Im} \left[U_{\alpha i} U_{\beta i}^{*} U_{\alpha j}^{*} U_{\beta j} \right] \sin \left(\Delta m_{ij}^{2} \frac{L}{2E} \right), \quad (2.8)$$

where $\Delta m_{ij} = m_i^2 - m_j^2$. So that more practical units are used for neutrino oscillation experiments, the terms inside the sinusoids are often written such that

$$\Delta m_{ij} \frac{L}{4E} \approx 1.27 \Delta m_{ij}^2 [\text{eV}^2] \frac{L[\text{km}]}{E[\text{GeV}]}.$$
(2.9)

A consequence of the mass splitting in the probability of oscillating from one flavour to another is that, as neutrinos have been observed to have oscillated, neutrinos must have a non-degenerate and hence non-zero mass which is not predicted in the Standard Model. The maximum sensitivity to the mass splittings occur when $2\pi E/L = \Delta m_{ij}$. The imaginary component in equation 2.8 has opposite signs for neutrinos and anti-neutrinos making it sensitive to CP violation.

The PMNS matrix in equation 2.2 can be expanded to four rotation matrices such that the effects of each of the mixing angles can be seen explicitly

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \cdot \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \cdot \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \operatorname{diag}(e^{i\phi_1}, e^{i\phi_2}, 1).$$

$$(2.10)$$

From left to right, the mixing due to the first matrix is observed in atmospheric and long baseline neutrino oscillation experiments. The effects of the second matrix are tested by reactor and long baseline neutrino oscillation experiments that have a different baseline to those which observe the effect of neutrino oscillations. These long baseline experiments are also sensitive to the CP violating term δ . The third matrix is observed in reactor and solar neutrino experiments. Neutrino oscillation experiments are not sensitive to the Majorana phases in the fourth matrix due to neutrino oscillation being a lepton number conserving process whilst the Majorana phases arise when a Majorana mass term is included which violates lepton number conservation. These phases cancel out when calculating the neutrino oscillation probabilities. The only known way of probing these phases is through neutrinoless double-beta decay experiments.

2.3 Mass of Neutrinos

Mass occurs in the Standard Model Lagrangian via the Dirac mass term

$$m\bar{\psi}\psi = m\bar{\psi}_L\psi_R + m\bar{\psi}_R\psi_L,\tag{2.11}$$

where the Dirac spinor ψ has been decomposed into its left and right chiral states and the fact that $\bar{\psi}_L \psi_L = \bar{\psi}_R \psi_R = 0$ has been used. This states that a particle must have left and right-handed chiral states to have mass in the Standard Model. This means that either neutrinos are massless or there exist right-handed sterile neutrinos that do not couple to the weak interaction which are present in the Standard Model to give neutrinos mass. A third option is that there is physics beyond the Standard Model, which gives rise to neutrino mass.

Neutrino oscillation experiments have allowed a measurement of the difference between the three neutrino mass states but not the absolute scale of them, as discussed in the next section. The experimental probes to the absolute scale are through beta decay experiments which are sensitive to the ν_e mass, observations of the Cosmic Microwave Background that provide an upper limit on the sum of the neutrino masses [25] and neutrinoless double-beta decay experiments that are sensitive to the effective Majorana mass, which is described in Section 2.5.

2.3.1 Mass Splittings

The differences between the mass states have been measured by neutrino oscillation experiments. These differences are much smaller than the absolute values of the masses. The solar neutrino experiments measure the mass splitting between m_1 and m_2 . The Mikheyev–Smirnov–Wolfenstein (MSW) effect takes into account the fact that ν_e interact via charged and neutral currents in the Sun due to the Sun's electron density, whereas other flavours only interact via neutral currents. Due to the MSW effect, from neutrinos propagating in the Sun we know the sign of Δm_{12}^2

$$\Delta m_{12}^2 = m_2^2 - m_1^2 > 0. \tag{2.12}$$

The other neutrino oscillation experiments that have been performed are approximate to neutrinos oscillating in a vacuum and hence the MSW effect is not observed and the sign of Δm_{13}^2 and Δm_{23}^2 is therefore unknown. This leads to the concept of the normal (NH) and inverted (IH) hierarchies of neutrino mass. In the normal hierarchy, the order of the neutrino mass is such that $m_3 > m_2 > m_1$, whereas in the inverted hierarchy $m_2 > m_1 > m_3$. To avoid complications due to the two different hierarchies, the mass difference between m_3^2 and $m_{1,2}^2$ can be reported as an average splitting

$$\Delta m^2 = m_3^2 - \frac{m_1^2 + m_2^2}{2}.$$
(2.13)

A positive value of Δm^2 relates to the NH and a negative value to the IH.

Figure 2.1 shows the mass splitting for the two hierarchies. It also shows the probabilities of each flavour state for each mass state. Note that Δm_{sol}^2 is equivalent to Δm_{12}^2 and Δm_{atm}^2 is approximately equivalent to Δm^2 .

There is also the possibility that the hierarchies are quasi-degenerate (QD). For the normal hierarchy this would be equivalent to stating $m_1 \leq m_2 \leq m_3 \simeq m_{\nu}^{QD}$ and for the



Figure 2.1: A pictorial representation of the normal hierarchy, left, and inverted hierarchy, right in terms of their mass states and their flavour states [26].

inverted hierarchy $m_3 \lesssim m_1 \lesssim m_2 \simeq m_{\nu}^{QD}$.

2.3.2 Oscillation Parameters

As explained in the previous section, the neutrino masses cannot be probed directly using oscillation experiments but their mass differences can. Table 2.1 shows the most recent best fit values and their associated uncertainty of all oscillation parameters. These values are from a global fit which was published in January 2016 [27]. The uncertainties on the mass splittings are too large for a hierarchy to be determined.

2.3.3 Particle Decay Constraints

The electron neutrino mass can be probed directly through measurements of beta decay. These experiments attempt to measure the electron neutrino mass by measuring the end point of tritium beta decay spectrum. Through conservation of energy and momentum the electron neutrino mass can be determined through the position of this end point. The Mainz [28] and Troitsk [29] experiments have the best 95% Confidence Limit (CL) on direct measurement of electron neutrino mass of

$$m_{\nu_e} = \sqrt{\sum_i |U_{ei}|^2 m_i^2} < 2 \text{ eV}.$$
 (2.14)

Table 2.1: Results of a global three neutrino oscillation analysis from January 2016 [27]. Best fit values are given for each oscillation parameter with one, two and three σ ranges for each hierarchy. The value for Δm^2 is positive for normal hierarchy (NH) and negative for inverted hierarchy (IH).

Parameter	Hierarchy	Best Fit	1σ range	2σ range	3σ range
$\Delta m_{12}^2 / 10^{-5} \text{ eV}^2$	NH or IH	7.37	$7.21 \rightarrow 7.54$	$7.07 \rightarrow 7.73$	$6.93 \rightarrow 7.97$
$\Delta m^2 \ / \ 10^{-3} \ \mathrm{eV^2}$	NH	2.50	$2.46 \rightarrow 2.54$	$2.41 \rightarrow 2.58$	$2.37 \rightarrow 2.63$
	IH	2.46	$2.42 \rightarrow 2.51$	$2.38 \rightarrow 2.55$	$2.33 \rightarrow 2.60$
$\sin^2 \theta_{12} / 10^{-1}$	NH or IH	2.97	$2.81 \rightarrow 3.14$	$2.65 \rightarrow 3.34$	$2.50 \rightarrow 3.54$
$\sin^2 \theta_{13} / 10^{-2}$	NH	2.14	$2.05 \rightarrow 2.25$	$1.95 \rightarrow 2.36$	$1.85 \rightarrow 2.46$
	IH	2.18	$2.06 \rightarrow 2.27$	$1.96 \rightarrow 2.38$	$1.86 \rightarrow 2.48$
$\sin^2\theta_{23} \ / \ 10^{-1}$	NH	4.37	$4.17 \rightarrow 4.70$	$3.97 \rightarrow 5.63$	$3.79 \rightarrow 6.16$
	IH	5.69	$4.28 { ightarrow} 4.91$	4.04→6.18	$3.83 \rightarrow 6.37$
			$5.18 { ightarrow} 5.97$		
δ / π	NH	1.35	$1.13 \rightarrow 1.64$	$0.92 \rightarrow 1.99$	$0 \rightarrow 2$
	IH	1.32	$1.07 \rightarrow 1.67$	$0.83 \rightarrow 1.99$	$0 \rightarrow 2$

The KATRIN experiment is expected to start taking tritium data in 2017. It is expected to have a sensitivity of $m_{\nu_e} < 0.2 \text{ eV}$ [30] at 95% CL. Figure 2.2 shows where these limits occur for each hierarchy as a function of the lightest mass eigenstate.



Figure 2.2: The mass of the electron neutrino as a function of the lightest neutrino state for normal, left, and inverted, right, hierarchies. The limits for Mainz and Troitsk are represented by straight lines as well as the expected limit for KATRIN. The values used to produce the curves are $\Delta m_{12}^2 = 7.53 \times 10^{-5} \text{ eV}^2$, $\Delta m^2 = 2.44 \times 10^{-3} \text{ eV}^2$, $|U_{e1}|^2 = 0.6803$, $|U_{e2}|^2 = 0.2354$ and $|U_{e3}|^2 = 0.0843$ [30].

The muon neutrino mass states can be probed directly through the measurement of

pion decay at rest $\pi^+ \to \mu^+ + \nu_\mu$ or $\pi^- \to \mu^- + \bar{\nu_\mu}$. The current 95% CL is [30]

$$m_{\nu\mu} = \sum_{i} |U_{\mu i}|^2 m_i^2 < 190 \text{ keV.}$$
 (2.15)

The tau neutrino mass state is probed directly through the measurement of τ^- decay kinematics. The current 95% CL is [30]

$$m_{\nu_{\tau}} = \sqrt{\sum_{i} |U_{\tau i}|^2 m_i^2} < 18.2 \text{ MeV.}$$
 (2.16)

2.3.4 Cosmological Constraints

As neutrinos are the second most abundant particle in the universe, the first being photons, a neutrino with mass will have an effect on the structure of the universe. The structure of the universe can be probed through the Cosmic Microwave Background (CMB) and Baryon Acoustic Oscillations (BAO). The CMB is radiation left over after recombination following the big bang. BAO are fluctuations in the density of visible baryonic matter.

The Planck Collaboration obtained constraints on the sum of the three neutrino mass states $\sum_{i} m_{i}$ [31] by combining their results from CMB data with the CMB data from the Wilkinson Microwave Anisotropy Probe (WMAP) [25] and the Atacama Cosmology Telescope (ACT) [32] experiments. They obtained a limit of $\sum_{i} m_{i} < 0.66$ eV at 95% CL [31]. This limit can be further improved by adding BAO data [31]

$$\sum_{i} m_i < 0.23 \text{ eV.}$$
(2.17)

2.4 Majorana Neutrinos

Observations of neutrino oscillation and constraints on neutrino mass demonstrate that neutrinos have a mass that is several orders of magnitude below the other fermions in the Standard Model. As described in Section 2.3, neutrinos in the Standard Model are massless as the Dirac mass term requires there to be both right and left-handed fermions and only left-handed neutrinos are observed in nature. One solution is to introduce a sterile right-handed neutrino that does not couple to the weak interaction. It is possible to detect the presence of sterile neutrinos in nature through neutrino oscillation since adding more mass and flavour eigenstates will affect the neutrino oscillation probabilities. There has been no observation of the effect of sterile neutrinos at present. Sterile neutrinos would also not explain why the neutrino mass is so suppressed. An alternative theory was proposed by Majorana in 1937 who wanted to describe a neutrino with mass using just a left-handed field. This theory described Majorana fields

$$\psi = \psi_L + \psi_R = \psi_L + \psi_L^C, \qquad (2.18)$$

where the L and R denote the left and right-handed fields and C denotes the charge conjugate. This means that a fermion described by a Majorana field is its own antiparticle and thus must be electrically neutral. The only fundamental particle that can be described by a Majorana field is the neutrino. The result of Majorana's theory led to a Majorana mass term in the lagrangian

$$L_{L}^{M} = -\frac{1}{2}m_{L}\nu_{L}^{\bar{C}}\nu_{L}, \qquad (2.19)$$

where m_L is the Majorana mass. The Majorana mass term allows mixing between neutrino (lepton number = +1) and antineutrino (lepton number = -1) states and hence does not conserve lepton number. Lepton number conservation is required by some cosmological models to meet the Sakharov conditions which led to the matter-antimatter asymmetry observed in the universe.

Unlike the Dirac mass term, $L_D = -m\bar{\nu}_R\nu_L$, the Majorana mass does not arise from Yukawa couplings to the Higgs Field in the Standard Model and hence $m_L = 0$ or it must be described by physics beyond the Standard Model. This description therefore does not negate the need for right-handed neutrinos in the Standard Model to describe neutrino mass. However, it could explain the suppression of neutrino mass through the see-saw mechanism.

2.4.1 See-Saw Mechanism

Although the left-handed Majorana mass is not allowed under Standard Model symmetries, the right-handed Majorana mass is. If we begin by assuming m_L is non-zero, we can write the neutrino mass term in the Standard model as follows

$$-\mathcal{L}_{mass}^{\nu} = \frac{1}{2} \left[\begin{pmatrix} \bar{\nu}_L^c & \bar{\nu}_R \end{pmatrix} \begin{pmatrix} m_L & m \\ m & m_R \end{pmatrix} \begin{pmatrix} \nu_L \\ \nu_R^c \end{pmatrix} \right] + h.c., \qquad (2.20)$$

where m is the Dirac mass. The mass matrix can then be diagonalised to give the two mass eigenvalues

$$m_{\alpha,\beta} = \frac{1}{2} \left[(m_L + m_R) \pm \sqrt{(m_L - m_R)^2 + 4m^2} \right].$$
(2.21)

If it is assumed that $m_L = m_R = 0$, then the Dirac mass is recovered: $m_{\alpha,\beta} = \pm m$. Conversely, if it is assumed that m = 0 then the Majorana masses are recovered $m_{\alpha} = m_L$, $m_{\beta} = m_R$. The see-saw limit occurs when it is assumed that $m_r = M >> m$, where Mis a heavy Majorana mass state, and, through Standard Model constraints, that $m_L = 0$. In this case the approximation

$$m_{\alpha,\beta} = \frac{M}{2} \left[1 \pm \sqrt{1 + 4\frac{m^2}{M^2}} \right] \approx \frac{M}{2} \left[1 \pm \left(1 + 2\frac{m^2}{M^2} \right) \right]$$
(2.22)

can be made leading to the mass eigenvalues

$$m_{\alpha} = M + \frac{m^2}{M} \approx M$$
 and $m_{\beta} = \frac{m^2}{M}$. (2.23)

where the negative sign of m_{β} has been absorbed by a phase. Here the eigenvalue, m_{α} , is very large and the eigenvalue, m_{β} , is very small due to the suppression from the heavy Majorana mass M. The geometric mean of m_{α} and m_{β} is |m| since $m_{\alpha}m_{\beta} = m^2$ leading to the see-saw effect as if m_{α} is large then m_{β} is small. This therefore provides a possible explanation for the light left-handed neutrino mass if there exists a heavy righthanded neutrino. For example, if the Dirac mass m = 1 MeV, which is of the order of the lightest charged lepton mass, then the heavy Majorana mass is of the order of 10^{15} eV. This would lead to effective neutrino masses in the meV range, which fits with current constraints. These heavy right-handed neutrinos may have existed in the earlier universe and, as the universe cooled, would have decayed to left-handed neutrinos and right-handed anti-neutrinos. If these decay rates to neutrinos and anti-neutrinos are slightly different this could have led to the matter-antimatter asymmetry observed in the universe today, which is one reason why determining whether or not neutrinos are Majorana is an important question in physics.

2.5 Neutrinoless Double-Beta Decay

The only known way to test whether neutrinos are Majorana or not is through the observation of neutrinoless double-beta decay $(0\nu\beta\beta)$. Double beta decay occurs in isotopes where single beta decay is energetically forbidden or it is highly suppressed due to a large change in the isotope's spin. In double beta decay two protons simultaneously change to two neutrons. From the decaying nucleus, two electrons and two electron antineutrinos are emitted. In neutrinoless double-beta decay, one of the $\bar{\nu}_e$ undergoes a helicity flip trans-



Figure 2.3: A feynman diagram of a process which may be responsible for neutrinoless double-beta decay where a Majorana neutrino is exchanged between two W^- bosons.

forming it to a ν_e . This can only occur if the neutrino has a non-zero Majorana mass. The $\bar{\nu_e}$ or ν_e then absorbs the other such that only the two electrons are emitted from the nucleus. This process is highly suppressed in comparison to the standard double beta decay emission due to the small neutrino mass and the energy required for the helicity flip. The process also violates lepton number by $\Delta L = 2$ due to the only leptons in the final state being two electrons.

The standard mechanism used to describe this process involves the exchange of a light neutrino between two W^- bosons as represented in the feynman diagram in Figure 2.3. The decay rate for this process is given by

$$(T_{1/2}^{0\nu})^{-1} = G^{0\nu} |M^{0\nu}|^2 \frac{\langle m_{\beta\beta} \rangle^2}{m_e^2}, \quad \langle m_{\beta\beta} \rangle = \left| \sum_{i=1}^3 U_{ei}^2 m_i \right|, \tag{2.24}$$

where $G^{0\nu}$ is the phase space factor which is dependent on the atomic number and the Q-value of the transition, $M^{0\nu}$ is the nuclear matrix element (NME) and m_e is the rest mass of the electron. The effective Majorana mass $\langle m_{\beta\beta} \rangle$ is dependent on the electron neutrino elements of the PMNS matrix U_{ei} and the neutrino masses m_i .

There are 35 isotopes in which double beta decay is observable. However, as the $G^{0\nu}$ is proportional to Q^5 , only the eleven isotopes with a Q-value above 2 MeV are considered in experimental searches. These isotopes are listed in Table 2.2.

There are non standard mechanisms that can describe $0\nu\beta\beta$ [33], when heavy neutrinos are responsible for giving neutrinos their light mass via the see-saw mechanism or through the exchange of super-symmetric particles. In the standard interpretation, the contribution



Figure 2.4: Values of the nuclear matrix element, $M^{0\nu}$, for several isotopes using different nuclear physics models. These are the Nuclear Shell Model (NSM); Quasi particle Random Phase Approximation (QRPA); Interacting Boson Model (IBM); Energy Density Functional (EDF) and the Projected Hartree-Fock-Bogoliubov model (PHFB). The error bars are due to theoretical uncertainties arising from varying model parameters. [35]

of these mechanisms to $0\nu\beta\beta$ is assumed to be negligible in comparison to the standard mechanism [34]. Non standard mechanisms also still require a Majorana mass term.

The value of $M^{0\nu}$ cannot be measured and has to be calculated theoretically. There are five different nuclear physics models which are used to calculate a value for $M^{0\nu}$ in the standard mechanism of $0\nu\beta\beta$. These are the Energy Density Functional (EDF) [36]

Table 2.2: The phase space factor $G^{0\nu}$ for the standard mechanism of neutrinoless doublebeta decay; *Q*-value and natural abundance for isotopes where double beta decay is observable and have a *Q*-value >2 MeV.

Isotope	$G^{0\nu}$	Q	nat. abund.
	$(10^{-14} \text{ yrs}^{-1})$	(keV)	(%)
⁴⁸ Ca	6.35	4273.7	0.187
$^{76}\mathrm{Ge}$	0.623	2039.1	7.8
82 Se	2.70	2995.5	9.2
96 Zr	5.63	3347.7	2.8
^{100}Mo	4.36	3035.0	9.6
¹¹⁰ Pd	1.40	2004.0	11.8
$^{116}\mathrm{Cd}$	4.62	2809.1	7.6
124 Sn	2.55	2287.7	5.6
¹³⁰ Te	4.09	2530.3	34.5
¹³⁶ Xe	4.31	2461.9	8.9
¹⁵⁰ Nd	19.2	3367.3	5.6



Figure 2.5: An idealised energy spectrum of standard double beta decay in black and neutrinoless double-beta decay in red normalised to the Q-value of the decay, $Q_{\beta\beta}$ [42].

also known as Generating Coordinate Method (GCM) [37], the Interacting Boson Model (IBM) [38], the Nuclear Shell Model (NSM) [39], the projected Hartree-Fock-Bogoliubov model (PHFB) [40] and the Quasi-particle Random Phase Approximation (QRPA) [41]. The calculations are complicated many body problems which describe the overlap of nuclear wave functions of the isotopes initial and final states. The small modifications to the Hamiltonians in the nuclear models lead to large changes in the wavefunctions, whilst the energy levels remain relatively stable [35]. This means that there are large theoretical uncertainties on the values of $M^{0\nu}$. Due to these theoretical uncertainties and the multiple methods used in determining $M^{0\nu}$, there are several possible values for $M^{0\nu}$ for each isotope as shown in Figure 2.4. Therefore experiments must use a variety of isotopes to factor out these theoretical uncertainties and distinguish between the different models in order to determine a value of $\langle m_{\beta\beta} \rangle$. However, any observation of neutrinoless double-beta decay would prove that neutrinos are Majorana particles and lepton number is violated.

The energy spectrum of $0\nu\beta\beta$ is a delta function at the Q-value of the decay of the isotope. This is due to the nuclear recoil of the isotope being negligible and all the available energy is transferred to the electrons whose energy deposition is detected. The energy spectrum of standard double beta decay is similar to that of single beta decay where the emitted neutrinos carry away some of the energy of the decay undetected. Figure 2.5 shows the energy spectra of the processes in an idealised case. In reality the $0\nu\beta\beta$ and

two neutrino spectra will overlap more due to a non perfect energy resolution broadening the two distributions.

In the case where an experiment has no background, the half life of $0\nu\beta\beta$ is given by the radioactive decay law

$$T_{1/2}^{0\nu} = \ln(2)aMtN_A(N_{\beta\beta})^{-1}, \qquad (2.25)$$

where a is the isotopes abundance, M is the mass of the isotope, t is the time of the measurement, N_A is Avogadro's number and $N_{\beta\beta}$ is the number of observed $0\nu\beta\beta$ events. In the background limiting case where the number background counts B is of the level $N_{\beta\beta} \approx \sqrt{B}$, the half life can be written as

$$T_{1/2}^{0\nu} \propto \frac{\ln(2)aMtN_A}{\sqrt{B}}.$$
 (2.26)

Backgrounds can split into two types: the backgrounds which are independent of the mass of the isotope and backgrounds which are not. An example of a background which is indepdent of M is ⁸B solar neutrinos as the source is from the Sun, not the isotope. The count rate of these backgrounds is proportional to energy resolution ΔE and t: $B \propto t \Delta E$. Removing the fixed constants, the half life sensitivity proportionality in the case where only these backgrounds are detected can then be written as

$$T_{1/2}^{0\nu} \propto \frac{M\sqrt{t}}{\sqrt{\Delta E}}.$$
(2.27)

Backgrounds which are dependent on M such as the two neutrino double beta decay have a background count proportionality of $B \propto Mt\Delta E$. So the half life sensitivity proportionality where only these backgrounds are detected can be written as

$$T_{1/2}^{0\nu} \propto \frac{\sqrt{Mt}}{\sqrt{\Delta E}} \tag{2.28}$$

One can then use equations 2.24 and 2.25 to examine how $\langle m_{\beta\beta} \rangle$ is related to M. In the case where background counts are proportional to isotope mass the sensitivity to $\langle m_{\beta\beta} \rangle$ for a fixed live time and energy resolution follows the proportionality

$$\langle m_{\beta\beta} \rangle \propto M^{1/4},$$
 (2.29)

whereas for the case where background counts are independent of isotope it is found that

$$\langle m_{\beta\beta} \rangle \propto M^{1/2},$$
 (2.30)



Figure 2.6: The effective Majorana mass $\langle m_{\beta\beta} \rangle$ as a function of the lightest neutrino mass $m_{lightest}$. Also shown is the 90% confidence limit from the KamLAND-Zen experiment with ¹³⁶Xe and the corresponding confidence limits with other nuclei [43].

which will be closer to the scenario for SNO+ as the dominant background for this experiment is from ⁸B solar neutrinos.

With the effective Majorana mass expressed in terms of oscillation parameters [35]

$$\langle m_{\beta\beta} \rangle = \left| c_{13}^2 c_{12}^2 e^{2i\alpha_1} m_1 + c_{13}^2 s_{12}^2 e^{2i\alpha_2} m_2 + s_{13}^2 m_3 \right|, \qquad (2.31)$$

where $\alpha i = \phi_i + \delta$, one can use the values for the oscillation parameters to map out the value of $\langle m_{\beta\beta} \rangle$ in terms of the lightest neutrino mass, m_1 in the case of normal hierarchy and m_3 for inverted hierarchy, as shown in Figure 2.6.

As the phases α_1 and α_2 are completely unknown, there are large bands in uncertainties for the two hierarchies. These phases can cause cancellations to the contribution to $\langle m_{\beta\beta} \rangle$ in equation 2.31. There can be complete cancellation for the normal hierarchy, if the value of m_1 is between 2 and 7 meV, but the inverted hierarchy is constrained to values larger than approximately 20 meV for all values of m_3 .

If neutrinoless double-beta decay is not observed with a value of $\langle m_{\beta\beta} \rangle < 20$ meV then this does not rule out inverted hierarchy however if it is observed with $\langle m_{\beta\beta} \rangle < 20$ meV then the hierarchy must be normal. If neutrinoless double-beta decay is observed with



Figure 2.7: Left: The effective Majorana mass $|m_{\beta\beta}|$ as a function of the effective neutrino mass m_{β} . Also shown is the expected 95% confidence limit from the KATRIN experiment on m_{β} and a mock 90% double beta decay experiment limit. Right: The effective Majorana mass $|m_{\beta\beta}|$ as a function of the sum of the three neutrino masses $\sum_k m_k$. Also shown is the current 95% confidence limit from cosmological experiments on m_{β} and a mock 90% double beta decay experiment limit. Taken from [35]

 $\langle m_{\beta\beta} \rangle > 20$ meV then the hierarchies are not distinguishable with the current cosmological bounds.

As the value of m_{min} is difficult to determine, it can be useful to plot $\langle m_{\beta\beta} \rangle$ in terms of the effective neutrino mass as measured by beta decay experiments as shown in Figure 2.7a. Here if the effective neutrino mass is measured to be above 30 meV or if $\langle m_{\beta\beta} \rangle$ is measured to be above 10 meV then it is difficult to distinguish between the two hierarchies.

One can also plot $\langle m_{\beta\beta} \rangle$ in terms of the sum of the three neutrino masses $\sum_i m_i$ as measured by cosmological experiments, shown in Figure 2.7b. Again if $\langle m_{\beta\beta} \rangle$ is measured to be above 10 meV or if the sum of masses is above 80 meV then distinguishing the two hierarchies is again problematic.

2.5.1 Double Beta Experiments

There are multiple techniques available for searching for neutrinoless double-beta decay. This section outlines current and future experiments arranged by their different techniques.

Calorimetry and Tracking - NEMO-3 and SuperNEMO

Located in the Fréjus Underground laboratory, the NEMO-3 experiment ran for eight years collecting five years worth of data. Cylindrical in structure, the technique uses plastic scin-

tillator blocks and gaseous drift chambers placed either side of a foil containing a candidate isotope. The drift chambers can track the particle's trajectory and the scintillator blocks allow a measure of the particle's energy. The energy deposited in the scintillator and angular distribution of the particles, acquired from the tracking in the drift chambers, is then used to distinguish neutrinoless double-beta decay events from other background events. The isotope foils are modular which allowed for multiple isotopes to be deployed. The isotopes deployed in NEMO-3 were ⁴⁸Ca [44], ⁸²Se [45], ⁹⁶Zr [46], ¹⁰⁰Mo [47], ¹¹⁶Cd [48], ¹³⁰Te [49] and ¹⁵⁰Nd [50]. The primary search for NEMO-3 was ¹⁰⁰Mo and with a 4.96 yr livetime and a mass of 6.914 kg. A lower limit on half life of $T_{1/2}^{0\nu\beta\beta} > 1.1 \times 10^{24}$ was obtained at the 90% CL which corresponds to a value of $\langle m_{\beta\beta} \rangle < 0.33 - 0.62$ eV [51].

The upgrade to NEMO-3 is SuperNEMO and uses the same technique as previously discussed with greater isotope mass, lower background levels and improved tracking and calorimetry. It will consist of 20 foil modules which will house a total of 100 kg of ⁸²Se. With 500 kg·yr exposure, SuperNEMO is expected to reach a sensitivity of $T_{1/2}^{0\nu\beta\beta} > 1 \times 10^{26}$ at the 90% CL which corresponds to a value of $\langle m_{\beta\beta} \rangle < 0.04 - 0.1$ eV [52].

Cryostat Bolometers - CUORICINO and CUORE

Located at the Gran Sasso Underground Laboratory, CUORCINO uses bolometers as its technique to search for neutrinoless double-beta decay. Bolometers measure temperature change in materials. They can be sensitive enough to measure the heat change caused by radioactive decays such as double beta decay.

CUORICINO used bolometers consisting of TeO₂ crystals. A total of 62 crystals was used. All except 4 of these crystals used natural Te, two were enriched to a purity of 75% ¹³⁰Te and two to a purity of 82.3% ¹³⁰Te. A total of 11.6 kg of ¹³⁰Te was used in the experiment and an exposure of 19.6 kg·yr was achieved. The crystals were cooled to a temperature of 8-10 mK. CUORICINO achieved an energy resolution of 7 keV at the ¹³⁰Te Q-value. A lower limit on the half life of $T_{1/2}^{0\nu\beta\beta} > 2.9 \times 10^{24}$ was obtained at the 90% CL which corresponds to a value of $\langle m_{\beta\beta} \rangle < 0.3 - 0.71$ eV [53].

CUORE is the upgrade to CUORICINO and employs the same technique and will consist of 988 TeO₂ crystals and have a total mass of approximately 210 kg of ¹³⁰Te. The experiment aims to achieve a half life sensitivity of $T_{1/2}^{0\nu\beta\beta} > 6.5 \times 10^{26}$ 90% CL corresponding to a value of $\langle m_{\beta\beta} \rangle < 26 - 40$ meV [54]. One tower, consisting of 52 TeO₂ crystals, of CUORE has already been used to take 9.8 kg·yr worth of data. When the results of this are combined with the CUORICINO data then a sensitivity of $T_{1/2}^{0\nu\beta\beta} >$ 4.0×10^{24} 90% CL corresponding to a value of $\langle m_{\beta\beta} \rangle < 0.27 - 0.76$ eV [55] is achieved.

High Purity Germanium Detectors - GERDA and MAJORANA

With High Purity Germanium detectors (HPGe), the source (⁷⁶Ge) is also the detector. Germanium is a semiconductor which can be doped such that a p-n junction is formed to allow recombination of electron-hole pairs. The nature of recombination varies by the type of incident radiation which excites the electrons. Pulse shape discrimination can then be used on the electrical readout of the HPGe to identify different particle types. A HPGe also has an energy resolution of approximately 2-5 keV [56], [57] which makes them very appealing for the search of neutrinoless double-beta decay.

Located at the Gran Sasso Underground Laboratory, GERDA reused a combination of enriched and natural HPGe from the ANG, Heidelberg-Moscow and IGEX experiments. The first phase of GERDA had a total of 17.7 kg of ⁷⁶Ge and an exposure of 21.6 kg·yr. A lower limit of $T_{1/2}^{0\nu\beta\beta} > 2.1 \times 10^{25}$ 90% CL was achieved which corresponds to a value of $\langle m_{\beta\beta} \rangle < 0.2 - 0.4$ eV [58]. The upgrade for the second phase of GERDA is currently in progress. It will aim to take an exposure of 100 kg·yr and aims to reach a sensitivity of larger than $T_{1/2}^{0\nu\beta\beta} > 10^{26}$ 90% CL [59].

MAJORANA is currently constructing a demonstrator module in the Sanford Underground Research Facility. It employs a similar technique to GERDA and is expected to run for an exposure of 100 kg·yr. With this exposure it will achieve a sensity of approximately $T_{1/2}^{0\nu\beta\beta} > 10^{26} 90\%$ CL which will translate to a sensitive of below $\langle m_{\beta\beta} \rangle < 0.1$ eV. MAJOR-ANA also plans a next generation experiment which will have an exposure of 1000 kg·yr. With this exposure it will achieve a sensity of approximately $T_{1/2}^{0\nu\beta\beta} > 10^{27} 90\%$ CL which will translate to a sensitive of approximately $\langle m_{\beta\beta} \rangle < 20$ meV [60].

Time Projection Chambers - EXO-200, nEXO and NEXT

A Time Projection Chamber (TPC) uses electric and magnetic fields to accelerate particles across the chamber which contains liquid or gas. The technique allows three dimensional track reconstruction and can achieve good position resolution.

EXO-200 is located at the Waste Isolation Pilot Plant in New Mexico, USA. The cylindrical TPC is filled with liquid Xe enriched to $80.6 \pm 0.1\%$ ¹³⁶Xe. The TPC is divided into two symmetric volumes seperated by a cathode grid. At each end of the TPC are 38 charge induction and 38 charge collection wire triplets which cross at 60°. The ends of the TPC each have approximately 250 Large Area Avalanche Photodiodes (LAAPDs) which

collect scintillation light that provides energy information of the interactions [61]. The energy resolution is approximately 1.5% at the Q-value of ¹³⁶Xe. Tracks approximately 1 cm apart can be distinguished and a position resolution of a few mm has been achieved. EXO-200 had an exposure of 100 kg·yr and achieved a lower limit of $T_{1/2}^{0\nu\beta\beta} > 1.1 \times$ 10^{25} 90% CL which corresponds to a value of $\langle m_{\beta\beta} \rangle < 0.19 - 0.45$ eV [62].

R&D has begun on the upgrade to EXO-200, nEXO. The experiment is planning to use 5 tonnes of enriched Xe and, after 5 years of running, plans to achieve a lower limit of $T_{1/2}^{0\nu\beta\beta} > 6.6 \times 10^{27}$ 90% CL which corresponds to a value of $\langle m_{\beta\beta} \rangle < 7 - 18$ meV [63].

NEXT is an experiment that employs a similar technology to EXO-200 and nEXO. The NEXT TPC will be placed under high pressure (15 bar) which will allow a greater source mass per unit volume. It is planning to have a source mass of 100 kg of enriched Xe. After an exposure of 300 kg·yr it is expected to achieve a lower limit of $T_{1/2}^{0\nu\beta\beta} > 6 \times 10^{25}$ 90% CL which corresponds to a value of $\langle m_{\beta\beta} \rangle < 0.08 - 0.16$ eV [64].

Liquid Scintillator - KamLAND-Zen and SNO+

Liquid scintillator detectors have a high light yield, are relatively inexpensive and allow for a large source mass which make them a good option for neutrinoless double-beta decay searches.

KamLAND-Zen is located in the Kamioka mine near Toyama, Japan. It is the upgrade to the KamLAND experiment. Located at the centre of the detector is a 3.08 m diameter spherical balloon containing 13 tons of Xe-loaded scintillator which contains 129 kg of ¹³⁶Xe. Surrounding this is an outer balloon containing 1 kton of liquid scintillator which provides shielding from external backgrounds. Approximately 1900 photomultiplier tubes (PMTs) surround the detector which detect scintillation light from interactions within the detector [65]. A recent result from KamLAND-Zen claimed a lower limit of $T_{1/2}^{0\nu\beta\beta} >$ 1.07×10^{26} 90% CL which corresponds to a value of $\langle m_{\beta\beta} \rangle < 61 - 165$ meV after 504 kg·yr of exposure [66].

Located in SNOLAB near Sudbury, Canada, SNO+ is the upgrade to the Sudbury Neutrino Observatory (SNO). Instead of using D₂0, which was used in SNO, for the active target mass, it is using 0.3% Te loaded scintillator for its first phase. Figure 2.8 shows the expected backgrounds and signal contribution at $\langle m_{\beta\beta} \rangle = 200$ meV for energies between 1.98 and 3 MeV. The dominant background for SNO+ is the double beta decay spectrum due to the energy resolution of the detector shifting events towards the Q-value of ¹³⁰Te. With 2.34 tons of Te and 5 years of data taking, SNO+ is expected to achieve a limit



Figure 2.8: A stacked plot showing the energy spectrum of the expected backgrounds and observation of neutrinoless double-beta decay with an effective Majorana mass of $\langle m_{\beta\beta} \rangle = 200 \text{ meV}$ in SNO+. Inset is a residual histogram showing the signal contribution above the expected background rate. The events shown are for a 3.5 m fiducial volume cut, 0.3% Te loading, and a resolution of 200 NHit/MeV. $T_{\beta\beta}$ is the observed kinetic energy [67].

of $T_{1/2}^{0\nu\beta\beta} > 9 \times 10^{25}$ 90% CL which corresponds to a value of $\langle m_{\beta\beta} \rangle < 55 - 133$ meV. The next phase of SNO+ aims for 3% Te loading and is expected to reach a sensitivity of $\langle m_{\beta\beta} \rangle < 19 - 46$ meV [67]. The SNO+ detector is described in detail in the following chapter.

Chapter 3

The SNO+ Experiment

SNO+ is a multi-purpose liquid scintillator detector based at the SNOLAB, the Canadian, international underground research facility. The primary focus of the SNO+ experiment is to search for neutrinoless double-beta decay, achieved by loading liquid scintillator with ¹³⁰Te. Additional physics goals for SNO+ involve measurements of low energy solar, reactor, supernovae and geo-neutrinos. A period of pure liquid scintillator data taking would be required to examine neutrinos with an energy less than the ¹³⁰Te Q value of 2527.01±0.32 keV [3] due to the irreducible two neutrino background. Anti-neutrinos could be investigated below the ¹³⁰Te Q-value by tagging the delayed neutron interactions from inverse beta decay. SNO+ utilises existing infrastructure from the Sudbury Neutrino Observatory (SNO) with several upgrades required due to the change in primary active medium from D₂O to liquid scintillator. The raw observables of SNO+ are the charge and hit times of its PhotoMultiplier Tubes (PMTs). The event position is based on time of flight measurements to the PMTs and the energy of an event is related to the number of hit PMTs. A more detailed analysis of these variables leads to a reconstructed position and energy.

3.1 The SNO+ Detector

Situated in Creighton Mine, near Sudbury, Canada, the SNO+ detector is 2092 ± 6 m underground with respect to its centre. The rock above the detector provides 5890 ± 94 metres water equivalent shielding and reduces the cosmic muon flux to a measured $(3.31 \pm 0.01 \text{ (stat.)}\pm0.09 \text{ (sys.)})\times10^{-10} \ \mu/\text{s/cm}^2$ above a zenith angle of $\cos\theta_{zenith} > 0.4$ [68]. A barrel shaped cavity hosts the major structures of the SNO+ detector that are sketched in Figure 3.1. At the centre of the detector is the Acrylic Vessel (AV) that will contain the



Figure 3.1: A schematic of the SNO+ detector in the SNO cavern (black). The acrylic vessel is shown in blue and is suspended from the deck above from the hold-up ropes (pink) and held in position by the hold-down ropes (red). The stainless steel struts of the PMT support structure are shown in grey [69].

liquid scintillator mixture. The AV is held in to place by a system of hold up and down ropes. Surrounding the AV is the PMT SUPport structure (PSUP) that hosts the PMT array. Each of these components is described in detail below.

The AV has a spherical body that has a 12.01 m inner diameter and is nominally 55 mm thick except for 10 of the 122 ultraviolet transmitting panels used to construct the AV. These have a nominal thickness of 11.4 cm [70]. Located at the equator of the AV, these are the rope grove panels and suspend the AV from the deck by ten loops of 19.5 mm thick Tensylon rope [69]. A spherical shape was chosen as it has the highest volume to surface ratio to minimise the contamination at the surface of the AV per unit mass of the active medium within it. The energy and position reconstruction are also simplified with a spherical geometry. The geometry of SNO+ makes it a graded-shield type detector.

This means that the centre of the AV has the least background activity originating from external sources with rates increasing as a function of radius as the self-shielding decreases. Hence, spherical volume cuts with respect to the centre of the AV will improve the ratio of signal to background in a given data set. The volume cut applied is known as the *fiducial volume*. A 6.8 m high, 1.46 m inner diameter neck completes the AV and allows access inside from the deck above for calibration sources and piping for filling and recirculation of the contents within the AV [70].

By filling the AV with liquid scintillator, the energy threshold of SNO+ will be approximately 200 keV with an expected energy resolution of 4.5% at the ¹³⁰Te Q-value [71]. Although light produced from interactions within liquid scintillator is largely isotropic, through time of flight measurements at hit PMTs, it is estimated that it is possible to achieve a position resolution of approximately 10 cm [72]. This is important for the fiducial volume cut described above. The density of the liquid scintillator used in SNO+ is 0.86 gcm⁻³ [73], which is less dense than the water surrounding the AV. To deal with the upward buoyant force on the AV a hold-down rope net, consisting of 20 interwoven 38.1 mm thick Tensylon ropes [69] and anchored to the cavity floor, has been installed.

Surrounding the AV is a 17.8 m diameter geodesic structure, the PMT SUPport structure (PSUP), that holds approximately 9500 R1408 Hamamatsu PMTs which each have a 20 cm diameter. A non-imaging light concentrator surrounds each PMT to increases the photocathode coverage to 54%.

The cavity where the detector is located has a diameter of 22 m at its waist and a height of 34 m. This cavity is filled with approximately 7000 tonnes of Ultra-Pure Water (UPW) for shielding the contents of the AV against the radioactivity of the PMTs. In order to make the cavity water proof and reduce the diffusion of radon from the rock the walls are lined with concrete and urylon [70].

To maintain the radiopurity and cleanliness requirements and to reduce backgrounds during the experiment's lifetime there has been a large investment in resources in a number of areas. This includes the installation of a processing and piping system for SNO+ to purify the liquid scintillator underground in a facility adjacent to the SNO+ detector. This will minimize the production of cosmogenically induced backgrounds. A cover gas system has been installed to meet the SNO+ requirements of 650 atoms/day in the air within the AV. It consists of three 240 l nitrogen filled bags. It is designed to compensate for pressure fluctuations on deck through the airflow between the bags and the detector and reduces the radon ingress into the detector from the mine air. The urylon liner on the cavity floor has been relined to repair splits found, which would allow the ingress of radon. As the AV has been exposed to the mine air, the inner surface of the AV has been cleaned to remove any particulates.

3.2 Electronics

The front end electronics are designed to acquire the hit time and charge of PMTs during events whose signal meets a set of trigger requirements. The electronics is split into 19 crates each processing and reading out up to 512 PMTs. One crate contains 16 PMT Interface Cards (PMTICs) and Front End Cards (FECs) with 1 pair referred to as a card. Connected via 32 m 75 Ω waterproof coaxial cable, the PMTIC supplies the high voltage and reads out the signal of up to 32 PMTs. Each PMT is read out by a dedicated channel out of the 32 that the PMTIC and FEC provides. The combination of channel, card and crate numbers allow each PMT to be identified by its logical channel number (lcn) where

$$lcn = 512 \times crate + 32 \times card + channel.$$
(3.1)

One channel contains two Time-to-Amplitude Converters (TACs) and 16 cells which store charge measurements. The data readout by one channel for one event contains one timing measurement (TAC) and three charge values: high gain long integration (QHL), high gain short integration (QHS) and QLX which uses low gain and can be set to long or short integration time¹. The short and long integration times are 70 ns and 400 ns where as the high gain is approximately 12 times greater than low gain. QHS is a measure of the pulse that crossed the discriminator threshold where as QHL spectra will also contain scattered and reflected light. The charge and time measurements are readout by 12 bit analogue to digital converters (ADCs) and hence have a value of 0 to 4095.

The signal from the PMTIC is first processed by one of four Daughter Boards (DBs), hosted on a FEC, each of which processes the signal of eight channels. The triggers of the 8 channels are initially summed by the DB. Two pairs of DBs have their triggers summed by one FEC resulting in two individual signals per trigger leaving each FEC. The crate trigger card will then sum the triggers for an entire crate. The triggers that are being summed are

• NHIT100 - A discrete square wave which is nominally 100 ns in width but can be varied over a small range. The amplitude corresponds to the number of PMTs that

¹QLX is primarily set to long integration time.

have fired during the 100 ns period. NHIT100 is the most commonly used trigger in SNO+.

- NHIT20 A discrete square wave which is nominally 20 ns in width but can be varied up to a maximum width of 50 ns. The amplitude corresponds to the number of PMTs that have fired during the width of the trigger pulse. The NHIT20 trigger can be delayed across different FECs by up to 80 ns. This is useful for varying the position that in time coincidences are studied within the detector and thus for examining backgrounds within the fiducial volume.
- ESUMHI A high gain sum of the raw PMT signals. Used for diagnostic purposes or creating a more complex trigger mask.
- ESUMLO A low gain version of ESUMHI.

Separate NHIT100, ESUMHI and ESUMLO triggers exist for the OutWard Looking (OWL) PMTs. These are useful as a veto from cosmogenic events.

Once the triggers have been summed for an entire crate by a CTC each signal is readout by an analogue master trigger card (MTC/A+). There are seven MTC/A+s each of which is responsible for a different detector wide trigger sum described above. If a MTC/A+ trigger is above threshold then a trigger is sent to the Digital Master Trigger Card (MTC/D). This also takes in triggers from the internal pulser and external sources. A trigger mask is set by an operator that designates the trigger conditions necessary to generate a Global Trigger (GT). This GT is sent to the crates on the following rising edge of the 50 MHz clock [74]. The XLink based XL3 on each enables communication to readout the data from trigged channels whilst also providing monitoring information on a crate-by-crate basis [75]. Any triggered channels in the last 400 ns will then have their charge and time readout and stored. Trigger words that denote trigger type are stored if they fired within ± 10 ns of the GT regardless of whether they are in the trigger mask or not.

There are 26 trigger words available for use within the trigger mask. These include the NHIT20, ESUMHI, ESUMLO and the OWL specific triggers described above. There are three NHIT100 triggers which can be set to three different thresholds NHIT100_HI, NHIT100_MED and NHIT100_LO. Multiple triggers exist for external sources such as calibration systems and triggers for handling Global Positioning System (GPS) timing.

The Trigger Utillity Boards, TUB and TUBII have many useful functions that provide diagnostic tools and features that are essential for running the detector. This includes



Figure 3.2: Diagram of the SNO+ trigger system. The PMT Interface Cards (PMTICs) supply the PMTs with high voltage and read out the signal passing it to the Daughter Board (DB) which calculates the charge and timing of the pulse. The Front End Card (FEC) then reads out the DBs and passes the trigger waveform to the Crate Trigger Card (CTC) which sums the triggers of an entire crate. The Analogue Master Trigger Card then sums up the triggers on all crates and this is passed to the Digital Master Trigger Card (MTC/D) which decides whether the trigger conditions gave been met to issue Global Trigger (GT) which will force all triggered channels to have their charge and time read out. The time taken for a signal to propagate through the electronics is denoted in red with the default values used to simulate the DAQ electronics in the Monte-Carlo in brackets.

additional clocks to the MTC/D's 10 MHz clock. Conversion between different logic families, Transistor-Transitor Logic (TTL) to Emmiter-Coupled Logic (ECL) for example. A discriminator with retriggering functionality. Monitoring tools for the event rate of the detector. It also provides external asynchronous triggers and communication with the detector for calibration systems [76]. A schematic of the electronics of SNO+ is shown in Figure 3.2.

3.3 Calibration of SNO+

There are several parameters that need to be calibrated so that the energy and position resolution is optimised. The electronics calibration provides the conversion of raw ADC counts to physical units of time and applies the pedestal subtraction to the charge as described in Section 3.3.1. The optical parameters, such as attenuation and scattering length of the media within the SNO+ detector, are constrained using optical calibration sources described in Section 3.3.2. Optical sources are also used to calibrate the position of structures, such as the AV and ropes, within the detector as well as determining the synchronicity and gain of each PMT. The multiple radioactive sources described in Section 3.3.3 are used to calibrate the energy response and geometric uniformity of the response of the detector; the light yield of the liquid scintillator mixture and its ability to discriminate between alpha and beta particles.

3.3.1 Electronics calibration

The Electronics CAlibration (ECA) converts raw ADC counts into appropriate units of charge and time. There are two calibrations that the ECA handles: the pedestal calibration sets a baseline for the charge measurements and the time slope calibration converts the TAC readout into nanoseconds. Each of the channel's two TACs and 16 cells described in Section 3.2 must calibrated.

The pedestal ADC value corresponds to the amount of charge present on a cell when no PMT signal is present. Each cell has a different pedestal value, hence all must be measured. This value is used to convert a cell's charge measurements from raw ADC counts to counts above pedestal (cap). Once a cell has made enough measurements, nominally ten, then the median and width of the data set is calculated and stored. The width is defined as the difference between 40^{th} and 70^{th} percentiles of the data set. The median value is the pedestal of the cell whereas the width is recorded for diagnostic purposes.

The time slope calibration is performed in an identical way to the pedestal calibration but instead of using one value of delay between the GT and pedestal, 150 ns typically for a pedestal run, 31 different delays ranging in steps from 10 - 20 ns are taken to sample the entire slope. Larger steps in delays, 20 ns, are used in the body of the slope whereas 10 ns delays are used in the non-linear regions to better profile these features. The median and width of each step is calculated as before. Points that have fewer than the minimum required measurements, nominally six, are removed and any points that are not monotonically increasing with the set delay are denoted as *curl* points as seen in Figure 3.3. The origin of the curl region is due to a known bug in the electronics. TAC ADC counts which can have two values of TAC time are flagged as *pre-curl*. The remaining points are fit with a first order polynomial function. If there are three or fewer remaining points then a χ^2 is performed on the fit otherwise and more commonly, a cubic fit is performed, using



Figure 3.3: An example of a time slope calibration taken suring SNO. The Time-to-Amplitude Converters (TAC) readout is plotted against the TAC time in nanoseconds. TAC times right of the green line are denoted as 'curl' as they are not monotonically increasing with TAC ADC counts. TAC ADC counts above the red dotted line are denoted as 'pre-curl' and can have two values of TAC time. These are flagged by the calibration. Adapted from [77].

the linear fit results as initial values, on the remaining and curl points. The results of the fit are stored and is used to convert TAC counts into nanoseconds [77].

3.3.2 Optical calibration sources

The optical calibration sources are used to synchronise the PMT array, positioning of the AV and obtain constants for the optical model of the medium within the AV. The stringent radiopurity requirements for SNO+ means that source deployment should be kept to a minimum as each time a source is deployed there is a risk of contaminating the liquid scintillator mixture. This was observed during the KamLAND experiment, as seen in Figure 3.4, where the rates ²²²Rn increased when a source was deployed in their liquid scintillator.

As a result the Embedded LED/Laser Light Injection Entity (ELLIE) was developed to allow for continuous optical calibration without deploying a source in the scintillator.



Figure 3.4: The activity of ²²²Rn as a function of time within the KamLAND detector [78]. Rates were calculated using ²¹⁴Bi-²¹⁴Po coincidences within a 5.5 m volume. The red and blue dashed lines indicate the start of calibration runs.

The ELLIE system has three subsystems each consisting of electronics located on the deck and the light injected through optical fibre fed from the deck to different nodal positions on the PSUP. The timing subsystem (TELLIE) calibrates the PMTs' timing and gain while the scattering and the attenuation of the medium inside the detector is measured by SMELLIE and AMELLIE subsystems respectively.

The TELLIE systems generate nanosecond pulses of light from a driver - LED system on deck with a peak wavelength of 503 nm. The intensity of pulses can vary between 10^3 to 10^6 photons per pulse. The light is injected to one of 92 PSUP positions via 47.75 m of PMMA optical fibre. The fibre has a wide emission angle of approximately 15° to allow for maximum coverage of the PMTs. As well as calibrating the PMTs' time offsets and gains TELLIE can also be used to monitor AV position through reflections and position of the rope nets supporting the AV and sources deployed within through shadowing.

The SMELLIE system generates light of 4 wavelengths: 375 nm, 405 nm, 440 nm and 500 nm from a system of lasers on deck and is injected through quartz fibre to 4 PSUP positions, each containing three directions 0° , 10° and 20° with respect to the centre of the detector. The light is collimated and has an opening angle of approximently 2.5° from the fibre. Analysis of PMT hit times and position should allow measurement of the scattering of the light in the medium within the detector. The different positions and directions allows systematics to be checked.

AMELLIE generates light using the same hardware as TELLIE but uses 405 and 520 nm wavelengths and is injected to 4 PSUP positions via 47.5 m of quartz fibre which



Figure 3.5: A sketch of the Embedded LED Light Injection Entity (ELLIE) calibration hardware on deck with an example of a fibre bundle that has been fed through from the deck into the cavity. Two examples of fibre injection points are also shown where the grey cones illustrating the regions of the detector that are directly illuminated.

has an opening angle of approximately 5° . Each injection point has two directions 0 and 20° . The beam illuminates a larger portion of the detector than SMELLIE beams which allows a relative attenuation measurement of the medium to be made by determining the intensity of the light detected.

The laserball is a deployed source used for optical calibration. The laserball will be used at least once each time the medium inside the AV changes. The purpose of the laserball is to calibrate the PMTs' timing and gain. It will verify results found by TELLIE whilst also calibrating out any time offsets between the different TELLIE fibres.

Sources that are deployed into the AV are kept in a sealed glove box above the AV neck in the Deck Clean Room. They are then lowered and positioned, in three dimensions, as required through a set of umbilical and side ropes, sketched in Figure 3.6. To ensure compatibility with the liquid scintillator, the side ropes are 0.5 mm Tensylon and the umbilical is made of Tygothane.

A system of six cameras has been installed at different positions on the PSUP to



Figure 3.6: Sketch of a deployed source within the AV. The position of the source can be manipulated in three dimensions through the ropes and umbilical [72].

monitor the position of sources deployed in the detector. An LED will placed on the source or the umbilical and the cameras will take pictures over a long exposure time. Then, through triangulation, the camera system it is expected to determine the source position within 1 cm. There is also a light source with the camera system which will be used in conjunction to monitor the rope net position again via triangulation.

3.3.3 Radioactive calibration sources

Numerous radioactive sources have been developed for deployment in SNO+. Some sources are untagged whilst others either use coincident decays or a PMT inside the source to tag the calibration sources decay. The ¹⁶N source, which produces a tagged 6.1 MeV γ , will be deployed in the water phase. Other gamma sources that are being developed for the scintillator phase are a ⁶⁰Co source with 1.2 MeV and 1.4 MeV γ s, a ⁴⁸Sc source with 1.0, 1.2 and 1.3 MeV γ s and a ²⁴Na with a 2.7 MeV and 1.3 MeV γ s. Beta sources that are being developed for the scintillator phase include ⁹⁰Y and ⁸Li. An AmBe (n, γ) source is
also under development.

To minimize risk of radiological contamination in the liquid scintillator, the deployment of radioactive sources will be restricted and possibly will only be deployed if an excess of events is observed in the signal region. The experiment aims to maximise the use of internal and external radioactive backgrounds to calibrate the energy response of the detector in order to avoid using deployed sources [79].

3.4 Liquid scintillator

SNO+ uses a Water-based Liquid Scintillator (WbLS) cocktail whereby liquid scintillator molecules and water are co-mixed by a surfactant. The scintillator cocktail that SNO+ is comprised of contains water, the liquid scintillator Linear Alkyl-Benzene (LAB), the wavelength shifters PPO and bis-MSB, the surfactant PRS and, in the double beta phase, tellurium. The major advantage of using a WbLS in SNO+ is that the surfactant allows high loading of the isotope 130 Te. Loading as high as 5% has been demonstrated [80], although high loading does impact the light yield and stability of the cocktail. The light yield of 3% loading is 40% the amount with 0.3% loading due to an increase in the proportion of PRS required in the cocktail to make it stable [81]. WbLS also improves the sensitivity to the directionality of events that are above the Cherenkov threshold. The primary wavelength shifter, PPO, is used to re-emit ultraviolet photons with a peak wavelength of 365 nm. The secondary fluor, bis-MSB, then shifts the photons to a wavelength of approximately between and 380 and 500 nm. This is to optimise the yield of photoelectrons given the wavelength dependence of the the quantum efficiency of the PMTs whilst minimising the absorption of the cocktail. Research and development is ongoing into the exact composition of the cocktail. A cocktail consisting of 0.3% Te, 5% PRS, 2% H₂O LAB-PPO ad 15 mg/L bisMSB has a measured light yield of (9725 ± 515) photons/MeV [82]. There is also ongoing research and development into higher light yield scintillator cocktails.

3.4.1 Linear Alkyl-Benzene

Linear Alkyl-Benzene (LAB) is the liquid scintillator that will be used in the SNO+ experiment. LAB was selected as it is a relatively safe scintillator for use underground and has good chemical compatibility with the AV. It has high flash and boiling points of 140 °C and (278 - 314) °C respectively and a low toxicity [73]. LAB has a light yield of approximately 10^4 photons per MeV [71]. It also has a high attenuation and scattering length of a combined value of approximately 48 m at a wavelength of 470 nm. LAB's fast

3.4.2 Tellurium

In 2013, the double-beta decay isotope selected for SNO+ was switched from 150 Nd to 130 Te. There are several advantages of using natural tellurium [83]

- The high, 34%, natural abundance of ¹³⁰Te which negates the need for costly and complex enrichment.
- The $2\nu 2\beta$ half life of ¹³⁰Te is $(7.14 \pm 1.04) \times 10^{20}$ y in comparison to the $(8.37 \pm 0.45) \times 10^{18}$ y half life of ¹⁵⁰Nd [84] making the intrinsic $2\nu 2\beta$ background of ¹³⁰Te approximately 100 times smaller than ¹⁵⁰Nd.
- It is inexpensive at \$31 USD/kg in comparison to other double-beta decay isotopes.
- Te has less optical absorption than Nd.

As the use of water increases the amount of uranium and thorium chain backgrounds in the cocktail which could be a problem for SNO+. However, the fast timing of scintillator allows Bi-Po α tagging which suppresses U and Th backgrounds that have energies at the Q value of ¹³⁰Te. It has also been shown that the external gammas are expected to be adequately attenuated in the outside the fiducial volume [83].

Tellurium is chemically bonded to LAB by first dissolving Teluric acid, $Te(OH)_6$, in UPW and re-crystallising using nitric acid above ground. Underground the crystals are dissolved in UPW at 80 °C and then recrystallised thermally. The surfactant PRS is used to chemically bond it with the liquid scintillator.

3.5 Backgrounds

Reducing and tagging backgrounds is fundamental to SNO+ in enabling sensitivity to $0\nu 2\beta$. As previously stated in Section 3.1, SNO+ is a graded shield type detector making the centre of AV the point of the lowest background rate with activity, originating outside the inner AV, increasing as a function of radius outwards. There are numerous backgrounds that SNO+ has to control and reduce. They are listed briefly below and described in detail in the following subsections in order of the position of the source of the background from those originating from the centre of the detector outwards

- **Tellurium** Natural tellurium loaded in the liquid scintillator has radioactive impurites.
- Scintillator Residual radioactive impurities in the scintillator cocktail from the production and processing of the LAB+PPO mixture.
- **External** Residual radioactive impurities which are outside the inner AV and deposit energy in the scintillator cocktail.
- Cosmogenic Cosmic ray muons or cosmic ray induced backgrounds.
- Solar Neutrinos Neutrinos originating from the sun. A significant background for neutrinoless double-beta decay.

Any combination of the above backgrounds occurring within the same event is known as pileup.

3.5.1 Tellurium Backgrounds

The telluric acid from which tellurium is extracted contains numerous impurities such as Ag, Y, Zr, Sb and Sn. These are removed through dissolving the telluric acid in nitric acid and then recrystallising the compound. This has been shown to remove impurities to negligible levels [85].

The two isotopes that are known to decay in natural tellurium are ¹²⁸Te and ¹³⁰Te which both decay via double beta decay with half lives of $(3.49 \pm 1.99) \times 10^{24}$ years and $(7.14 \pm 1.04) \times 10^{20}$ years to the stable isotopes of ¹²⁸Xe and ¹³⁰Xe respectively. The Q-value of ¹²⁸Te is 0.87 MeV and thus will not be a background for the search of $0\nu 2\beta$ decay in ¹³⁰Te. The $2\nu 2\beta$ decay of ¹³⁰Te is an irreducible background in the search for neutrinoless double-beta decay in SNO+.

3.5.2 Scintillator Backgrounds

The scintillator cocktail is purified underground in a plant adjacent to the SNO+ detector, see Figure 3.7, so that the number of cosmogenically induced backgrounds is reduced. The different radioactive backgrounds that are in the liquid scintillator are removed in several stages. Heavy metals are removed through multi-stage distillation. This process will also improve the liquid scintillator cocktail's transparency to UV light. Micro-filtration removes particulates such as dust. Stream stripping and N₂ filtration removes Kr, Ar, Rn and O₂. Ra, K and Bi are removed through water extraction. Bi and Pb are also removed in metal scavengers.



Figure 3.7: A flow diagram of the SNO+ liquid scintillator processing and purification facility. The liquid scintillator is first stored on surface before being shipped underground for purification and eventually filling the inner AV [86].

The purification system is expected to reduce the radioactivity in the LAB+PPO mixture to a level seen in Borexino, Table 3.1, as both purification systems used in the respective experiments were built by the same company.

SNO+ is sensitive to is the decay chains of uranium and thorium. Natural uranium is comprised of 99.3% 238 U, 0.7% 235 U and trace 234 U which itself is part of the 238 U decay chain. Natural thorium is mononuclidic and comprised of the isotope 232 Th.

The decay chains of ²³⁸U and ²³²Th are in the appendix A. It is assumed that the decay chains are in equilibrium in the liquid scintillator. This may not be true, for instance, if ²²²Rn enters the liquid scintillator and through its decay chain produces an excess of isotopes below this in the ²³⁸U decay chain. The decay chains could also be out of equilibrium due to differences in solubility of the isotopes.

The other backgrounds in the liquid scintillator include the beta decays of 39 Ar, 85 Kr

Isotope	Borexino Level (g/g)
$^{238}\mathrm{U}$	1.6×10^{-17} [87]
232 Th	6.8×10^{-18} [87]
$^{210}\mathrm{Pb}$	6.11×10^{-25} [88]
$^{210}\mathrm{Bi}$	3.78×10^{-28} [88]
$^{85}\mathrm{Kr}$	$< 2.40 \times 10^{-25}$ [88]
^{40}K	$< 1.30 \times 10^{-18} [88]$
$^{39}\mathrm{Ar}$	2.75×10^{-24} [88]
$^{14}\mathrm{C}$	1.00×10^{-18} [88]

Table 3.1: The levels of isotopes within the liquid scintillator as measured by the Borexino experiment.

and ¹⁴C which have end points of 0.565 MeV, 0.687 MeV and 0.156 MeV respectively. The isotope ⁴⁰K beta decays 89.3% of the time with an end point of 1.311 MeV, 10.72% of decays will be electron capture emitting 1.46 MeV gamma and very rarely (branching ratio = 0.001%) ⁴⁰K will decay via positron emission with an end point of 0.483 MeV.

3.5.3 External Radioactive Backgrounds

There are numerous naturally occurring radioactive sources in the PMTs, ropes, water and AV of the SNO+ detector. These will produce charged particles and gammas when they decay. While the water will likely shield the inner AV, through attenuation, from charged particles originating from the PSUP, water and ropes, decays which produce high energy gammas are still a concern. In this section, gammas with an energy greater than 2 MeV will be discussed as these are backgrounds for the double beta phase of the experiment. Gammas with a lower energy will be a background in a pure liquid scintillator phase.

Each PMT has an estimated 100 μ g of uranium and thorium within them [70]. Only ²¹⁴Bi of the ²³⁸U and ²⁰⁸Tl of the ²³²Th decay chain produce high energy gammas in their decays. There are no high energy gammas produced in the decay chain of ²³⁵U. When isotopes emit beta particles that have an end-point in their spectrum which is less than the Q-value of the parent isotope then the rest of the energy is emitted as the total energy of all coincident gammas. The Q-value of ²¹⁴Bi is 3.272 MeV. The maximum energy of the beta particles emitted in its predominant decay modes are denoted in Table 3.2. Approximately 21.8% of the gammas emitted in coincidence of ²¹⁴Bi decays will have a sum total energy > 2 MeV, all gammas produced in decays of ²⁰⁸Tl will have a sum energy greater than this value.

The UPW which surrounds the AV and PSUP is purified from water taken from the mine supply, resulting in its saturation of mine air and particulates. A deaerator first removes the O_2 and N_2 . Large particulates are then removed via a multimedia filter and 10 micron filters. This is followed by charcoal filters to remove organic matter. To remove divalent ions and prepare the water for reverse osmosis, used at a later stage of purification, cation exchange resins in zeolite softeners are used. The water is then mixed with EthyleneDiamineTetraacetic Acid (EDTA) to form complexes with ions. Any EDTA complexes or molecules with a weight greater than 200 u are then removed via reverse osmosis. Further organic purification is achieved through the use of UV light and ion-exchange. The water is regassed with N_2 to ensure the water is at atmospheric pressure to reduce the probability of PMT breakdowns which occur at low pressures. The UPW

is then chilled and injected into the cavity. The purification process is shown visually in Figure 3.8. After purification the UPW contains $< 35 \times 10^{-14}$ g/g ²³⁸U and 3×10^{-14} g/g ²³²Th.

The ropes made of Tensylon contain 0.1 ppb ²³⁸U, 0.17 ppb ²³²Th and 250 ppb natural K. Natural potassium consists of 93.3% ³⁹K, 0.0117% ⁴⁰K and 6.7% ⁴¹K. The isotopes ³⁹K and ⁴¹K are stable. The ⁴⁰K decays to the stable isotopes ⁴⁰Ar and ⁴⁰Ca via electron capture and beta decay respectively. The electron capture has a branching ratio of 10.72% and will result in a 1.46 MeV gamma. The beta decay has a branching ratio of 89.28% and has an end point of 1.311 MeV.

The AV acrylic contains < 2.3 ppb of 40 K and < 1.1 ppt of both uranium and thorium. Due to 222 Rn being electrostatically attracted to the surface of the AV, there is an unknown quantity of 210 Pb and 210 Po embedded near the surface of the AV. This is due to their large half lives of 22.3 years and 138 days in comparison to other isotopes below 222 Rn in the decay chain.

Depending on the medium within the AV, ²¹⁰Pb and ²¹⁰Po will leach out of the acrylic at different rates. The leaching rate (LR) is described by the formula

Parent Isotope	E_{max} (MeV)	BR (%)
²¹⁴ Bi	3.270	19.10
	1.892	7.35
	1.727	3.12
	1.540	17.57
	1.505	16.96
	1.423	8.14
	1.380	1.588
	1.275	1.177
	1.259	1.431
	1.253	2.450
	1.151	4.345
	1.066	5.66
	0.822	2.78
	0.788	1.244
²⁰⁸ Tl	1.801	49.1
	1.524	22.2
	1.291	24.2
	1.038	3.18

$$A(t) = A_0 e^{-t/\tau}.$$
 (3.2)

Table 3.2: The maximum energy E_{max} of the emitted beta and its branching ratio (BR) of the predominant (BR $\geq 1\%$) decay modes of the isotopes ²¹⁴Bi and ²⁰⁸Tl.



Figure 3.8: A flow diagram of the SNO+ water processing and purification facility. The water is from the mine supply where it is purified before entering the cavity. Adapted from [70].

Medium	Isotope	Temperature (°C)	$LR (d^{-1})$	au (d)
	²¹⁰ Pb	25	$(3.9\pm0.5)\times10^{-3}$	256 ± 32
$\mathbf{I} \mathbf{A} \mathbf{D} + \mathbf{D} \mathbf{D} \mathbf{C} (\mathbf{r} (\mathbf{f}) + \mathbf{H} \mathbf{O} + \mathbf{T}_{\mathbf{r}} (\mathbf{O} \mathbf{H}) (0, 9 (\mathbf{f}))$	²¹⁰ Po	25	$(3.5 \pm 1.0) \times 10^{-3}$	288 ± 83
LAD+1 $R_{5(5/0)}$ +1120+16(011)6(0.370)	210 Pb	12	$(1.83 \pm 0.29) \times 10^{-3}$	$546\pm^{103}_{74}$
	²¹⁰ Po	12	$< 1.53 \times 10^{-3}$	> 655
LAB+PRS(5%)+H ₂ O(0.25%)	²¹⁰ Pb	25	$(8.7 \pm 0.9) \times 10^{-3}$	115 ± 11
IAR	²¹⁰ Pb	25	$(8.0\pm0.3)\times10^{-4}$	1250 ± 47
LAD	²¹⁰ Pb	12	$(1.4 \pm 0.2) \times 10^{-4}$	7200 ± 900
	²¹⁰ Pb	25	$(2.2\pm0.2)\times10^{-3}$	455 ± 45
$H_{a}O(UPW)$	²¹⁰ Po	25	$(2.1 \pm 0.6) \times 10^{-3}$	476 ± 143
$\Pi_2 O(OPW)$	210 Pb	12	$(4.8 \pm 1.2) \times 10^{-4}$	2083 ± 500
	²¹⁰ Po	12	$(3.8 \pm 0.6) \times 10^{-4}$	2632 ± 400
	$^{210}\mathrm{Pb}$	25	$(1.3 \pm 0.3) \times 10^{-2}$	74 ± 14
UDW + 0.95M FDTA	²¹⁰ Po	25	$(1.3 \pm 0.4) \times 10^{-3}$	77 ± 22
OFW + 0.25WI EDTR	$^{210}\mathrm{Pb}$	12	$(1.24 \pm 0.09) \times 10^{-4}$	81 ± 6
	²¹⁰ Po	12	$< 3.2 \times 10^{-3}$	> 310
UDW + 0.097M EDTA	²¹⁰ Pb	12	$(6.4 \pm 0.6) \times 10^{-3}$	$156\pm^{16}_{13}$
$\mathbf{OPW} + 0.027 \mathbf{M} \mathbf{ED1A}$	²¹⁰ Po	12	$(3.7 \pm 1.5) \times 10^{-3}$	$270\pm_{78}^{185}$

Table 3.3: The leaching rates (LR) and leaching time constants (τ) of the isotopes ²¹⁰Pb and ²¹⁰Po in various medium [90]. The measurements were taken at two temperatures 25 and 12 °C.

Where A_0 is the initial activity on the surface of the AV and A(t) is the activity at time t. The leaching time constant τ is the inverse of LR. Measurements of τ are listed in table 3.3. Seven metric tons of EDTA is required to achieve a concentration of 0.027M of EDTA in the UPW during the water phase of SNO+. A period of 6 months of UPW+0.027M EDTA in the AV would reduce the amount of ²¹⁰Pb and ²¹⁰Po by approximately 70% and 50%. A nine month period would lead to an approximate reduction of 83% and 63% of ²¹⁰Pb and ²¹⁰Po [89].

3.5.4 Cosmogenics

The overburden the mine provides reduces the rate of cosmic ray muons passing through the AV to approximately 52 muons per day with $\cos \theta_{zenith} > 0.4$. These must have a minimum energy of approximately 3 TeV to reach the SNO+ detector from the surface. Muons can also be atmospheric neutrino induced through interactions in the rock; and these have a minimum energy at production of 2 GeV to be detected at SNO+ [68]. Muons can cause background events in the detector directly and indirectly. Direct events are due to the muon itself ionising the scintillator. When muons interact they can also produce high energy neutrons which may be absorbed by atoms such as ¹²C that are within the liquid scintillator. The daughter particles, known as muon-induced backgrounds, are unstable and decay causing the indirect background events.

The high energy of muons passing through the detector make them easy to identify and reject. The identification of muonic events, combined with their low rate, also allows

Decay	Isotope	$T_{1/2}$	E_{max} (MeV)	BR (%)
β^{-}	$^{12}\mathrm{B}$	0.02 s	13.4	100
	$^{11}\mathrm{Be}$	$13.80~\mathrm{s}$	11.5	100
	11 Li	$0.09 \ { m s}$	20.8	100
	⁹ Li	$0.18 \mathrm{~s}$	13.6	100
	⁸ Li	$0.84 \mathrm{~s}$	16.0	100
	⁸ He	$0.12 \mathrm{~s}$	10.6	100
	⁶ He	$0.81 \mathrm{~s}$	3.5	100
β^+	$^{11}\mathrm{C}$	$20.38 \min$	0.96	99.79
EC	$^{11}\mathrm{C}$	$20.38 \min$	1.98	0.21
β^+, γ	^{10}C	$19.30~\mathrm{s}$	$1.91(\beta) \ 0.72(\gamma)$	98.5
	^{10}C	$19.30~\mathrm{s}$	$0.88(\beta) \ 1.02(\gamma)$	1.5
β^+	⁹ C	$0.13 \mathrm{~s}$	16.0	100
	⁸ B	$0.77 \mathrm{\ s}$	13.7	100
EC	⁷ Be	$53.3 \ d$	$0.48 (\gamma, 10.44\%)$	100

Table 3.4: The half life $(T_{1/2})$ decay, branching ratio (BR) and maximum energy of the decay (E_{max}) of muon-induced backgrounds produced through interactions with ¹²C within the scintillator [91].

those muon-induced backgrounds which have a half life of order seconds to be removed with a few minute veto. Examining the half lives in Table 3.4, this will remove all muoninduced backgrounds associated with ¹²C apart from ¹¹C and ⁷Be. However, both these backgrounds have a Q value (E_{max}) that is lower than that of ¹³⁰Te and so will only be of significant concern during a pure liquid scintillator phase.

Other muon-induced backgrounds occur from neutron absorption by other isotopes in the liquid scintillator and UPW. Neutron capture by Hydrogen produces Deuterium which releases a 2.2 MeV gamma.

3.5.5 Solar Neutrinos

The ⁸B neutrino flux is an irreducible background for SNO+ during the double beta phase. However constraints on the expected flux can be made using previous SNO measurements. The ⁸B neutrinos interact via elastic scattering in the liquid scintillator. The flux rises as a function of energy peaking at approximately 6 MeV. The cross section also rises linearly with a function of energy. Conversely, the kinematics of the interaction are such that the neutrino does not transfer all its energy to the electron. The event rate therefore peaks at a lower energy because of this and is slightly decreasing throughout the $0\nu 2\beta$ signal region.

3.5.6 Pileup

Pileup is multiple backgrounds decaying within one event trigger window. The event trigger window is 400 ns so only high rate backgrounds are likely to pileup. The event will have an energy equal to the sum of the decay energies. Given that background A has decayed, the probability background B will decay within the event trigger window is

$$P(k,\lambda) = \frac{\lambda^k e^{-\lambda}}{k!},\tag{3.3}$$

where k is the number of pileup events and λ is the average number of decays of the background B in an event trigger window. It is also important to consider both time orderings of pile up.

3.6 Monte-Carlo

The Monte-Carlo simulation SNO+ uses is Reactor Analysis Tools (RAT) [92]. Originally developed for the Braidwood collaboration [93], it is based on a KamLAND-like detector

for use with generic liquid scintillator experiments. RAT is primarily coded in C++ with a python interface available. Geant4 libraries [94] are utilised for simulation of physical processes and geometry while ROOT libraries are used for data manipulation and storage. RAT is also being developed on distinct branches by the DEAP and CLEAN dark matter experiments based at SNOLAB [72].

Chapter 4

PMT Timing and Gain Calibration System (TELLIE)

The in-situ calibration of PMT timing and gain is handled by the Timing component of the Embedded LED Light Injection Entity (TELLIE). This chapter explains the requirements of the PMT calibration for SNO+ in Section 4.1 and describes the hardware and software that has been developed primarily by the University of Sussex, prior to this thesis, in sections 4.2 and 4.3 respectively. Section 4.4 describes the procedure and results of the characterisation and quality assurance of TELLIE, conducted by myself, undertaken at the University of Sussex prior to its installation underground at SNOLAB. Results from the commissioning of TELLIE installed in the SNO+ detector are presented in Section 4.5. All analysis of the commissioning was performed by myself apart from the PMT calibration. As a result of the experience gained during the commissioning runs of TELLIE, data quality requirements, designed by myself and described in Section 4.6, have been developed and encoded in software to ensure that future TELLIE runs are of a high standard for analysis and calibration.

4.1 Requirements of PMT Calibration

To understand the requirements of the PMT calibration it is necessary to consider various aspects of the SNO+ detector and the analysis of the data it collects. The experimental observables of SNO+ are: the number and distribution of PMTs which have a large enough signal to cross a fixed discriminator threshold to register a hit; the integrated charge of the PMT pulse and the time at which the hit occurred. From these variables a reconstructed position and energy is obtained. The position reconstruction is achieved via time-of-flight calculations to the hit PMTs. Reconstructed energy is calculated through the number of hit PMTs and the reconstructed position, hence also requires accurate timing measurements. The time profile of detected events is dependent on several factors: the decay time of the scintillator signal, 4.6 ± 0.5 ns for beta particles emitted from a ⁹⁰Sr source in the LAB/PPO mixture; the intrinsic PMT resolution, 1.7 ns for SNO PMTs [95] and the overall synchronisation of the PMT array. If the synchronisation is to be the sub-dominant effect then time offsets of 1 ns or less between all PMTs is required for simultaneous signals of equivalent intensity.

The Electronics Calibration (ECA) accounts for channel-dependant signal amplitudes and delays in the electronics. As discussed in Section 3.3.1, there are two outputs of the ECA

- the number of ADC counts for each cell which corresponds to no detected signal. This is known as a pedestal.
- conversion between TAC counts to the point in time at which the global trigger occurred. This converts ADC into a time in nanoseconds.

However, there are also unknown timing offsets due to PMT cable delays which is external from the ECA and must be corrected for by the PMT Calibration (PCA) to ensure synchronisation.

The charge a photoelectron avalanche deposits in a PMT varies due to statistical fluctuations. Figure 4.1 illustrates the time difference between large and small pulses. Signals with smaller amplitudes will cross the discriminator threshold at a later time than those with larger amplitudes. Therefore, due to amplitude fluctuations, there is a systematic variation on when the signal crosses the discriminator threshold with respect to when the photon arrives at the PMT. This phenomena is known as timewalk and has to be accounted for by the PCA. The high gain short-integrated charge (QHS), rather than the long-integrated, is used for timewalk corrections so that it is the initial pulse which triggered the discriminator that is calibrated and not any subsequent pulses which may occur during the event.

SNO+ expects between 200 - 300 detected photoelectrons per MeV_{ee}, where *ee* denotes electron equivalent, for energy deposited in the liquid scintillator. With 9500 PMTs a typical event with 2.5 MeV_{ee} produces 500 - 750 photoelectrons, therefore only (4 - 7)%of PMTs will typically be hit. The energy range of the signal interactions used for most of the physics goals of SNO+ is less than 10 MeV_{ee} and even at these higher energies the probability for a PMT hit is less than 0.32 therefore most PMT hits will be single



Figure 4.1: A sketch of integrated charge (QHS) vs. time. The charge deposited in a channel is measured in counts above pedestal (cap), where a pedestal is the average number of counts a channel readouts when no charge is deposited. With respect to the centre of the peak, larger pulses cross the discriminator threshold (dashed horizontal line) at an earlier time than smaller pulses.

photoelectrons. However, as the energy gets larger, the number of multi-photoelectron hits grows, approximately 4% of PMTs for energies of 10 MeV_{ee}. The aim for SNO+ is to measure the gain of PMTs within the dynamic range of 3-4 photoelectrons. This will aid energy reconstruction and assess its stability. It is the role of the PCA to measure and correct the gain and timing variation of the PMTs.

The main requirement for TELLIE is that it can achieve the timing calibration of the PMT array to a precision of 1 ns within 24 hours operating at a frequency at 10 Hz or less to minimise dead time during physics data taking whilst achieving constant monitoring of the PMT calibration. To calculate the run time needed to achieve this precision one begins by assuming that the LED pulses of TELLIE are Gaussian. The measured precision of the mean of a Gaussian, $\Delta \mu$, is given by

$$\Delta \mu = \frac{\sigma}{\sqrt{N}},\tag{4.1}$$

where, in the case of the SNO+ detector, N is the number of hits of a PMT and σ is the

standard deviation of the time profile of the LED pulse. However the light pulses are not truly Gaussian and can be better described by a combination of 2 Gaussians of fixed mean with differing σ to reflect the short rise time and longer fall time. The rise and fall times of the light pulses are defined in this thesis as the difference in time where the amplitude of the light pulse is at 10% and 90% of its peak on the rising and falling edges of the light pulse respectively. It has been shown empirically [96] that $\Delta \mu$ for the TELLIE pulses is equivalent to

$$\Delta \mu = \frac{\text{FWHM}}{\sqrt{N}},\tag{4.2}$$

where FWHM is the full width half maximum of the TELLIE pulses. As N is the product of the number of TELLIE pulses, N_{pulses} , and the probability of the PMT hit, P_{hit} , to achieve the required $\Delta \mu$ one requires the following N_{pulses}

$$N_{pulses} = \frac{\text{FWHM}^2}{\Delta \mu^2 P_{hit}}.$$
(4.3)

Therefore to minimise N_{pulses} one must minimise the FWHM of the TELLIE pulses. The only remaining unknown is P_{hit} which depends on the number of photons injected per pulse and the angle of the PMT with respect to the direction of the injected light. Monte Carlo is used to calculate P_{hit} by simulating TELLIE pulses from the fibre injection points at various intensities with the AV filled with the different media of each phase of the experiment. The mean hit probability, μ , is then extracted from the Monte Carlo and Poisson statistics can be used to calculate the probability of whether a PMT hit is caused by single or multi-photoelectrons. The probability of k photoelectrons being created at a PMT is therefore given by

$$P_k = \frac{\mu^k}{k!} e^{-\mu}.$$
 (4.4)

Therefore the probability of a single photoelectron hit is

$$P_1 = \mu e^{-\mu}, \tag{4.5}$$

and the probability of a multi-photoelectron hit is

$$P_{k>1} = 1 - P_0 - P_1 = 1 - (1+\mu)e^{-\mu}.$$
(4.6)

Therefore, using equations 4.3 and 4.5 we can now calculate the required N_{pulses} . The total run time is therefore just the product of N_{pulses} , the pulse frequency, 10 Hz, and the

number of injection points (92). The total run times for different numbers of photons per pulse are given in Table 4.1. From this Table it is clear that LED pulses with a FWHM of less than 6 ns are more than sufficient for the required run time for PMT calibration. The centre spot in Table 4.1 is defined as the angle from the fibre direction at which the intensity of the emitted pulses is at 20% of its maximum. Based on measurements described in Section 4.2.3, this angle is 14.5°. The fringe region is defined by PMTs whose angle is between 14.5° and 29°.

	Photons	Total Hi	otal Hit Probability Run Time Per Injection Point Total Run Tim		Run Time Per Injection Point		ın Time
Phase:	per pulse:	Centre:	Fringe:	Centre (s):	Fringe (s):	Centre (hr):	Fringe (hr):
UPW	1000	0.027	0.003	16.470	169.671	0.416	4.289
	10000	0.237	0.026	1.897	17.564	0.048	0.444
	100000	0.884	0.201	0.509	2.237	0.013	0.057
LAB+PPO	1000	0.039	0.003	11.618	169.033	0.294	4.273
	10000	0.313	0.026	1.437	17.031	0.036	0.431
	100000	0.914	0.213	0.492	2.111	0.012	0.053
LAB+PPO+Te	1000	0.036	0.002	12.421	180.573	0.314	4.564
(0.3%)	10000	0.296	0.024	1.519	18.427	0.038	0.466
	100000	0.904	0.199	0.497	2.259	0.013	0.057

Table 4.1: The number of simulated photons per pulse in a TELLIE light pulse; probability that a PMT will acquire a hit; per injection point and total TELLIE running times for the different phases of SNO+. The hit probabilities were calculated from Monte Carlo simulation of a TELLIE channel consisting of a 504 nm LED at a given intensity measured in photons per pulse [97]. The run times are calculated by assuming a pulse width of 6 ns and a desired accuracy of 1 ns running at 10 Hz. The centre and fringe areas of PMTs are defined as being between 0° and 14.5° and 29° respectively where the angle is defined as being between the incident fibre direction and the PMT position.



Figure 4.2: A photo of the components that host the electronics of the timing component of Embedded LED Light Injection Entity (ELLIE) housed in the ELLIE rack on the deck of the SNO+ detector. The electronics handles the control and light pulse generation of the TELLIE system. They consist of a power supply, a control box and twelve boxes with each containing eight light emitting channels. The boxes are daisy chained via a ribbon cable to the control box which handles communication between the DAQ system and the boxes. The control box is connected to the DAQ system via a USB port. The control box also hosts two BNC connections. The left BNC connection accepts trigger signals from the SNO+ DAQ system and the right outputs trigger signals from TELLIE to the SNO+ DAQ system. The power supply is located at the top of the ELLIE rack.



Figure 4.3: Left: A photo of an LED that has been drilled and has had an optical fibre attached. The ruler is in centimetres. Right: A photo of two TELLIE channels. Each channel consists of a driver board, an optical fibre, one of the four electronic channels housed on a PIN diode readout board and a coupler made from common brass plumbing fitments with custom end caps to make the coupler light tight. The coupler houses an LED, fibre connection and a photodiode which monitors the light output of the LED.

4.2 Hardware

The components that host the electronics of the TELLIE system are a power supply, control box and twelve *boxes* as indicated in Figure 4.2. Each box contains eight *channels*. The components of one channel are an LED, LED driver board, optical fibre, a photodiode and one of four channels housed on a PIN diode readout board. The light output of each channel has a unique intensity and time profile. All the channels in one box are connected to one motherboard which can communicate with the control box through a serial port on the front of the box. Each channel has one push-pull LED driver board which applies the voltage across an LED coupled to an optical fibre. Figure 4.3a shows an LED which has had a hole drilled into the LED lens to allow the 1 mm diameter optical fibre to be attached to it. One end of the fibre is approximately 1 mm from the surface of the semiconductor which emits the light. The other end of the fibre is terminated at the front of the box via an ST connection. A PIN diode board, which hosts four channels, provides measurements of the light output from each of the four LEDs. One of the four photodiodes connected to the PIN diode board; an LED and its connecting fibre are all housed in a common, brass plumbing fitment which has custom end caps to make the fitment light tight. This is known as a coupler and can be seen in Figure 4.3b. The internal electronics of a complete box can be seen in Figure 4.4.

The communication between a TELLIE channel and the SNO+ DAQ is handled by a separate control box which is connected to each box through two daisy chained ribbon



Figure 4.4: A picture of the internal electronics of a TELLIE box consisting of eight channels, housed on one motherboard, each containing one LED driver, coupler, fibre, terminating at the front of the box at an ST connector and a channel on one of two PIN diode boards. Two identical power sockets and a serial port for communication are visible on the front of the box.

cables. The upper set contains five boxes and the attenuation calibration system AMEL-LIE, the lower seven. The DAQ computer is connected to the control box via a USB cable. The control box converts the USB signal to serial port internally. The control box has two trigger ports. One trigger port issues and the other accepts ECL trigger signals down a BNC cable to and from the SNO+ DAQ system. The impedance of the cable is 50 Ω to match the impedance of the ports.

The TELLIE power supply is housed at the top of the ELLIE rack. Any number of boxes can be powered at one time due to a regulated power supply adjusting the overall current depending on the required output.

To simplify the connection and mapping of fibres which have been installed on the PSUP to the TELLIE channels a patch panel is used. Figure 4.5 shows a photograph of the front patch panel before any fibres are connected to it. The patch panel is secured to the bottom of the ELLIE electronics rack and has 250 ST connectors either side. A 2 m



Figure 4.5: A photograph of the front patch panel which is secured to the bottom of the ELLIE electronics rack. The orange optical fibres are part of the scattering subsystem of ELLIE. All optical fibres which are connected to the electronics of the ELLIE subsystems, known as patch fibres, are connected via an ST connection at the front side of the patch panel. There are also ST connections at the back of the patch panel where the dry end of the optical fibres installed on the PSUP are connected to. There are at total of 250 ST connections either side of the patch panel.

optical fibre known as a patch fibre is used to connect the fibre which is terminated at the face of a TELLIE box to the patch panel. The optical fibres installed on the PSUP are connected to fixed positions on the patch panel and are protected by a metal fibre guide. The patch fibres are then used to map the optical fibres to the TELLIE channels. Each time fibres are moved there is a risk of damaging them. Hence mapping with patch fibres minimises the risk of damaging the installed optical fibres and if any patch fibres are damaged during remapping then they can be easily replaced on the SNO+ deck.

4.2.1 LED selection

The LED spectrum is optimised at a peak wavelength, λ , that maximises the quantum efficiency of the PMT whilst also minimising the absorption probability of the liquid scintillator mixture to optimise direct transmission to the PMTs. A higher wavelength is also desirable to minimise Rayleigh scattering which is proportional to λ^{-4} . From examining Figure 4.6, the optimum wavelength was found to be close to 500 nm. The average peak wavelength of all 96 LEDs was measured to be 503 nm. Details of this measurement are described in Section 4.4.1.



Figure 4.6: A plot of the absorption probability of 0.3% Te-doped LAB+PPO in arbitrary units (vertical line shading), PMT quantum efficiency (horizontal line shading) and the average intensity in arbitrary units of all 96 LEDs as a function of wavelength [95].

To achieve the timing calibration requirements described in Section 4.1, LED pulses with a width of a few nanoseconds are necessary to reduce the amount of data required to calibrate the PMT array.

The LED photon emission rate can be described using the approximation for the low excitation LED modulation characteristics. If an active LED has an active region of volume V_a and is switched on at time t = 0 with a supplied constant current I then the carrier concentration n_a will increase as electrons flow into V_a . Given the average time taken for a charge carrier to recombine and generate a photon is τ_e , the emitted rate of photons per volume is $dn/dt = n_a/\tau_e$. If the current across the pn-junction, I, is taken to be constant once the voltage has exceeded the threshold, and neglecting the stationary charge densities either side of the active region of the LED and their intrinsic capacitance, then the emitted rate of photons is

$$\frac{dn}{dt}(t) = I \cdot \frac{1}{q_e V_a} \cdot \frac{\tau_l}{\tau_e} \cdot (1 - e^{-t/\tau_l}), \qquad (4.7)$$

where q_e is the charge of an electron and τ_l is the average time taken for a charge carrier to

exit the active volume. This equation implies that current can be used to reach a desired photon emission rate in a given time. In reality this current is small in comparison to the current required to increase the voltage for the capacitance of the diode. A reverse current can be applied to sweep out remaining charge carriers in the pn-junction on a time scale shorter than τ_l to stop emission from the LED and hence shorten the fall time of the LED. Details on how this is achieved by the driver is discussed in Section 4.2.2.

An LED of resistance 9 Ω , in comparison to a typical LED resistance of 40 to 100 Ω , was chosen to increase the current across the LED. This differential resistance was measured in the near linear region of forward current against forward voltage above the LED threshold voltage. It was guaranteed by the manufacturer that all LEDs would be from the same batch.

4.2.2 Push-Pull Driver

The LED driver, based on the IXLDO2SI commericial chip, uses MOSFET transistors to first create a forward current across the LED and activates it. This is followed by a reverse current which sweeps charge carriers out of the active region of the diode and stop light emission. This method of switching currents gives the driver its Push-Pull name. The main components of the circuit are shown in Figure 4.7 with the rest of the circuit shaded out in grey boxes.

For the LED, D1, to emit light then the capicitor C within it must be first at threshold voltage. Once above threshold, then a fraction of the current across the LED will create a forward modulation current of a few mA across the pn junction of diode D and activate the diode. The capacitance of C is 200 pF, therefore a maximum forward current of 0.4 A leads to a maximum slew rate of dV/dt = I/C = 2 V/ns. This is fast enough to create a one nanosecond, bipolar current pulse through diode D and meet the requirements of optical pulses at a timescale of a few nanoseconds. The resistance of the LED, R_s, is 9 Ω .

The operating current through MOSFET transistor Q3 is controlled by the MOSFETs Q1 and Q2. Initially Q1 is open and Q2 is closed. When Q1 closes and Q2 opens a forward voltage is applied across the LED. When Q1 is returned to its open state and Q2 is closed then a temporary reverse current is created which will sweep charge carriers out of the LEDs active region and switch the LED off. The choice of the value of resistors R1 and R2 is a comprise between a low resistance required for a high voltage slew rate and a high resistance to limit the current through the transistors and create forward and reverse voltages across the diode. Taking these factors into account the values of R1 and R2 were



Figure 4.7: A circuit diagram of the main components of the Push-Pull driver and LED. D1 is the LED which is broken down into an equivalent circuit in the dashed box consisting of a series resistor R_s , diode D and capacitor C. The MOSFET transistors Q1 and Q2 are responsible for applying a forward and reverse current through MOSFET Q3 and across the D1. The resistors R1 and R2 were selected to maximise forward and reverse currents and slew rate across D1 whilst limiting the current through the transistors. The currents IOP and IPW set the LED pulse amplitude and widths respectively. IBI is a baseline current with respect to IPW. The LED is triggered via the TRIG input and PDN can power down the current across Q1 and Q2. The circuit is operated at a voltage $V_{CC} = 7$ V. The rest of the circuit is represented by grey boxes [95].

chosen to be 3.3 Ω .





The Push-Pull board contains the drivers for the MOSFETs and the circuitry that set the values of the currents IPW and IOP that control the pulse width and amplitude respectively. The current IBI is a baseline current with respect to IPW and compensates for intenal delays. IBI is fixed by the hardware for every Push-Pull board whereas IOP and IPW can be controlled for each board individually through software and communication with the control chip located in the control box. The LED pulses are triggered through the TRIG input and PDN can power down the current across Q1 and Q2 between pulses. No heat sink is required as the circuit is active for 200 ns per pulse. The maximum rate the circuit can operate is 10 kHz. The voltage at the common collector of the circuit is $V_{CC} = 7$ V.

To ensure that the driver-LED pair met the requirement of being able to produce pulses of light which have a FWHM of less than 6 ns, a single photon counting technique was used to record the time profile of the pulse [98]. The technique uses delayed coincidences between the trigger of the TELLIE system and the signal from a low background PMT which has a certified low dark-count rate [99]. Figure 4.8 shows a schematic of the experimental setup used. Before recording the time profile of the LED pulse, one must first confirm that the PMT is recording only single photoelectron hits. This is achieved by using neutral density filters to reduce the intensity of the incoming light pulse. The driver is set to a level at which the intensity of the LED pulses it generates is of the order of 10^3 photons per pulse. More neutral density filters are then added until the signal observed from a photoelectron generated in the PMT is comparable to its noise rate. The ADC spectrum is then taken which identifies that the PMT is detecting single photoelectron hits. Figure 4.9 shows an example of an ADC spectrum taken while the PMT is detecting single photoelectrons. The trigger from TELLIE is used as the signal to record data rather than the PMT signal. The first, larger, peak is when there is no incoming photon at the PMT at the point at which the trigger fired. This is referred to as a noise hit. The second broader peak is due to single photoelectron hits. Thus, to distinguish the single photoelectron hits from the noise hits while recording the time profile of an LED pulse, one uses delayed coincidences between the TELLIE trigger and the PMT signal. The time profile of light pulses with a higher number photons per pulse can be recorded by adding extra neutral density filters between the end of the optical fibre and the PMT.

Single photon counting is the most accurate method for acquiring the time profile of the light pulses. However, this method is extremely resource intensive and it proved infeasible to record the time profiles of all 92 drivers used in the TELLIE system using



Figure 4.9: A histogram of the ADC counts measured by the single photon counting set up in Figure 4.8. The value of the ADC count is proportional to the charge collected by the PMT which is detecting single photons. The PMT is triggered by the trigger signal generated by TELLIE when it is emitting light pulses. The peak at approximately 40 ADC is due to the trigger causing the readout of the PMT while the signal is absent. The shoulder is the charge collected upon detection of a single photon [98].

Table 4.2: The measured full width half maximum (FWHM), rise and fall time of light pulses at various intensities using the single photon counting method. All measurements made with the driver running at 1 kHz. The measurement of the number of photons is performed after 45.5 m of optical fibre and without a 2 m patch cable.

Number of Photons	Full Width Half Maximum (ns)	Rise Time (ns)	Fall Time (ns)
10^{3}	4.4	1.6	6.6
10^{4}	4.4	2.0	6.0
10^{5}	5.2	2.0	6.4
10^{6}	7.0	4.2	6.8

this technique. Only one of the drivers had its time profile recorded to confirm that the FWHM of the light pulses they generate is small enough to meet the requirements.

Table 4.2 gives the measured values which describe the time profiles of a single TELLIE fibre at various intensities. These measurements were taken using 45.5 m of optical fibre. When the TELLIE system is operating in SNO+, an additional 2 m of optical fibre is used to connect the ST connection of the channel on the box to the patch panel described in Section 4.2. The time profile of the TELLIE channel at an intensity of 10^3 photons per pulse before and after the full length optical fibre can be seen in Figure 4.10. The broadening of the time profile observed after the light is emitted through the optical fibre is due to modal dispersion within the fibre. Modal dispersion occurs due to different stable light transmission paths, or modes, in the optical fibre. A simplified model of modal dispersion can be viewed as different modes having different path lengths, see

64



Figure 4.10: Time profile of light pulses emitted by a TELLIE channel without (red) and with (black) transmission through 47.5 m of optical fibre. The same driver and settings were used for each measurement with its intensity set to 10^3 photons per pulse. The spectra are normalised to their amplitudes. The FWHM of the pulse after the fibre is low enough to meet the SNO+ requirements [100].



Figure 4.11: A simplified diagram of two different stable light transmission paths, or modes, within an optical fibre. The different path lengths of the modes causes different propagation times through the optical fibre and hence a broadening of the time profile of the initial light pulse. The effect of this broadening is known as modal dispersion.

Figure 4.11, and hence different total times to propagate from one end of the fibre to the other end, leading to a broadening of the time profile of the light pulse. In actuality, the electromagnetic field of the light in the fibre is more complex than this straight line approximation but the conclusions are the same. Despite this broadening the FWHM of the TELLIE channel after the optical fibre is well within the required 6 ns necessary to achieve a 1 ns precision in the timing calibration of the SNO+ PMT array over a 24 hour period with TELLIE pulsing at a rate of 10 Hz.



Figure 4.12: A projection of the PMT Support Structure (PSUP) and the location of installed fibres at the time of the commissioning runs. The black lines indicate the struts which make up the geodesic sphere. The different dashing and number of lines identify the different types of strut which vary in length and thickness. Between the struts are panels which host the PMTs. The struts are connected via three different types of hubs at 91 nodal positions indicated by numbers which provide the positions that the TELLIE fibres are mounted to. Each fibre is labelled with a fibre number which is the same as the nodal position it is connected to. The red circles indicate fibre positions which were installed via ladder in March 2012 while the cavity was not filled with water. The blue circles indicate fibre positions which were installed in October and November 2014. The upper blue circles had fibres installed by climbing on top of the PSUP where as the lower blue circles were installed via boat. The remaining fibres will be installed via boat as the water level increases in the cavity.



Figure 4.13: A drawing of a hex cell with an attached fibre mounting plate. A hex cell houses a PMT and its concentrator. The mounting plate are where the ends of two fibres are fixed into place [95].

4.2.3 Optical Fibre

The optical fibre was chosen to maximise illumination coverage area of the PMTs. It also must not add a significant amount of time to the rise time of the pulses. Polymethyl methacrylate (PMMA) step-index fibres, of model Mitsubishi ESKA SH4002 [101], were selected due to their 1 mm diameter and large 0.5 numerical aperture which maximises the illuminated area. The drawback of a large numerical aperture is that there is an increased amount of modal dispersion which will lead to longer time profiles. It was found that the rise time of light emitted from PMMA fibre is sufficient for the timing calibration.

A total of 110 duplex fibres were constructed for the TELLIE system. A duplex fibre consists of 2 fibres which improves the mechanical robustness of the cable and reduces the number of fibres to install on the PSUP. Out of the 110 duplex fibres constructed, 92 will be installed in the final system. One duplex fibre is installed at the top of the AV neck pointing downwards into the detector. The remaining 91 duplex fibres will be installed on the PSUP nodes as seen in Figure 4.12. A node is a hexagonal panel which connects the triangular faces of the geodesic geometry of the PSUP. On the triangular faces are panels which contain PMTs and their light concentrators which are housed in a hex cell. A double fibre connector is attached to a mounting plate, as seen in Figure 4.13, on the side of a

hex cell in the hexagonal panel located at a PSUP node. The hex cells at the PSUP nodes were chosen as the position of fibre installation as this will not shadow any PMTs and there is a wider gap between neighbouring PMTs which eases the installation of the fibres. Together, the PSUP node positions also conveniently provide complete optical coverage of the entire PMT array from the light emitted from all terminations of the optical fibres on the PSUP. The other end of the fibres have an ST connection at the dry end on the deck of the SNO+ detector.

Quality assurance tests were undertaken by the Lisbon group on all 220 fibres to ensure they meet timing, intensity and opening angle requirements. The least desirable fibre pairs were selected as spares and two fibres which have good parameters were kept for fibre ageing measurements. The timing measurements were made with a PMT connected to a 1 GHz oscilloscope which measured the rise time, fall time and Full Width Half Maximum (FWHM) of the time profile generated by a Push-Pull board and LED. The rise and fall times are defined as the difference in time at 10% and 90% of the peak amplitude on the rising and falling edges respectively. The Push-Pull board used was not the final version used for TELLIE but was very similar. A constant voltage LED source was used in conjunction with a power meter to initially and rapidly check the light intensity of the fibre¹.

Precise measurements of light intensity at different angles was made in the Lisbon lab using a bench top set up as seen in Figure 4.14. The bench top set up consists of a dark box which contains a PMT and a motor controlled table top which can hold up to three duplex fibres at a time. External to the dark box is a DAQ computer, an LED and its constant power supply, the PMT high voltage power supply, a picoammeter for measurement of the current from the PMT. The PMT is placed at a fixed 51 mm in the X direction away from the fibre termination points using the coordinate system in Figure 4.14.

The PMT was connected to a light guide with a $2.8 \times 15 \text{ mm}^2$ slit lying along the Y × Z axes. The data from the picoammeter or oscilloscope is collated by the DAQ computer depending on whether intensity or time profile measurements are being taken. The DAQ computer also controls the motorised table top where the fibres are fixed. The table top can move in the X and Y directions but the X position was kept constant. The Y position was moved in 1 mm steps. The LED, pulsed with a constant voltage, had been drilled and had a fibre attached to it as described in Section 4.2. This was to keep consistency with the light source that will be used in the final system. The LED was attached to the

¹This check was repeated once a fibre had been installed in the detector to ensure the fibre was not damaged during its installation into the detector.



Figure 4.14: A schematic of the bench top set up at Lisbon used for intensity measurements of light emitted from optical fibre at different angles. Three duplex fibres which each host two fibres can be mounted in the dark box at a time. The PMT, powered by a high voltage power supply, has its X position fixed at 51 mm away from the termination point of the optical fibre. A measurement of the current of the PMT was recorded by a DAQ computer connected to a picoammeter at 1 mm steps in the Y direction. The position of the fibres was changed through a motorised Table top which was controlled by the DAQ computer. A drilled LED with a fibre inserted was connected to each of the fibres individually via ST connection. The LED was powered by a power supply at a constant voltage.

fibres via an ST connection. Three duplex fibres each containing two fibres were mounted to the table top. Once a new set of fibres had been mounted the dark box was sealed and a background measurement taken. This allowed six fibre measurements to be taken sequentially for each high voltage power cycle. The intensity of light of the LED was set to a level that was greater than the background light level previously measured and the background reading was subtracted from the final readings.

The optical fibres were selected based on three criteria: the maximum output intensity of the optical fibre; the half width angle at 20% peak intensity which was measured in air and corrected to an angle in water; and the rise time of the light pulse. The results and empirical criteria of the quality assurance tests are shown in Table 4.3. One fibre failed the intensity criterion and the duplex fibre which contained it was excluded from installation. No fibres failed the angular criterion although two single fibres with values close to the limit were set aside as spares. Four single fibres had timing or intensity measurements which were close to the accepted limit; these were also selected to be spares.

69



Figure 4.15: The intensity of light emitted as a function of angle of emission, with respect to the centre of the fibre, in water. All measurements were made in air and had their angles corrected to a value in water taking into account the refractive indices of the media. The minimum to maximum range and RMS of the intensity measurements made using the Lisbon bench top set up in Figure 4.14 of all 220 constructed PMMA optical fibres is shown in the yellow and green respectively. The black line is a higher resolution measurement made at Sussex performed on one fibre using the coincidences techniques described in Figure 4.8 [95].

The variation in intensity of the 220 fibres as a function of angle in water as measured by the Lisbon bench top set up is shown in Figure 4.15. The yellow band is the minimum to maximum range of intensity of all the fibres whereas the green band is the RMS. The intensities were measured in air but the angle with respect to the centre of the fibre was converted into water by taking into account the refractive indices of the media. Also shown in the figure is an independent higher resolution measurement performed on a single fibre at Sussex which has a similar set up but uses coincidence techniques to identify light produced by TELLIE. As described in Section 4.2.2, the TELLIE control chip provided a trigger signal to indicate that TELLIE had produced a light pulse which allows the removal of background events. Modal dispersion of light within the fibre causes the three peak interference pattern that is observed in the high resolution measurements.

Table 4.3: The results and criteria of the optical fibre quality assurance tests. To pass the fibres must meet the denoted criteria for maximum intensity, I_{max} ; the half width angle at 20% of the peak intensity and rise time of the time profile of the fibre. The mean and RMS are the results of all 220 fibres constructed for the TELLIE system [95].

Characteristic	Mean and RMS	Criterion
I_{max} (nA)	120.2 ± 21.9	> 60
20% Angle (deg.)	14.5 ± 0.2	> 13.5
Rise time (ns)	1.9 ± 0.1	< 2.25



Figure 4.16: Left: A picture of a PIN diode readout board which consists of 4 channels. The PIN diodes are contained inside brass couplers that can just be seen at the top of the photo and are described in Section 4.2. The probes are attached to the sample points used for PIN diode channel calibration. The blue potentiometer below the top probe point adjusts the amplitude of the PIN diode pulse and the potentiometer to its left adjusts the time of sample point on the PIN diode pulse of that channel. Right: A picture of the scope readout of the two connected probes. CH1 (yellow) is the pulse shape of the PIN diode and CH2 (blue) is the sample point of the PIN diode pulse of that channel. The voltage of the pulse shape is measured at the point in time at which the sample point intersects the pulse shape. Note that the voltage scale of CH1 and CH2 are 0.05 V and 1 V respectively hence there is only one sample point which is on the rising edge of the pulse shape of the PIN diode.

4.2.4 PIN Diodes

To obtain a measurement of the intensity of light in each coupler PIN diodes are used. A PIN diode readout board hosts four independent PIN diodes. Each PIN diode channel consists of a Hamamatsu S5971 Si PIN photodiode inside the brass coupler, described in Section 4.2, and the circuitry hosted on the board which shapes the output of the PIN diode and samples it. The photodiode was chosen for its fast response of 100 MHz, low cost and high reliability [102]. The amplitude of the pulse shape of the channel can be adjusted using the blue potentiometer between the two probe points shown in Figure 4.16a and the sample point of pulse shape is adjusted by the potentiometer to the left of the amplitude potentiometer. The signal on the probes can be seen in the oscilloscope trace in Figure 4.16b. The voltage of the pulse shape, which is proportional to the amount of photons detected by the photodiode, is measured at the time of intersection between the PIN pulse and the sample point shown as CH1 (yellow) and CH2 (blue) in Figure 4.16b respectively. Note that the scale of the volts per division of CH1 and CH2 is 0.05 V and 1 V respectively. Therefore, the sample point shown in the figure occurs on the rising and not the falling edge of the pulse from the photodiode. The PIN diode readout board digitises the voltage measurement at the sample point as a 16 bit integer (0-65535) number with 0 being the minimum intensity that it is possible to measure. The control chip of the control box then reads out the integer number after each LED pulse and stores it on a buffer. An average and corresponding RMS of all the PIN readings on the buffer is then computed by the control chip at the end of a sequence. The average and RMS values are then read by the SNO+ DAQ and stored in a CouchDB database.

4.3 Software

The TELLIE hardware is manipulated through a control chip in the control box which sends and receives data to the 12 TELLIE boxes via a pair of ribbon cables. The communication with the control chip is handled by a python based piece of software which runs on the SNO+ DAQ machine, connected to TELLIE via USB. The parameters which can be adjusted through the software are

- Channel Number The TELLIE channel or channels selected to be fired.
- **Pulse Width (IPW)** Sets the amount of time the potential difference across the LED is applied for.
- Pulse Height (IOP) Sets the amplitude of the potential difference across the LED.
- Number of Pulses The number of pulses in a sequence before the PIN diode is read out and a new sequence can begin.
- Delay Between Pulses The amount of time between each pulse in a sequence.
- Fibre Delay A time delay applied to individual channels before initiating a sequence to account for channel to channel variations due to internal TELLIE electronic delays and the optical fibres. Different fibre delays can be applied simultaneously to multiple channels.



Figure 4.17: A diagram denoting the time delays due to light and signal propagation from when a TELLIE driver control board is triggered to when the SNO+ DAQ system registers an event [95]. The total delay is estimated to be approximately 615 ns. There is also a delay of approximately 70 ns between when the trigger signal is emitted by TELLIE and when it is detected by the SNO+ DAQ. The trigger delay should therefore by set to approximately 545 ns. A delay can be applied to when the trigger to the SNO+ DAQ is sent to register the TELLIE event. This can be set in steps of 5 ns between 0 and 1275 ns.

• Trigger Delay - A time delay of when the trigger signal is produced with respect to when each light pulse is emitted. The same time delay is applied to all channels and accounts for the amount of time between when the light is detected by the SNO+ DAQ system and when the trigger signal is received by the SNO+ DAQ so a global trigger can be produced and the event recorded. Figure 4.17 shows the various sources of light and signal propagation times within the TELLIE system and the SNO+ detector. The total length of time between light being emitted and the corresponding light to be detected by the SNO+ detector has been estimated to be approximately 615 ns. There is an estimated delay of 70 ns between when the trigger signal is emitted by TELLIE and when it is received by the SNO+ DAQ to generate the global trigger to record the event. Therefore, the approximate value the trigger delay should be set to is 545 ns.

The control chip manipulates 8-bit integers to assign values to the above parameters. The
	Parameter Name	Hi/Lo Range	Internal Calculation	Range
	Pulse Width (IPW)	Hi 0 - 63	$\mathrm{Hi}\times 256 + \mathrm{Lo}$	0 16383
		Lo 0 - 255		0 - 10303
	Pulse Height (IOP)	Hi 0 - 63	$Hi \times 256 + Lo$	0 - 16383
		Lo 0 - 255		0 10000
	Number of Pulses	Hi 0 - 255	$\mathrm{Hi} \times \mathrm{Lo}$	0 - 65025
		Lo 0 - 255		0 00020
	Delay Between Pulses	ms 0 - 255	$ms + 4 \times \mu s$	(0 - 256 02) ms
		$\mu { m s}$ 0 - 255		(0 200.02) 1115
	Fibre Delay	ns 0 - 255	$0.25 \times \mathrm{ns}$	(0 - 63.75) ns
	Trigger Delay	ns 0 - 255	5 imes ns	(0 - 1275) ns

Table 4.4: The range of different parameters that can be set for a TELLIE channel or channels for a sequence of pulses. Most parameters accept two 8-bit integer inputs, denoted in the Hi/Lo Range column, except Fibre and Trigger Delay which accept one. These inputs are used to calculated either an integer value or time which is set for the sequence. Although frequency can be technically set to be infinity, a hard coded limit of 10 kHz has been set so that the TELLIE channel operates safely.

range of values that the parameters can have and how the control chip calculates them from one or two 8-bit integers is shown in Table 4.4. Although the frequency of the pulses in a sequence can be technically be set to infinity, or as close to infinity as physically possible in reality, a maximum frequency of 10 kHz has been deemed safe for TELLIE to operate at and this limit has been encoded into the software.

There are two modes of running TELLIE: an automated mode and an expert mode. The automated mode is used during physics data taking. The channels will be flashed one at a time in an order which maximises optical coverage of the PMT array at a rate of 10 Hz with each channel having pre-set parameters optimised for PMT calibration. These parameters and pattern order are hosted on a CouchDB database which the python based software reads and writes to. The pattern order is decided according to which fibre is connected to which channel to maximise illumination coverage as a function of time. The expert mode is used for commissioning or specialised calibration runs with full customisation of each of the channel's parameters possible using the DAQ GUI. All parameters actually used during a TELLIE sequence are stored in the CouchDB database along with a temperature measurement of each box's motherboard and a PIN readout.

4.4 Characterisation and Quality Assurance of TELLIE channels

Before the TELLIE system was installed into the SNO+ detector it was important to check that the TELLIE channels met the PMT calibration requirements. These requirements are:

- The peak wavelength of each LED must be between 500 nm and 520 nm so that absorption of the light is minimised and PMT quantum efficiency is optimised and uniform across the system.
- The driver boards need to be able to emit light pulses which consist of 10³ photons per pulse or less so that PMT pulses collected for their calibration are resultant from single photoelectrons.
- The maximum photon output of the driver boards must be at least 10⁵ photons per pulse with 10⁶ photons per pulse being desirable. This is so multi-photoelectron events can also be studied.
- The PMT time offsets must be measured with a precision of 1 ns or less within a run time of up to 24 hours. To achieve this a TELLIE channel must have a time profile with a FWHM of less than 6 ns.

It is also important that the channel's output is also characterised for the operational range of IOP and IPW currents so the time profile and photon output is understood with the given parameter inputs. Details of the driver board characterisation can be found in Section 4.4.2.

4.4.1 LED Quality Assurance

The intensity spectrum as a function of wavelength was measured for 120 LEDs including the 96 used in the TELLIE system. This was done to check that all wavelengths had a peak wavelength between 500 and 520 nm so that transmission of LED through the active medium is maximised and the quantum efficiency of the PMTs is optimised, as shown in Figure 4.6, and ensure the peak wavelengths of the LEDs are uniform across the system. It also allows all LED spectra to be stored in a database used for simulating TELLIE runs.

The spectra of each LED was recorded with an Ocean Optics model Maya 2000 Pro spectrometer, operating with an integration time of 40 ms per scan and averaged over 20 scans. Each LED was connected in series with a 550 Ω resistor and was supplied a



Figure 4.18: A graph showing the mean and RMS of the wavelength dependent intensity spectrum of 120 LEDs. All spectra were normalised so that the integral of each spectrum is equal to one.

constant 2.48 V. Each LED was housed in its coupler and the connecting fibre was attached to 45.5 m PMMA fibre and a 0.2 m patch fibre was used to convert the ST connection to SMA for the input of the spectrometer.

All LEDs met the 500 to 520 nm peak intensity requirement. The minimum and maximum peak wavelengths observed were 500.0 and 510.1 nm. The mean peak wavelength of the LEDs tested was (505.5 ± 2.6) nm. The average and RMS of the spectra obtained can be seen in Figure 4.18 where all spectra were normalised such that their integral was equal to one. There was a measured factor of 19.7 difference between the LED with the minimum and the LED with the maximum number of photons detected by the spectrometer during its integration time. This was not viewed as a concern due to the ability of the driver to control pulse intensity through its IOP and IPW settings. If the minimum intensity of a TELLIE channel is deemed too high for PMT calibration then data obtained from these measurements and intensity measurements described in Section 4.4.2 can be used to match the driver with a dimmer LED. To date, no instance of this has been found.

4.4.2 Push-Pull Quality Assurance

This section describes how each driver was checked to ensure it met the requirements of the TELLIE system. First the requirements of the driver boards are described and then a description of the methodology used in the quality assurance tests is given. This is followed by details on how the PMT used to measure the light output was calibrated. The results of the IOP and IPW quality assurance tests are then given in the following sections.

Requirements

To ensure that driver boards met the requirements for PMT calibration and to characterise their behaviour as a function of IOP and IPW currents, each was subject to a series of tests. To pass the tests the boards need to be able to produce a minimum photon output of 10^3 photons per pulse or lower and a maximum photon output of 10^5 with 10^6 photons per pulse being desirable. Each Push-Pull board must also be able to produce pulses with their respective LEDs which will allow the time offset of each PMT to be measured to a precision of 1 ns. This must be achieved on a 24 hour time scale with the channels pulsing at a rate of 10 Hz, the rate at which PMT Calibration will be performed during physics data taking. This will correspond to a channel having a time profile with a FWHM of less than 6 ns with the channel operating at intensity of 10^3 photons per pulse.

Methodology

The first test was a basic functionality check to see whether the driver boards are able to produce LED pulses. Three out of 104 driver boards failed this test. The first step in characterising the driver boards was to examine the photon output of the driver boards as a function of IOP setting. This was primarily to check that no double pulsing is occurring at high IOP values which has been observed in earlier prototypes of the Push-Pull boards. The same reference LED was used for each board to ensure only driver boards behaviour was being examined. The reference LED was then also used to measure boards' behaviours as a function of IPW value so that the photon output of the driver boards is recorded independent of the LED it is coupled with. A higher resolution scan across IPW settings was performed with the complete channel after the PIN board had been tuned as described in Section 4.2.4. This was to measure the photon output per pulse and the time profile of the channel as a function of IPW. This data is also used to convert PIN readout into number of photons emitted after the optical fibre.



Figure 4.19: A flow diagram of the test bench set up used to characterise the light output of each TELLIE channel as a function of IOP and IPW values and to tune each channel's PIN board used to monitor each channel's light output during operation once installed in SNO+.

A flow diagram of the bench top set up for the driver characterisation is shown in Figure 4.19. One TELLIE box was tested at a time, this was powered through the TELLIE power supply whose voltage output had been tuned down to account for fewer boxes being connected. The control box relays parameter inputs to the TELLIE box and reads out the channel's PIN. A DAQ computer communicates with the control box so users can adjust parameter settings and the PIN readout can be recorded. The light produced from a channel was transmitted through 47.5 m of PMMA optical fibre as this is the length used when installed in the SNO+ detector. The light output from the fibre was measured by a Hamamatsu H10721P-110 PMT which was operated at a voltage of 0.6 V or 0.8 V depending on the intensity of the light output being injected into the fibre. The voltage was monitored with a multimeter at the output of the power supply. The integrated value, FWHM, rise and fall times of the time profile of the pulse was measured with a Tektronix DPO 3054 oscilloscope, labelled as oscilloscope 2 in the flow diagram, averaging over 16 waveforms and readout by the DAQ. The PIN board was initially tuned before a full channel characterisation using a Tektronix DPO 3054 oscilloscope, labelled as oscilloscope 1 in the flow diagram, to set the PIN diode's pulse amplitude and sampling point. The PIN boards were tuned to maximise the range of photons per pulse they are sensitive to. A PIN calibration scan was performed to ensure a PIN diode had been tuned correctly before a full channel characterisation. This was a scan with IOP maximally set and IPW raised in steps of 100 from 0 until no light is produced. The readout from the PIN diode was then plotted as a function of photons per pulse. If the PIN diode was returning a minimal value whilst light was still being emitted from the fibre or if a maximal value was returned at a low intensity then the PIN diode readout board was retuned and another PIN calibration scan was performed. The PIN diode readout boards were tuned such that the range of intensities they are sensitive to was maximised whilst ensuring sensitivity at the setting which produces the fewest number of photons per pulse.

PMT Calibration

The PMT response was characterised using a ThorLabs PM100USB light-meter. Using a TELLIE channel, the number of photons per pulse was measured by taking an average of 5 readings of two second integrations with the channel pulsing at a rate of 1 kHz. The number of photons per pulse obtained from these measurements was then compared with the integrated signal from the PMT to extract the gain of the PMT. Seven applied voltages on the PMT were used ranging from 0.56 V to 0.85 V. The relationships between the applied PMT voltage and the resultant gain were consistent between two channel intensity settings.

The number of photons per pulse, N, can then be calculated through integrating the measured pulse shape

$$N = \frac{\int V dt}{R \cdot e \cdot G} \cdot QE, \tag{4.8}$$

where $\int V dt$ is the integrated pulse shape, R is the 50 Ω termination resistance at the oscilloscope, e is the charge of the electron and G is the gain of the PMT. The quantum efficiency, QE, of the PMT was calculated to be

$$QE = \frac{SE_{\gamma}}{e} = \frac{Shc}{e\lambda} = 19.1\%, \tag{4.9}$$

using the cathode radiant sensitivity, S, shown in Figure 4.20, which is given as 77.4 mAW⁻¹ at 503 nm in the Hamamatsu data sheet [103]. The energy of the incident photon, E_{γ} , is calculated in terms of Planck's constant, h, the speed of light, c and the wavelength of the incident photon $\lambda = 503$ nm.

IOP Optimisation

To ensure the IOP is set such that the photon output of the channel is maximised the number of photons was measured as a function of IOP value ranging between 10583 and 16383 in steps of 100. With the IOP maximally set to produce the highest number of photons per pulse, the IPW setting is primarily used to control pulse intensity whilst



Figure 4.20: Cathode radiant sensitivity of the PMT [99] as a function of the wavelength of the incident photon [103]. The solid and dashed lines are the sensitivities of the -110 and -210 models of the PMT respectively. The -110 model was used for the quality assurance measurements.

the option to reduce intensity further is available by lowering the IOP further. A visual inspection of the pulse shape at its maximal setting is also used to ensure that only a single peak is observed in the time profile of the pulse. In previous prototypes of the driver board, two peaks have been observed in the time profile of the pulse.

All channels had a photon output that increased as a function of IOP number and only single peaks in the time profiles were observed. Figure 4.21 shows the average and range of photon output of all channels as a function IOP. All channels are set to the maximal IOP setting of 16383.

IPW Characterisation

For a given IPW setting the number of photons per pulse, PIN readout and the full width half maximum (FWHM), rise and fall times of the pulse shape were recorded. The rise and fall times are defined as the difference in time between 10% and 90% of the peak amplitude of pulse on the rising and falling edges of the pulse shape. These measurements were made primarily to verify that the channels used in TELLIE met the timing and



Figure 4.21: A graph showing the average, standard deviation and range of number of photons per pulse as a function of the pulse height parameter IOP of all 96 driver boards used in the TELLIE system. The IPW parameter was set to 0 to produce the highest number of photons per pulse.

photon intensity requirements described in Section 4.4.2 and ensure consistency amongst the drivers. This data is also used during TELLIE calibration runs to allow users to select the appropriate IPW setting and achieve the desired photon intensity. The PIN readings are used to monitor the photons per pulse over the course of a calibration run ensuring that the light being injected into the detector is stable.

The mean minimum photon output of the TELLIE channels was 232 photons per pulse and has a range of 20 to 510 photons per pulse with IOP set to the maximal value of 16383 and IPW set to its lowest light producing value. The mean maximum photon output of the TELLIE channels was 1.13×10^6 photons per pulse and has a range of 0.45×10^6 to 1.85×10^6 . Therefore all driver boards met the required photon intensity range of 10^3 to 10^5 photons per pulse. A total of 27 driver boards had less than the desirable 10^6 photons per pulse.

Figure 4.22 shows the minimum and maximum range and the mean and rms of the FWHM of all 96 boards. At an intensity of 10^3 photons per pulse the range of the FWHM of the driver boards is 1.5 ns. As shown in Table 4.1, only pulses with a photon intensity of the order 10^3 photons per pulse have the requirement that their FWHM is 6 ns or less to achieve a total run time of less than 24 hours. Higher intensity pulses have a



Figure 4.22: The mean, standard deviation and range of the full width half maximum (FWHM) of the time profiles of the photon pulses of all 96 channels used in TELLIE as a function of the number of photons per pulse. The FWHM was measured by the experimental set up described in Figure 4.19. The IOP value was set to its maximum value of 16383.



Figure 4.23: The mean, standard deviation and range of the rise times of the time profiles of the photon pulses of all 96 channels used in TELLIE as a function of the number of photons per pulse. The rise times were measured by the experimental set up described in Figure 4.19. The IOP value was set to its maximum value of 16383.

shorter run time due to a greater probability that an incoming photon will cause a PMT hit meaning that a fewer number of pulses are required to achieve enough statistics for PMT calibration. Due to systematic effects arising due to PMT and oscilloscope response times and the oscilloscopes sampling rate, the recorded time profiles were more broad than the measurements recorded using the single photon counting technique described in Section 4.2.2. Due to the limitations that arise due to systematics from the oscilloscope and PMT response time, the true range of FWHM at this intensity will be less than this value. While this is undesirable these systematic effects were the same for all channels and thus the results can still be used to check for consistency in the time profiles of all the TELLIE channels. However, if we neglect these systematic effects and assume the true range of FWHM of the time profiles of light pulses at an intensity of 10^3 photons per pulse is 1.5 ns then if one assumes that driver who had its FWHM measured at 4.2 ns using the single photon counting technique in Figure 4.10 has the shortest FWHM of all 96 drivers used in TELLIE then all drivers will still meet the required less than 6 ns FWHM at 10^3 photons per pulse as the driver with the widest pulse shape will have a FWHM of 5.7 ns. Therefore, this measurement is accurate enough that we can confidently say that all driver boards meet the FWHM requirements of the quality assurance tests.

Similarly the rise and fall time was also measured by the oscilloscope to ensure consistency between all the channels used in the TELLIE system. As shown in Figure 4.23, the range of the rise times of the channels is less than 2.2 ns at 10^3 photons per pulse. The range of fall times at 10^3 photons per pulse is less than 3.7 ns, Figure 4.24. From these conservative measurements of the spread of the time profiles of the TELLIE channels, all channels had a consistent enough time profile to be used within the TELLIE system.

4.5 TELLIE commissioning

The TELLIE system was installed at SNOLab in August 2013. Two major commissioning periods have taken place since the installation. The water levels and main purposes of the commissioning periods were:

- February/March 2014 The water level in the cavity was below the AV for initial tests of the stability of submerged PMTs. These tests are designed to ensure we do not see high failure rates of PMTs once the detector has been filled with water and switched on at high voltage for the initial SNO+ water phase.
- December 2014 The water level in the cavity at 6.77 m relative to the cavity floor,



Figure 4.24: The mean, standard deviation and range of the fall times of the time profiles of the photon pulses of all 96 channels used in TELLIE as a function of the number of photons per pulse. The fall times were measured by the experimental set up described in Figure 4.19. The IOP value was set to its maximum value of 16383.

inside the AV at 1.28 m relative to the bottom of the AV. The effect of the higher water level outside the AV in the cavity to the water level inside the AV is to mimic the buoyancy the AV will experience when it is filled with the liquid scintillator that is less dense than water. This allows the hold down rope net system to be monitored to check the load distribution is uniform and that the net is not slipping. It is also allows the AV position to be monitored to ensure minimal displacement and tilt.

The extended periods of high voltage running allowed an opportunity to commission the TELLIE and scattering module (SMELLIE) subsystems of the optical calibration system in parallel with achieving the above key experimental milestones.

At the time of the February and March commissioning periods, 32 out of the 92 nodes had TELLIE fibres installed. The positions of the installed fibres are denoted by red circles in Figure 4.12. These fibres were installed by ladder in March 2012 when there was no water inside the cavity. In October and November 2014 the nodes denoted by blue circles in Figure 4.12 had TELLIE fibres installed. The nodes at the top of the PSUP were installed by climbing on top of the structure. The other nodes were installed via boat. The remaining nodes will also be installed via boat as the water level rises inside the cavity. At the time of the December commissioning period, 59 out of the 92 nodes had TELLIE fibres installed.

4.5.1 Goals

The goals of the TELLIE commissioning period were:

- Stability of TELLIE The intensity of light injected into the detector will be monitored over the course of the data taking. This is to ensure that when TELLIE is used in a mode which requires single photoelectron hits at the PMTs, such as a PMT calibration run, that the intensity does not increase to an amount that would lead to multi-photoelectron hits at the PMTs leading to a bad calibration of the PMT array. The delay between the emission of light pulses will also be monitored.
- PMT Calibration As only 32 out of 92 nodes had TELLIE fibres installed during the February and March 2014 commissioning runs, not all PMTs will be illuminated due to the 14.5° direct light apertures of the fibres and a complete calibration of the PMT array is not possible. However, data from PMTs which obtain enough hits were used to test the PMT calibration code and performance standards established for SNO were used to indicate whether a PMT passed or failed. The partially filled AV and the different water level outside the AV in the cavity during the December 2014 commissioning period meant that this data was an unsuitable test of the PMT calibration code due to water level refractions and reflections.
- Test of the SNO+ DAQ system The trigger efficiency will be measured to ensure all TELLIE events are tagged by the appropriate trigger bit and no TELLIE events can be mistaken for physics data.
- Gated charge spectrum To measure a true single photoelectron spectrum with a SNO+ PMT the PMT must be triggered before any TELLIE light interacts with the PMT. This removes the need for the PMT to collect enough charge to cross its threshold before it is registered as a hit and will cross check that the thresholds have been set correctly and the PMTs are recording single photoelectron and not multi-photoelectron hits. To achieve this the external asynchronous (EXTA) trigger input, which receives triggers from TELLIE to trigger the detector to record the TELLIE events, is split in three directions and fed into the detector. Two short cables are used as input into an external trigger known as EXT5 and the external pedestal trigger (EXTPED). The EXT5 trigger is used to trigger the global trigger and record the events. The EXTPED trigger is used to force specified crates to read out. A longer cable is used to delay the incoming trigger signal from TELLIE to the EXTA input so that the light produced by TELLIE arrives at the PMTs located in



Figure 4.25: A two dimensional histogram showing the number of hits that a PMT registers during a TELLIE run. The PMTs have been binned as a function of distance from the emission point of the optical fibre within the detector and the time at which a PMT hit occurs. Data was taken during the February commissioning runs when the AV contained only air and only a fraction of PMTs below the AV were submerged in water.

the crate after the crates have been forced to begin reading out by the EXTPED trigger signal.

4.5.2 Transition of TELLIE Light in SNO+

Many different reflections can occur during the transit of TELLIE light in SNO+. These reflections can be observed by viewing the TELLIE events in space and time. To examine the different regions in space and time the PMT hits of TELLIE runs can be plotted in these two dimensions.

The time and spatial information for a standard TELLIE run can be seen in Figure 4.25. This data was taken during the February commissioning where the AV only contained air and only a small fraction of PMTs located below the AV were submerged. It can be seen here that there is an initial reflection off the AV which illuminate PMTs which are close to the fibre position. This data may be able to be used for determining the position of the AV [100]. The majority of TELLIE light is detected directly opposite the fibre. It is this light that is used to calibrate the PMTs. After the direct light has been detected, PMTs across the detector are illuminated. This is due to 35° reflections off the PMT, into the light concentrators that surround the PMTs that then reflect the light outward into the



Figure 4.26: A two dimensional histogram showing the number of hits that a PMT registers during a TELLIE run. The PMTs have been binned as a function of distance from the emission point of the optical fibre within the detector and the time at which a PMT hit occurs. Data was taken during the December commissioning runs when the cavity and the AV was partially filled with water. The data points highlighted in the red circle occur due to initial reflections off the side of the AV closest to the fibre. These data points are used for the AV location code. The data points highlighted in black is due to the direct light illuminating PMTs at the opposite side of the detector. These data points are used for PMT calibration. The data points highlighted in the dark blue ellipse are mainly due to reflections off the PMT concentrators. There is a peak in reflections off the concentrators when light is entering the PMT at a 35° incident angle.

detector [104]. A number of photons also reflect off the AV side opposite the fibre position and are detected again near the original fibre position.

For comparison, Figure 4.26 shows the various reflections, refractions and scattering that occurs to the light that is emitted from the optical fibre during the December commissioning runs. This is when the AV is filled partially with water. The different regions of reflections are shown in this plot. The data points highlighted in the red ellipse are used for positioning the AV [100]. Also visible the water level reflections highlighted in the blue ellipse in Figure 4.26 are not observed in 4.25. The light highlighted in the black ellipse is used to calibrate the PMTs. The light highlighted in the white ellipse is from the PMT concentrator reflections.



Figure 4.27: A profile histogram of R_{NHit} , defined in equation 4.10, as a function of the time of the event relative to the first triggered TELLIE event of the run. The plot shows the average intensity rise of 30 PMT calibration runs from the March 2014 commissioning period. The runs were taken at 500 Hz and are binned in 5 second bins. The average initial NHit of the runs was 52.7 ± 11.7 . The black line shows the result of a linear fit to the data. The gradient of the linear fit is $6.0 \pm 0.1 \times 10^{-5} \text{ s}^{-1}$ and its intercept is 0.0545 ± 0.0002 . The reduced chi-squared of the fit is 5.99.

4.5.3 Results

Stability of TELLIE

To monitor the relative change in intensity as a function of time of TELLIE light pulses a variable R_{NHit} is defined as

$$R_{NHit} = \frac{N_{Hit} - N_{Hit}^{Init}}{N_{Hit}^{Init}},\tag{4.10}$$

where N_{Hit} is the number of PMTs that are hit by a particular light pulse and N_{Hit}^{Init} is the average N_{Hit} of the first 30 TELLIE light pulses of a run. Figure 4.27 shows how R_{NHit} increases as a function of time since the beginning of the run. The results are an average over 30 PMT calibration runs taken in March 2014. All runs were taken with a pulse delay set to 2 ms (500 Hz). The average initial NHit of the runs was 52.7 ± 11.7 . The data was fit with a linear function. The parameters of the linear fit are a gradient of $6.0 \pm 0.1 \times 10^{-5}$ s⁻¹ and intercept of 0.0545 ± 0.0002 . The value of the intercept is a measure of the initial intensity rise due to the channel initially switching on. The cause of the intensity rise is thought to be a combination of the driver and LED heating as a



Figure 4.28: A histogram of R_{NHit} , defined in equation 4.10, of all 30 PMT calibration runs taken at 500 Hz from the March 2014 commissioning period. The average initial NHit of the runs was 52.7 ± 11.7 . The total number of events recorded in the runs was 4.69×10^6 . The mean value of R_{NHit} of all events is 0.065 and the value of the RMS is 0.246.

channel produces light pulses. As an LED increases in temperature then a larger amount of energy is transferred through interactions of phonons with the charge carriers within the semiconductor of the LED. This leads to a greater light output than one would expect than just the contribution from electrical power alone as the recombination rate of charge carriers increases.

The intensity rise is not obviously apparent in individual runs due to a wide distribution of R_{NHit} . The distribution of R_{NHit} is shown for all 30 runs in Figure 4.28. The RMS of the distribution is 0.246. When this value is compared to the intensity increase of $R_{NHit} =$ 0.082 after 200,000 light pulses it is clear that the intensity rise due to a temperature increase within the TELLIE channel is a sub-dominant effect in comparison to the spread in intensity of all light pulses in a TELLIE run. Due to the circuit being active for less time, it is also expected that the lower pulsing rate of 10 Hz for a standard PMT calibration run will lead to a lower temperature rise of the TELLIE channel and hence lower rise in intensity. The overall intensity rise during a run can also be reduced by simply lowering the number of pulses a channel emits before switching channels.



Figure 4.29: The black points show the value of ECA calibrated time versus the measured QHS at PMT 1851. The red points show the interpolation points used in the PMT calibration code. The first six interpolation points with lower values of QHS show timewalk at the PMT whereas the four interpolation points higher in QHS show the offset due to the PMT cable.

PMT Calibration

The February 2014 commissioning run was used to test the PMT calibration code. At the time of the commissioning run 32 out of 92 duplex fibres had been installed as shown in Figure 4.12. Due to time constraints and hardware issues, 18 of the duplex fibres were used for PMT calibration. As a result there was not enough optical coverage to calibrate the entire PMT array. Although more fibres were installed at the time of the December 2014 commissioning run, these runs were unsuitable for PMT calibration due to water level refractions and reflections.

The timing component of the PMT calibration code accounts for PMT cable delays and the time walk effect described in Section 4.1. The direct light is selected to perform the timing calibration. This is defined as the 8.5 ns window with the largest integral. The value of 8.5 ns was used for SNO PMT Calibration and is based on a 5 ns time walk and a laserball pulse width of 2.5 ns. Data occuring within 10 ns of the peak time is then used for the analysis. The region of charge used for the analysis is found by first selecting the smallest charge range which contains 70% of the data. This range is then extended to contain 90% of the data with the upper limit restricted to be no more than 50 counts above pedestal (cap) above the previous upper limit to ensure the low charge time walk



Figure 4.30: The results of the gain fit of the PMT calibration code for PMT 2790. The peak, denoted with a blue line, is the centre of a sliding window which contains the most hits. The high half point, denoted with a black line, is the centre of the sliding window, moving to the right, when its integral is equal to half the integral at the position of the peak value. The threshold is the centre of a differently sized sliding window, moving to the left, which contains half the entries of the peak window.

effect is included in the data set. Once this charge region has been selected then it is split into sections used for the interpolation points. There are three widths used for the sections which increase with increasing charge. The width of the sections are user defined with the intermediate sections being twice as large as the low charge sections. For each section the interpolation points are calculated. The x value is the average charge for that section and the y value is the median of the time values. The last four points are linearly fit to extract the PMT cable delay. Cuts are then used to decide on whether a PMT channel passes or fails based on the gradient of the fit, the RMS of the interpolation points and the time between consecutive points known as **Tstep**. Figure 4.29 shows an example of the timing calibration for one PMT channel. Two further cuts are made based on the fraction of hits outside the data used for the analysis. **Fout** is the fraction of hits that occur before the data region. These hits are mainly due to noise. **Flate** is the fraction of hits that occur after the data region. A less stringent cut is placed on this due to scattered and reflected light.

The gain fit of the PMT calibration code extracts three parameters from both the QHS and QHL charge spectra from each PMT. These parameters are extracted using a sliding window technique described in detail here [105]. The three parameters extracted for each charge spectra are

- Peak The centre of the window which contains the most hits. If the peak occurs at 5 cap or less than this is identified as a noise peak and a peak finding algorithm in ROOT is used to find new peak candidates and the sliding window technique is repeated.
- High Half Point The centre of the window where the integral is half the integral of peak window.
- Threshold Depending on the the width of the peak window, a window of size one, two or three cap is used to find the point at which the number of entries within the window is half or less than the number of entries at the peak position.

These parameters are then compared to the parameters obtained from the previous gain fit to identify whether a PMT channel passes or fails the current gain fit. An example of a PMT channel which passes the gain fit calibration is shown in Figure 4.30.

The PMT Calibration code requires that a PMT channel has at least 1000 hits for the gain fit and 1000 hits and at least 10 events in one interpolation bin to perform the time walk correction. A total of 3224 channels had enough statistics for the gain fit of which 2633 passed SNO gain fit cuts and 3164 channels had enough statistics to perform the time walk correction of which 2771 channels passed SNO cuts. Tables 4.5 and 4.6 contain the cuts and the values used for the calibration.

A pass rate of 81.7% for the gain fit calibration and 87.6% for the timing calibration shows a successful demonstration of using TELLIE to calibrate the SNO+ PMT array. This is compounded by the fact that old SNO PMT Calibration constants were used for the gain fit calibration. It is therefore to be expected that the parameters obtained using the gain fit will have changed over a period of over 8 years. Also it is expected that the pass rate of the timing calibration will improve once all fibres are installed so that optical coverage is maximised and the laserball data has been taken so delays arising from variations in TELLIE channels can be accounted for.

Test of the SNO+ DAQ System

The commissioning runs using the TELLIE system were used to check the efficiency of the SNO+ trigger system. The number of trigger signals sent by TELLIE in a run is known as the number of pulses, N_{pulses} , and is saved to a database. This can be then compared

Cut Description	Criterion	
Low Channel Occupancy	NHits of channel < 1000	
RMS of time of events	> 5.0 ns	
RMS of time of events	> 3.5 ns	
High charge points were not fitted		
Tstep warning	> 1.5 ns	
Flate warning	> 0.4	
Fout warning	> 0.1	
Fitted gradient too big	> 0	
Fitted gradient too small	> -0.1	
Too many interpolation bins failed	> 2	

Table 4.5: The cuts used in the timewalk PMT calibration and the criteria required for a PMT channel to fail the cut.

Table 4.6: The cuts used in the gain fit PMT calibration and the criteria required for a PMT channel to fail the cut.

Cut Description	Criterion	
Low Channel Occupancy	NHits of channel < 1000	
Too little hits in 100-bin QHS or QHL window	NHits of window > 100	
QHS Threshold (QHSTH) too high	> 60 cap	
QHL Threshold (QHLTH) too high	> 60 cap	
QHSTH too low	< 0 cap	
QHLTH too low	< 0 cap	
QHSTH difference too large	$ \text{QHSTH} - \text{QHSTH}_{prev} > 10 \text{ cap}$	
QHLTH difference too large	$ \text{QHLTH} - \text{QHLTH}_{prev} > 10 \text{ cap}$	
QHS Peak (QHSPK) difference too large	$ \text{QHSPK} - \text{QHSPK}_{prev} > 10 \text{ cap}$	
QHL Peak (QHLPK) difference too large	$ \text{QHLPK} - \text{QHLPK}_{prev} > 10 \text{ cap}$	
QHS High Half Point (QHSHP) difference too large	$ \text{QHSHP} - \text{QHSHP}_{prev} > 10 \text{ cap}$	
QHL High Half Point (QHLHP) difference too large	$ \text{QHLHP} - \text{QHLHP}_{prev} > 10 \text{ cap}$	

to the number of EXT-ASYNC triggers during a run, N_{EXTA} , to calculate an efficiency ϵ

$$\epsilon = \frac{N_{EXTA}}{N_{pulses}}.$$
(4.11)

To examine the trigger efficiency, 50 PMT calibration runs were used each consisted of 2×10^5 pulses emitted by TELLIE and trigger signals sent from TELLIE to the MTC/D. The pulse delay of the runs was set to 1 ms however the true pulse delay is approximately 1.21 ms due to signal propagation time in the electronics of TELLIE. The trigger efficiency of the runs is $\epsilon = 0.999935$.

It is important to identify any events that contain TELLIE light so they are not misidentified as a physics event. To examine why TELLIE trigger signals are missing a window is identified between sequential TELLIE triggers in a sequence of pulses if the difference in time between the two TELLIE triggers is greater than 1.6 times the true pulse delay. The number of events that are suspected to be missing in the window, N_{miss} , can be calculated using

$$N_{miss} = ||(t_i - t_{i-1})/\Delta t|| - 1$$
(4.12)

where t_i and t_{i-1} denotes the times, from the 50 MHz clock, of the current and previous events containing a TELLIE trigger bit and Δt is the true delay between pulses. A total of 653 suspected missing TELLIE events were found in all runs using this technique which is equivalent to the total missing observed in the runs. This confirms that the windows are correctly identifying periods in time where a certain number of TELLIE events occurred but are either not tagged as TELLIE events or not stored. These windows can therefore be removed from the physics analysis so that TELLIE events are not misidentified as physics events.

The DAQ system test identified issues within the SNO+ DAQ with the data-stream, how data flows from the PMT and trigger output to the resultant data files, and the event builder, the software which constructs individual events in the data files. A significant effort has been made within the software and hardware to resolve the issues which result in missing triggers. This is yet to be tested and TELLIE can be used to investigate whether this issue has been resolved or not.

The types of trigger found in these windows are listed in Table 4.7. Note that the total number of events or orphans or windows with no events or orphans is greater than the suspected missing events. This is due to some windows containing one or more event.

Trigger Type	Number of Triggers
No triggered events or orphans	357
Orphan	48
ESUMHI	284
NHIT_20+NHIT_20_LB	82
PULSE_GT	20
NHIT_100_LO	14
OWLE_LO	1
NHIT_20_LB	1
SYNC	1
NHIT_20+NHIT_20_LB+NHIT_100_LO	1
Total	819

Table 4.7: The types and number of triggers recorded when a TELLIE event has been missed. Note that Orphan is not a recorded trigger type but denotes when the builder has failed to store an event in the data file. PULSE_GT is a global trigger generated by a pulser within the SNO+ trigger system. SYNC is a trigger signal which is used to synchronise the DAQ system with the GPS system on surface. Other trigger types are defined in Section 3.2.



Figure 4.31: A histogram showing the difference in time of a non TELLIE event occurring and the previous TELLIE event during a time window where there is a suspected missing event during a total of 50 PMT calibration runs. A cut has been made such that only windows which contain one non TELLIE event are shown. The spike observed in the distribution occurs at a time when one would expect the missing TELLIE event. The true pulse delay of the runs was approximately 1.21 ms.

Orphans occur when the event builder fails to store an event in the data file. These are identified by looking for an increment in the global trigger ID of more than one between sequential events. Orphans are a known problem in SNO+ and an effort is ongoing to stop them from occurring. PULSE_GT is a global trigger that is generated by a pulser located in the SNO+ trigger system. The SYNC trigger is used to synchronise the DAQ system with a GPS system located on surface. The other trigger types in the table are described in Section 3.2. One can reduce dead time further by attempting to tag any events which contain light emitted from TELLIE within these time windows thus saving any physics events which also occur during this window. Figure 4.31 shows the time distribution of non TELLIE events occurring during a time window where a TELLIE event is suspected to be missing. A cut has been made so only windows where one event occurs are shown. There is a clear spike in the distribution at 1.21 ms which corresponds to the true pulse delay of the TELLIE runs. It is therefore likely that these events are associated with TELLIE light and can be tagged as such. There is a 20 ns lockout time from when a trigger signal is received by the SNO+ DAQ when no other trigger signals can latch to the event. It is therefore likely that these events were followed by an EXTA trigger signal within 20 ns therefore the trigger bit associated with TELLIE will have been lost and thus the event



Figure 4.32: Screenshots of the SNO+ event display of the gated charge spectrum run. The two displays are of the summation of TAC counts during the run.

will appear to be missing. A total of 275 events were found in a 40 μ s window centred around 1.21 ms. This therefore means that potentially 42% of TELLIE events that have been missed can be tagged and removed from the physics data set.

Gated charge spectrum

In order to obtain a gated charge spectrum the data from a crate must begin to be read out before any light pulses reach the PMTs within the crate. Trigger signals known as pedestals are used to force a crate to have its data read out. Figure 4.32 shows screenshots of the SNO+ event display during the gated charge run. The left screenshot shows the three dimensional view of the SNO+ detector. The color represents the value of the TAC count as shown in the right screenshot where time is increasing from right to left on the x-axis. The left screenshot shows how the PMTs connected to a crate are arranged in vertical segments around the detector. The crate which was chosen to be readout was crate 5 and fibre 62 was used to emit the light pulses from as it is located on the opposite side of the detector to the crate.

Figure 4.33 shows the time profile of the gated charge spectrum. The larger initial peak of the time profile is from the pedestal triggers forcing crate 5 to readout. The following three peaks are due to PMTs, which are not located in crate 5, registering a hit from the light emitted from the TELLIE fibre. The difference in time between the pedestals and the emitted TELLIE light is greater than 70 ns which is the amount of time that a QHS spectrum is integrated over. Therefore only the QHL spectra which are integrated over 400 ns can be used to compare the gated and ungated charge spectra.

The complete QHL charge spectra of PMT 2779 are shown in Figure 4.34 where the black line denotes the gated charge spectrum and the red line is the ungated charge

96



Figure 4.33: A histogram of the time at which PMTs were triggered during the gated charge spectrum run. The first peak corresponds to the PMTs in crate 5 receiving pedestal triggers. The following three peaks corresponds to the time profile of TELLIE light.



Figure 4.34: A histogram of the charge collected on PMT 2779 during the gated (black) and ungated (red) charge spectrum runs. The large initial peak in the gated charge spectrum run is the charge collected on the PMT due to noise when no photoelectron is produced. Both histograms are normalised so there integrals are equal to 1.

spectrum. Both histograms are normalised so their integrals are equal to one. The peak in the gated charge spectrum is due to the PMTs being readout when no photoelectron has been detected which is also known as a noise hit. To compare the shape of the histogram



Figure 4.35: A histogram of the charge collected on PMT 2779 during the gated (black) and ungated (red) charge spectrum runs. The histogram is integer binned from 20 to 200 cap to remove the noise peak in Figure 4.34 Both histograms are normalised so there integrals are equal to 1.

to ensure the gated and ungated charge spectra obtained are both single photoelectron spectra a cut value to remove hits below 20 cap is placed and the resultant histograms were renormalised as before. Figure 4.35 shows the renormalised histograms after this cut has been placed. The similar distributions observed in the figure confirm that a similar intensity of light emitted from TELLIE was used.

The PMT Calibration code was then used to extract the threshold values of the QHL charge spectrum of the PMTs in crate 5 obtained in the gated and ungated charge spectrum runs. If a PMT passed the gain fit calibration in the gated and ungated runs then the difference between the thresholds obtained was then calculated using

$$\Delta T = T_{gated} - T_{ungated} \tag{4.13}$$

where T_{gated} and $T_{ungated}$ are the thresholds obtained using the PMT calibration code for the gated and ungated runs. The difference in thresholds are shown in Figure 4.36. A total of 62 PMTs located in crate 5 passed the gain fit in both the gated and ungated charge spectrum runs. The mean difference of the thresholds between the two types of runs is -14.3 ± 0.9 cap. The average threshold of the gated charge spectra is 29 ± 1 cap and the average threshold of the ungated charge spectra in crate 5 is 42.6 ± 0.9 cap.



Figure 4.36: A histogram of the difference between the thresholds, ΔT , obtained from the gain fit PMT Calibration code for gated and ungated charge spectra.

As previously stated, the gain fit calibration code extracts the threshold by finding the point at which the number of entries in the sliding window is equal to half that at the peak of the distribution. It is therefore expected that lower thresholds will be obtained from the gated charge spectra as there is a greater number of noise hits in the data set. However the data set clearly highlights how the discriminator thresholds of the PMTs introduce a bias to higher threshold values when using the current methodology of the gain fit.

4.5.4 Conclusion

The electronics of TELLIE have been successfully installed at the SNO+ detector and it has been demonstrated to be useful for various purposes. This includes its primary purpose of calibrating the PMT array where it achieved a pass rate of 81.7% for the gain fit and 87.6% for the timing calibration when 18 of the 92 nodes were utilised, old SNO PMT Calibration constants were used and the delays between channels have not been corrected for. The data collected by TELLIE has been used to calculate a trigger efficiency of 99.9935% of the SNO+ DAQ system. It is also been demonstrated how time windows can be identified which contain TELLIE light and therefore can be removed from physics analysis. The potential of tagging events as from TELLIE when one has been missed has also been shown. TELLIE has also allowed gated charge spectra to be acquired which contains single photoelectron spectra. This data has been used to demonstrate the bias the discriminator thresholds of the PMTs introduce when running the gain fit calibration code. The stability of TELLIE has also been investigated and an intensity rise over the course of a TELLIE run has been observed. However, when one considers the spread in intensity of light emitted from TELLIE then this effect is negligible. One can also reduce the effects of the intensity rise by simply reducing the maximum number of pulses in one sequence to allow time for the electronics within TELLIE to cool down.

4.6 Data Quality processor

A data quality processor was designed to check whether the data collected using the TEL-LIE system is of a good enough quality to be used for PMT calibration or other analyses. The need for a data quality processor was apparent after examining the data collected from the commissioning runs. During the February commissioning period, a faulty fuse connection caused voltage fluctuations across the system leading to communication issues with the control chip and different numbers of photons per pulse injected into the detector to what was observed in the characterisation of the system at Sussex. The data quality processor performs a series of checks which a run will pass or fail. The results of the checks will allow a user to determine whether the data obtained in the run is suitable for analysis. The data quality processor will be used on any run containing TELLIE events immediately after the run is complete. The results of the data quality processor and a selection of plots will be uploaded to a monitoring website. This will allow users operating the TELLIE system to be alerted quickly if they are taking data which is unsuitable for analysis. A user can then adjust the parameters of TELLIE or debug the system such that when data taking is resumed with the TELLIE system in operation it will then be of good enough quality for data analysis.

4.6.1 Pulse intensity checks

Three intensity checks are performed on data collected while a TELLIE channel or channels are active. These checks are used either for diagnostic purposes or to determine whether the data is suitable for use for calibration. Two values are used to perform three checks on the data. The values used are the maximum NHit of all events in the data and the average NHit of all events in the data collection period.

The maximum NHit obtained is checked against an upper limit which enables users to identify unstable channels that have emitted high intensity pulses of light which have a far greater initially than what is required by the calibration. Any unstable channels identified can then have their drivers removed which then can be fixed and replaced. The fibre can then be connected to a channel that is stable so that the injection point remains active. As there are 91 injection points and 96 channels this means up to five unstable channels can be swapped. If multiple channels are failing this check this can indicate an issue with either the power supply or the control chip.

The average NHit value is the mean of all events containing a TELLIE trigger bit. It is checked against a minimum and maximum threshold value which can be defined by a user dependent on what the data collected is used for. If TELLIE has been used to collect data for PMT calibration then the minimum value is set above the value of NHit that is expected where PMTs are triggered on noise alone and not physics events so that it is known that light is being detected in the detector. The maximum value is chosen to minimise the number of events which contain PMTs which have been triggered by multiple photons as single photoelectron spectra are required for PMT calibration.

4.6.2 Trigger check

The trigger check is used to ensure that all events containing TELLIE light have been tagged with a TELLIE trigger bit. This ensures that all possible data is being used to calibrate the detector and thus minimise the required run time for calibration. It also identifies the risk of events containing TELLIE light being mistagged as physics events if the calibration data is being taken simultaneously with physics data. The values calculated by the data quality processor are the trigger efficiency, ϵ , and the number of suspected missing triggers, N_{miss} . The trigger efficiency is a ratio of the number of events containing a TELLIE trigger bit, N_{trig} , and the number of pulses requested by a user N_{pulses}

$$\epsilon = N_{trig}/N_{pulses}.\tag{4.14}$$

The number of suspected missed triggers, N_{miss} , are calculated by

$$N_{miss} = (t_i - t_{i-1})/\bar{t} - 1, \tag{4.15}$$

where t_i is the time of the current event, t_{i-1} is the time of the previous event and \bar{t} is the mean time between events. N_{miss} is rounded to the nearest integer. A cut is placed on N_{miss} such that it must be less than 5 to be counted as suspected missing. This is to stop sequence boundaries being misidentified as missed triggers.



Figure 4.37: The difference between the mean pulse delay of a run and the pulse delay of an event. All TELLIE data from Decemember 2014 was used.

4.6.3 Pulse delay check

The pulse delay is calculated through the difference in time between consecutive triggered TELLIE events, t_{i-1} and t_i

$$\Delta t = t_i - t_{i-1} \tag{4.16}$$

This value has to fall between ± 10 ms of the expected pulse delay, 99.9% of the time, to pass the check. The value of Δt is approximately 0.21 ms slower than the set pulse delay due to processing time. A value of ± 10 ms was determined by plotting the difference in pulse delay of the event and the mean pulse delay of the run, see Figure 4.37.

4.6.4 Fibre check

The fibre check is to ensure the correct fibre is emitting light pulses. This works by first finding the PMT with the maximum number of hits, NHit. It could be the case that the PMT is triggering on noise and not light pulses. To check this two neighbours on the same panel of the PMT are selected. If both of these PMTs have NHit < 80% then it is assumed that the PMT is noisy and the PMT with the next greatest NHit is selected. The noise check is then repeated until a PMT which is not deemed noisy is found. The



Figure 4.38: An example of a run that passes the fibre check. The plot shows a projection of the PSUP with the number of hits each PMT has been triggered on the z-axis. The circle denotes where the fibre check has calculated the centre of the direct light. The cross denotes where the fibre check has calculated the centre of the reflected light and hence the fibre position.

position of this PMT, \bar{a} , is then used to split the other PMTs into direct and reflected categories:

$$\bar{a} \cdot \bar{b} > 0$$
 Direct
 $\bar{a} \cdot \bar{b} < 0$ Reflected (4.17)

where \bar{b} is the position of the other PMT. An average reflected position \bar{r} is then found by

$$\bar{r} = \frac{1}{N_{tot}} \sum_{i} N_i \bar{b}_i \tag{4.18}$$

where N_i is the NHit at the PMT located at \bar{b}_i and $N_{tot} = \sum_i N_i$. The closest fibre to this reflected position is then assumed to firing.

Figure 4.38 shows the projection of the PMT hits on the PSUP which is what the fibre check uses in its calculations. The fibre check has calculated the centre of the direct light to be located at the position of the circle. The centre of the reflected light has been calculated to be located at the position marked with a cross. This is where the fibre check

believes the fibre is located.

4.6.5 Timing checks

The first timing check is the run time, t_{run} , check. This calculates the run time by subtracting the last event time, t_N , from the first event time, t_0 . To pass the check the following inequality must be true

$$t_{run} = t_N - t_0 \ge N_{pulses} \Delta t, \tag{4.19}$$

where N_{pulses} is the number of pulses in the run and Δt is the set pulse delay.

Further timing checks are made on the TAC count of a run. There should be three peaks in the histogram which correspond to AV reflections and the prompt light directly opposite the fibre. These checks assume that the majority of direct light reaches the AV at 90° incidence. This is a reasonable assumption to make as the light emission is peaked at 0° from the fibre, see Figure 4.15, and the fibres are installed such that they are facing the centre of the detector.

To make it easier to identify the three peaks, a 15° area cut is made around the assumed fibre position and directly opposite the fibre position. The 15° area is with respect to the centre of the detector, the PMT and fibre positions, see Figure 4.39. The two AV reflection peaks will occur in the area cut around the fibre and a prompt peak will occur in the area directly opposite the fibre. ROOT's TSpectrum class is then used to identify the peaks in the TAC histograms. If there are three peaks then the number of peaks check will pass. The prompt peak must occur between 1024 and 3072 ADC counts to pass the prompt peak time check. This is to ensure that there is enough time before and after for AV reflections and approximately a 50 ns buffer. To ensure that the time profile is what is expected for a TELLIE run there is a relative peak time check which ensures that the AV reflection peaks occur between 307 and 512 before and after the prompt peak. These numbers correspond to an event lasting 400 ns with 4096 TAC counts. It is approximately 18 m from one side

Table 4.8: The expected TAC counts and the lower and upper bounds of TAC counts for a given medium of the position of the reflection peaks with respect to the direct light peak. The lower and upper bounds correspond to a 10 ns buffer.

Medium	Lower Bound	Expected	Upper Bound
Air	307	410	512
Water	442	545	647
Scintillator	512	614	717



Figure 4.39: Potential light paths in the detector and area cuts made for the timing checks. The outer circle is the PMT Support Structure (PSUP) and the inner circle is the Acrylic Vessel (AV). The centre of the AV is denoted with the point labelled 0. The fibre position is denoted with the point labelled X. The arrows show potential light paths in the detector, one path is travelling directly opposite the fibre, one is reflected off the near AV surface and one off the far AV surface. The dashed lines show the points on the PSUP which are cut between so that the timing checks can be performed in the direct and reflected light regions. A 15° area cut is made with respect to the fibre position or directly opposite the fibre position and the centre of the detector. A 15° segment is denoted in red.

of the detector to the other, approximately 6 m for the initial reflection off the AV and 30 m for the reflection off the AV on the opposite side of the detector. This is a path length difference of approximately 12 m for both peaks. In air, 12 m corresponds to a light transit time of approximately 40 ns or 410 TAC counts, if we allow a 10 ns (3 m) buffer, this brings the lower limit down to 307 TAC counts and an upper limit of 512 TAC counts. The corresponding TAC counts with 10 ns buffers for the other media used in SNO+ are given in Table 4.8.

There is also a peak amplitude check which ensures that the prompt peak is at least twice as large in amplitude as the AV reflection peaks. This number was determined through observations of the time profiles within the area cut.

Figure 4.40 shows an example of a run which passes all timing checks. It can be seen that the initial reflection is much larger than the second reflection due to scattering and refractions which occur during the lights transit across the detector. It can also be observed that the direct light peak is larger than the reflection peaks due to the low probability



Figure 4.40: An example of a run that passes the time checks described in Section 4.6.5. The blue histogram is the reflected light contained in a 15° area cut around the fibre with respect to the centre of the detector. The red histogram is the direct light contained in a 15° area cut directly opposite the fibre with respect to the centre of the detector. The triangles denote the peaks found by ROOT's TSpectrum class which are used to perform the timing checks.

that reflection will occur at the surface of the AV.

4.6.6 Results of checks

The data quality chacks were performed on 55 runs from the February 2014 data set. The results of these checks can be seen in Table 4.9.

Both of the NHit checks were 100% successful. This implies that TELLIE was not pulsing at anomalous intensities and the correct intensity was used for use for PMT calibration.

The pulse delay check was also 100% successful which means the requested pulse delay was being correctly set by the TELLIE system.

None of the runs passed the trigger check. This is because the check requires 100% trigger efficiency. For reasons described in Section 4.5.3 100% trigger efficiency was not achievable for these runs.

A total of 20% passed the fibre check. There were issues with the wrong channel

firing during these runs which is why the check fails. The problem was traced to a faulty connection at the fuse which was leading to power fluctuations. These power fluctuations caused the control chip to fire the wrong channel. It was due to these fibre checks that a possible hardware issue was identified and shows why data quality checks are necessary for monitoring the stability of the TELLIE system.

A total of 70.9% of runs passed the number of timing peaks check. Any that failed this check will automatically fail the other timing checks which is why these values are lower. The checks failed if the prompt peak was set to arrive too late in the event so the second reflection was missing. This was identified as an issue and future operators of TELLIE have been notified that the prompt peak should be set so it arrives between the appropriate TAC count boundaries to ensure all TELLIE light is recorded during the event time. It could also be due to the wrong fibre firing so the area selected by the area cut was not encapsulating the correct reflections to perform the timing check in.

40% of runs failed the run length check. It was found to be due to run operators physically stopping the run before the pulse sequence had been completed. In the future the SNO+ DAQ system will handle the run boundaries such that the run will not be stopped until the pulse sequence is finished so that this problem is avoided in the future.

4.7 Conclusion

The Timing component of the Embedded LED Light Injection Entity (TELLIE) was designed for SNO+ primarily to calibrate the gain and timing of the PMT array. TELLIE consists of a rack of electronics that is housed on the deck above the SNO+ detector. These electronics produce nanosecond pulses of light which are transmitted down optical fibre from the deck to one of the 91 nodes located on the PMT support structure (PSUP).

Check	Runs Pass (%)	
Average NHit	100	
Fibre	20	
Max NHit	100	
Peak Amplitude	67.3	
Peak Number	70	
Peak Time	60	
Prompt Time	67.3	
Pulse Delay	100	
Run Length	60	
Trigger	0	

Table 4.9: The results of the data quality checks on February 2014 data.

The light illuminates PMTs opposite the injection point and the data collected is used to calibrate the PMT. This system allows for continuous calibration of the detector's PMTs and reduces the need to deploy sources into the active mass which risks introducing various radioactive backgrounds.

Several quality assurance tests were conducted prior to the installation of the system at the detector. The requirements of these tests were:

- A peak wavelength between 500 and 520 nm which maximises the quantum efficiency of the PMTs and minimises the absorption of light by the scintillator cocktail. The measured peak wavelengths ranged between 500 and 510.1 nm with a mean of (505.5±2.6) nm.
- A pulse intensity ranging from 10³ to 10⁵ photons per pulse so single photoelectron and multi-photoelectron hits at the PMTs can be analysed. The highest minimum output from a driver-LED combination was 510 photons per pulse. The lowest maximum output from a driver-LED combination was 4.5 × 10⁵ photons per pulse. Therefore all drivers met the required intensity range.
- The time profile of the light pulses after transmission down a fibre had to be less than 6 ns. All driver-LED combinations met this requirement.

After TELLIE was installed at SNO+ there were commissioning periods in February/March 2014 and December 2014. The goals of these commissioning periods were:

- The stability of TELLIE was monitored. It was found that the intensity rise over the course of a run was less significant than pulse to pulse fluctuations. This means that the probability of multi-photoelectron hits at PMTs will not change significantly.
- The PMT array was calibrated. The gain calibration checks had a 81.7% pass rate. The timing calibration checks had a 87.6% pass rate. There were 3224 PMT channels which obtained enough statistics to perform the calibration checks on.
- The SNO+ DAQ system was tested. A total of 0.65% TELLIE triggers were lost by the DAQ system. As a result of this finding there has been a substantial effort by the DAQ group to resolve this issue.
- Measure a true single photoelectron spectrum. This was achieved successfully and data was used to calculate that the gain threshold could be reduced by an average of 13.6 ± 1.3 counts above pedestal.

Data quality checks were implemented to ensure data taken with the TELLIE system is of a high enough standard for PMT calibration. The checks implemented were an average NHit check to ensure single photoelectron statistics are being collected, a max NHit check to ensure TELLIE is stable, a check to confirm the correct fibre is firing, the delay between pulses is checked to make certain that this parameter was correctly set, the run length is checked to make sure the correct number of pulses were fired, the trigger count is checked to ensure no TELLIE events are lost and there are multiple checks on the time profile of the TELLIE events.
Chapter 5

Methodology For Extracting Limits On Neutrinoless Double-Beta Decay

This chapter details the methodology used in this thesis and implemented in the python based SNO+ analysis software called **echidna** to extract confidence limits. In Section 5.1 the likelihood technique used in echidna to extract confidence limits is detailed. The software also accounts for systematic uncertainties in background rates and measured physical parameters by the SNO+ detector. The techniques used for this are described in Section 5.1. A method for parametrising the reconstruction of energy of events from Monte-Carlo simulation is described in Section 5.2. Methods which parametrise the energy resolution, scale and shift systematics are described in sections 5.2.1, 5.2.2 and 5.2.3 respectively. How echidna handles background rate systematics is described in Section 5.2.5. Background events originating outside the scintillator cocktail are handled differently to backgrounds originating from inside the scintillator, this is explained and justified in Section 5.2.4. Section 5.3 describes the minimisation procedures used by echidna.

5.1 Likelihood technique

At an energy of the Q-value of ¹³⁰Te and after background rejection cuts have been made, a low number background events are expected to be observed over the lifetime of the experiment. Due to this, Poisson statistics will be applicable to the experiment and the appropriate likelihood technique must be selected to reflect this. For a binned Poisson statistics histogram the relevant likelihood function is

$$\mathcal{L}(y;n) = \prod_{i} \mathcal{P}(n_i|y_i) = \prod_{i} e^{-y_i} y_i^{n_i} / n_i!,$$
(5.1)

where y_i and n_i are the number of events predicted by a model and observed in a data set in the i^{th} bin of a histogram respectively. The likelihood ratio test can be used to examine the goodness of fit of different models in question. If m is the true, unknown values of nthen one can construct the likelihood ratio

$$\lambda = \mathcal{L}(y; n) / \mathcal{L}(m; n). \tag{5.2}$$

This ratio can then be used to construct a likelihood chi-squared which is defined by

$$\chi^2 = -2ln\lambda = -2ln(\mathcal{L}(y;n)) + 2ln(\mathcal{L}(m;n)).$$
(5.3)

By replacing the unknown m with its bin-by-bin model-independent maximum likelihood estimation which is equivalent to n, the Poisson likelihood chi-square is found to be [106]

$$\chi^2 = 2\sum_i y_i - n_i + n_i ln(n_i/y_i).$$
(5.4)

This is the χ^2 used by echidna to extract the confidence limits for neutrinoless double beta decay sensitivities. If systematics are neglected at this point then equation 5.4 only requires two inputs per bin to calculate a χ^2 , an observed number of events n_i and an expected number of events y_i . In a real experiment n_i and y_i can be either expected background, b, or expected background plus an expected signal, b + s. When calculating sensitivities there are two possible ways of calculating a χ^2 that are of interest

- 1. $n_i = b + s, y_i = b$.
- 2. $n_i = b, y_i = b + s$.

Option 1 will calculate a χ^2 where the amount of signal in your data is varied and it is compared to a model that only contains background. In a real experiment n_i will be a fixed number therefore this option does not represent what is observed after a complete data set has been obtained. Even though this option may not be physically true some meaning can be interpreted from this method. This option allows us to obtain the confidence level that the result $m_{\beta\beta} = 0$ can be excluded if an excess is observed above an assumed background level. An issue with using this option is that the signal's systematic error is included, and

Confidence Level	$\Delta \chi^2$		
(%)	m = 1	m = 2	m = 3
68.27	1.00	2.30	3.53
90	2.71	4.61	6.25
95	3.84	5.99	7.82
95.45	4.00	6.18	8.03
99	6.63	9.21	11.34
99.73	9.00	11.83	14.16

Table 5.1: The values of $\Delta \chi^2$ which correspond to a median confidence limit for simultaneous estimation of *m* parameters [108].

not varied, in the fit so the previous conclusion is not completely accurate. Option 2 will calculate a χ^2 which would be obtained in a real experiment where only the background events have been observed and there is no contribution from the signal. A limit at a given confidence level can then be obtained for the half-life of neutrinoless double-beta decay in ¹³⁰Te by varying the signal contribution above the background level with all systematic uncertainties accounted for. This is the method that is used in this thesis.

If a high statistics data set, known as an "Asimov data set" is used to calculate the χ^2 for a given hypothesis, the median confidence level for an experiment [107] can be obtained from the value of

$$\Delta \chi^2 = \chi^2 - \chi^2_{min},\tag{5.5}$$

where χ^2_{min} is the value from the hypothesis which has the lowest value of χ^2 . The values for $\Delta \chi^2$ that correspond to given confidence levels are given in Table 5.1 [108].

The likelihood technique defined in equations 5.1 to 5.4 neglects the effect of systematic uncertainties and hence assumes perfect knowledge of signal and background events in the parameter space one is extracting a confidence limit from. Clearly once systematic uncertainties have been included in the analysis then this will lead to a lower predicted sensitivity. If it is assumed that the systematic uncertainties are Gaussian then one can add a nuisance parameter to account for this by defining the likelihood as

$$\mathcal{L} = \prod_{i} \mathcal{P}(n_i | y_i) \mathcal{G}(X | \tilde{X}, \sigma_X),$$
(5.6)

where \mathcal{G} is a normalised Gaussian with a mean of \tilde{X} and sigma, σ_X , which has been added multiplicatively to equation 5.1 and represents the systematic uncertainty in parameter X. Here X has been measured by an analysis, independent from the limit setting analysis, to have a mean value of $\tilde{X} \pm \sigma_X$. The log-likelihood ratio can then be calculated arriving at [107]

$$\chi^{2}_{total} = \chi^{2} + \chi^{2}_{Pen} = -2ln\lambda - 2ln\lambda_{Pen} = \chi(X)^{2} + \frac{(X-X)^{2}}{\sigma^{2}_{syst}},$$
(5.7)

where χ^2 is as stated in equation 5.4. χ^2_{pen} is sometimes referred to as a **penalty term** [109] as this can only be zero or positive and hence can only increase, or 'penalise', the total χ^2 . The addition of this term to the test statistic allows the current value of the parameter X to be floated away from its mean expectation value \tilde{X} during the likelihood fit. One can therefore minimise the chi-squares obtained from varying X away from \tilde{X} for each scaling of the signal and a new value for the sensitivity to a signal can be obtained with this systematic taken into account.

5.1.1 Verification of the penalty term



Figure 5.1: A histogram showing the Gaussian distributions used to verify that the penalty term has been implemented correctly. The same non-normalised Gaussian distribution was used for the signal and the background and has a mean value of 2.5 MeV and a sigma of 0.112 MeV. The background in black was normalised to 100 counts. The histogram in blue shows the scale of the signal at the 90% Confidence Limit (CL) when no penalty term was applied. The histogram in red shows the scale of the signal at the 90% CL when a penalty term was applied where the background rate has a systematic uncertainty of $\sigma = 20$ counts and was varied in steps of one count between 50 and 150 counts.

To verify that the penalty term had been implemented correctly in echidna a mock data set was generated consisting of two identical Gaussian PDFs with a mean value of 2.5 MeV and a RMS of 0.112 MeV. Each data set contains 100,000 randomly distributed events from these PDFs. One was used as the background and scaled to 100 counts and



Figure 5.2: The minimised χ^2 value obtained for each scaling of the signal. No penalty term and a fixed background was used for the points in blue. A penalty term was used for the points in red where the background rate had a system uncertainty of $\sigma = 20$ counts and was varied between 50 and 150 counts in steps of 1 count. The PDFs used are shown and described in Figure 5.1. The dashed black line corresponds to one σ . The dashed blue and red lines are where the cases of no penalty and penalty intersect with the one σ line at 10 and 22.36 respectively.

the other data set was used as the signal and was scaled in steps of 0.25 counts from 0.25 to 25 counts. The one sigma sensitivity to the signal is then extracted when the χ^2 value for a given signal count is equal to or greater than one. The background was then given a rate uncertainty of 20 counts and the case with and without the penalty term, and hence the systematic uncertainty, were examined. When no penalty term is included the background is fixed to 100 counts. With the penalty term the background count was allowed to vary between 50 and 150 counts in steps of one count. The background PDF and the corresponding one sigma CL signal PDFs without and with the penalty term are shown in black, blue and red in Figure 5.1.

If the penalty term has been implemented correctly in echidna then the extracted one sigma sensitivity to the signal should correspond to the statistical and systematic error added in quadrature. The statistical error is equivalent to $\sigma_{stat} = \sqrt{N} = 10$ counts and the systematic error is set to $\sigma_{syst} = 20$ counts. Therefore the value of sigma with the penalty term added should be

$$\sigma = \sqrt{\sigma_{stat}^2 + \sigma_{syst}^2} = \sqrt{100 + 400} \approx 22.36.$$
 (5.8)

Figure 5.2 shows the obtained χ^2 value for each signal scaling with and without the penalty term in blue and red respectively. The black line corresponds to one sigma. The extracted one sigma values for the cases with or without the systematic uncertainty are 22.5 and 10 counts respectively. A finer scan between 22.25 and 22.5 signal counts reveal that once the penalty term has been added the extracted one sigma CL is 22.36 counts as expected thus confirming that the penalty term has been implemented correctly.

5.2 Parametrised Monte Carlo

A parametrised Monte Carlo was used to mimic the effects of different energy resolutions, scales and shifts. This was used instead of full Monte Carlo simulation due to the inordinate amount of computing time required to produce individual data sets which can represent variations in the above parameters. The parametrised Monte Carlo allows variables to be changed 'on the fly' during minimisation routines which is required for sensitivity studies.

5.2.1 Energy Resolution

In order to mimic different light yields and energy resolutions one must use a technique which will convert the true quenched energy deposited in the scintillator to an appropriate reconstructed energy. Ordinarily one would simulate the expected backgrounds in the detector and reconstruct the events observed in the detector to extract a value for reconstructed energy. The issue with this is that, in the region of interest, there are tails of the energy spectra of backgrounds which require a large number of statistics to reproduce the expected energy response of the detector. This high statistics data set would also be needed to be reproduced for different light yields leading to an inappropriately large requirement in computing time and storage space.

Two alternative, less computing intensive techniques were explored in this thesis by assuming that the energy response of the detector is either Gaussian or Poissonian. One can see in Figure 5.3 that for energies of interest for neutrinoless double-beta decay, 130 Te has a Q-value of 2527.01 ± 0.32 keV, and the relationship between the amount of quenched energy deposited in the scintillator and the number of PMT hits observed is expected to be linear. Therefore one can use the expected NHit per MeV as a measure of the energy resolution as it is only the gradient of the relationship between NHit and quenched energy



Figure 5.3: The number of hit PMTs (NHit) as a function of the quenched energy deposited in the scintillator from ²¹⁰Tl beta decay in the rat 4.5.0 production data set. The data was fitted with a linear function which has a resultant gradient and intercept of 204.44 ± 0.71 NHit/MeV and 8.7 ± 1.5 NHit respectively. The reduced χ^2 of the fit is 1.57.

deposited that will change for the energies of interest.

If it is then assumed that the energy response of the detector is Gaussian then one can calculate an energy resolution, σ , of the Gaussian using

$$\sigma = \sqrt{E/L},\tag{5.9}$$

where E is the quenched energy deposited in the scintillator and L is the light yield measured in NHit/MeV. The expected light yield during the 0.3% Te-Loaded phase is 200 NHit/MeV. With this energy resolution one can then construct a Gaussian PDF representing the detector response to an event depositing a true quenched energy \bar{E}

$$P(E) = \frac{1}{\sigma\sqrt{2\pi}} exp\left(-\frac{(E-\bar{E})^2}{2\sigma^2}\right).$$
(5.10)

One way to use this technique to recreate the reconstructed technique is to produce a Gaussian or Poissonian PDF from the \bar{E} of each event. However, this is clearly too computer intensive for a data set containing a large number of events. A solution to this is to first bin \bar{E} in a histogram and produce one Gaussian PDF per bin where \bar{E} now becomes the mean value of the bin and the PDF is renormalised to the total number of events in that bin. A problem with this technique is that it may introduce a bias, particularly at the tails of the PDFs of the background spectra due to the steep slope of the energy spectra. For example the true values of \bar{E} within a bin in the spectrum of ¹³⁰Te $2\nu 2\beta$ decay as you approach its Q-value are more likely to be at a energy lower than the mean value of that bin. Hence, the Gaussian PDF produced will be biased towards higher energies than if individual Gaussian PDFs were produced for each event. To minimise this effect the background and signal spectra are first binned in small steps of 500 eV before applying the technique. This reduces the bias on the mean for events with a true energy at an edge of a bin to 250 eV. The maximum bias on the σ of the PDF for these events which have a true quenched energy of \bar{E} is

$$\sigma_{bias} = \sqrt{\frac{\bar{E} + 250 \text{ eV}}{L}} - \sqrt{\frac{\bar{E}}{L}}.$$
(5.11)

For events with an energy of 2527 keV near the Q-value of Te¹³⁰ and L = 200 NHit/MeV the σ_{bias} is 5.56 eV. This is negligible in comparison to the level of uncertainties investigated in this thesis. For example, a 0.1% uncertainty in energy scale at 2.5 MeV corresponds to an energy of 2.5 keV.

Weights are calculated from the Gaussian PDF for bins within 5σ of the current bin. After the technique has been applied the resultant spectra are then rebinned in to larger step sizes which are then used for the limit setting routines. As the largest background in the region of interest for neutrinoless double-beta decay is from $2\nu 2\beta$ decay then any bias introduced by using this technique to recreate detector energy response will lead to a more conservative limit on the half life of neutrinoless double-beta decay.

The alternative distribution used to mimic energy reconstruction is a Poisson PDF. This technique uses the Poisson statistics of the number of hit PMTs N. The probability of observing N for an event which has a mean number of hit PMTs \bar{N} is therefore

$$P(N) = \frac{\bar{N}^N e^{-\bar{N}}}{N!}.$$
 (5.12)

As \bar{N} is related to the \bar{E} of the event by the light yield $\bar{N} = L \cdot \bar{E}$, if the \bar{E} distribution is binned in a histogram in the same way as before the Gaussian approximation is applied, the resultant reconstructed \bar{E} PDF can be obtained by using the weights obtained from the PDF

$$P(E) = \frac{(\bar{E} \cdot L)^{E \cdot L} e^{-\bar{E} \cdot L}}{(E \cdot L)!}.$$
(5.13)

Due to the small binning required when producing background PDFs at different res-

olutions the process is very computational heavy. As a result of this the above methods are not appropriate when performing minimisation routines. An interpolation method is used in this case. All background PDFs are produced at different resolutions in steps of one NHit/MeV around the prior resolution value. For a given resolution value the minimisation routine loads two PDFs below the value and two above the required value. For each energy bin a cubic spline interpolation is performed using scipy's python libraries. The energy bin in the resultant spectrum is then weighted according to the value calculated by the interpolation at the requested resolution.

5.2.2 Energy Scale

Energy scale represents how the PMT hits in the detector is converted into a measured energy by the reconstruction. It can be constructed as the ratio between the gradient of the line in Figure 5.3 and the calibrated value of light yield. In this case, if the light yield of the detector has been calibrated to be 200 NHit/MeV, the energy scale is therefore approximately 1.022 ± 0.004 . The energy scale therefore represents a miscalibration of the light yield. Many factors can effect the energy scale. These include variations in the number of online PMTs during physics running, degradation of the liquid scintillator and changes to trigger thresholds in the DAQ. The energy scale is expressed numerically by a coefficient *a* and, when the effect is taken individually, can be written as

$$E = a \cdot E',\tag{5.14}$$

where E' is the quenched energy deposited in the scintillator and E is the energy measured by the detector.

The energy scale is applied by the echidna software by first interpolating a binned PDF of the quenched energy deposited. If the mean energy of a bin in the original PDF is \bar{E}' and a scale factor a is applied to the spectrum, then the weighting of the bin, w, in the scaled spectrum of mean energy \bar{E} is calculated by evaluating the interpolation $f(\bar{E})$ at position

$$w = f(\bar{E}/a). \tag{5.15}$$

Proportionate weighting of other dimensions such as radius is retained across all bins which have energy \bar{E}/a .

For similar reasons as in the case of energy resolution, fine binning is required so events are not artificially scaled to higher or lower values due to the fact that a mean binned energy is used. The range of energy used should be greater than $a \cdot E_{low}$ and $a \cdot E_{high}$ where E_{low} , E_{high} are the upper and lower edges of the region of interest used in sensitivity studies. This is due to edge effects arising from the interpolation at the outermost bins.

5.2.3 Energy Shift

The energy shift can be viewed as the intercept divided by the true value of energy resolution. For example in Figure 5.3 the intercept is 8.7 ± 1.5 NHit therefore if it is assumed that the calibrated light yield is approximately 200 NHit/MeV then the energy shift is approximately 43.5 ± 7.5 keV. The energy shift can be viewed as representing the average number of noise hits in an event. The level of noise hits effectively shift the energy of all events by the same absolute amount of energy. Similar to energy scale, energy shift can be effected by variations in the number of online PMTs during physics running and changes to trigger thresholds in the DAQ. Energy shift can be expressed numerically by an addition of value *b* and, when this is the only effect on energy, can be written as

$$E = b + E', \tag{5.16}$$

where E' is the quenched energy deposited in the scintillator and E is the energy as measured by the detector.

The energy shift is applied by the echidna software by first interpolating a binned PDF of the quenched energy deposited deposited. If the mean energy of a bin in the original PDF is \bar{E}' and the spectrum is shifted by a value b, then the weighting of the bin, w, in the scaled spectrum of mean energy \bar{E} is calculated by evaluating the interpolation $f(\bar{E})$ at position

$$w = f(\bar{E} - b).$$
 (5.17)

Proportionate weighting of other dimensions such as radius is retained across all bins which have energy $\bar{E} - b$.

Similarly to energy resolution and scale, fine binning again is required so events are not artificially shifted to higher or lower values due to the fact that a mean binned energy is used. The range of energy used should be greater than $E_{low} - b$ and $E_{high} + b$ where E_{low} , E_{high} are the upper and lower edges of the region of interest used in sensitivity studies. This is due to edge effects arising from the interpolation at the outermost bins.

5.2.4 Energy Reconstruction Of Backgrounds Originating Outside The Scintillator

Backgrounds which originate from outside the scintillator, known as external backgrounds, must be handled differently to internal backgrounds in this thesis due to the nature of their interactions in the scintillator. External backgrounds are any background whose isotope originates from the inner surface of the AV outwards. The backgrounds of interest usually emit gamma rays as they decay and these gammas deposit energy via compton scattering in the scintillator. There can be multiple scatters within the scintillator which results in difficulty in quantifying a single radial position of the interaction both in terms of true and reconstructed position. There is currently no output that quantifies the true position of these events in the RAT software. Due to this, reconstructed position is always used for external backgrounds. A significant portion of external events have interactions near the inner surface of the AV but are misreconstruced inside the fiducial volume with a low energy. Figure 5.4 shows reconstructed energy as a function of true energy, where true energy is the true quenched energy deposited in the scintillator. This figure shows how the majority of events are misreconstructed to a lower energy. For a region of interest of 2.2 to 2.8 MeV and a fiducial volume cut at a reconstructed radius of 3.5 m, the highest true energy event that was misreconstructed into the region of interest was 2.909 MeV to a reconstructed energy of 2.774 MeV; the lowest true energy event was 1.959 MeV with a reconstructed energy of 2.316 MeV. The largest difference between reconstructed and true energies was 0.3808 MeV between a true energy of 2.524 and a reconstructed energy of 2.143. The value of σ at L = 200 NHit/MeV at this true energy is 0.1123 MeV making the largest misreconstruction into or within the region of interest 3.39 σ away in energy. Based on these observations it is assumed that the poor reconstruction occurs when true energies are misreconstructed to lower energies rather than the much more dangerous opposite case. Figure 5.4 also shows that below approximately 1 MeV reconstructed energies of external background events become unreliable as they could have a much higher true quenched energy in the scintillator.

Table 5.2 gives the misreconstruction rates for 214 Bi and 208 Tl events originating from different source locations in the detector. The misreconstruction rates were calculated using the equation

$$M = (1 - R/T) \cdot 100 \tag{5.18}$$

where M is the misreconstruction rate and R and T are the reconstructed and true energies of events in the region of interest between 2.2 and 2.8 MeV and after a fiducial volume



Figure 5.4: A two dimensional histogram showing reconstructed energy as a function of the true quenched energy deposited in the scintillator. All external background files in the rat 4.5 production set were used and were not weighted according background rate. A fiducial volume cut at a reconstructed radius of 3.5 m was made. The red lines indicate the region of interest used in this thesis between 2.2 and 2.8 MeV.

cut at a reconstructed radius of 3.5 m. Examining these rates in Table 5.2 it is clear that the closer to the inside of the AV the original isotope is, the more likely it is to be misreconstructed to the wrong energy. This is likely due to the effect of interactions occurring near the AV which are being misreconstructed inside the fiducial volume to a

Table 5.2: A table of the different external backgrounds arranged by their original isotope location within the detector. The number of events which have a reconstructed (Reco) or true quenched energy deposited (True) between 2.2 and 2.8 MeV is listed. The misreconstruction rate (Misreco) as calculated by 5.18 is also listed for each background. All external backgrounds in the rat 4.5 production were used and they were not weighted by the expected rate of each background component. A fiducial volume cut at a reconstructed radius of 3.5 m was made.

	Background Type	$^{214}\mathrm{Bi}$			²⁰⁸ Tl		
Dackground Type	Reco	True	Misreco (%)	Reco	True	Misreco (%)	
	AV	1	33	$97.0 {\pm} 3.1$	29	1621	$98.2 {\pm} 0.34$
	AV Innerdust	0	428	100	2	175	$98.9 {\pm} 0.81$
	AV Outerdust	3	36	$91.6 {\pm} 5.0$	121	4480	$97.3 {\pm} 0.25$
	External Water	53	71	25 ± 14	105	262	$59.9 {\pm} 4.6$
	Hold Down Ropes	15	24	$38{\pm}21$	319	504	36.7 ± 4.5



Figure 5.5: A histogram showing the number of events with an energy between 2.2 and 2.8 MeV as a function of volume elements defined by the cube of reconstructed radius, R, normalised to the radius of the AV, R_{AV} . Reconstructed events which have an energy in the energy range are in red and true quenched energy deposited in the scintillator events in the the energy range are in blue. Reconstructed position is used for both the true and reconstructed events. The entire external background non-normalised RAT 4.5 production data set was used.

low energy. It should also be noted that due to the low number of events that occur in the region of interest, some misreconstruction rates have large errors associated with them, especially for ²¹⁴Bi backgrounds.

Figure 5.5 shows the number of reconstructed energy and true quenched energy deposited in the scintillator events as a function of reconstructed volume element which have an energy between 2.2 and 2.8 MeV. There is an increase in true events towards the centre of the detector which occur in a volume elements $(R/R_{AV})^3 \leq 0.03$, where Ris reconstructed radius and R_{AV} is the radius of the AV. This increase is likely due to events which actually occur near the AV being misreconstructed as events occurring in the centre of the detector with a low energy which is why no reconstructed events occur in this energy range. Reconstructed events only start appearing in this energy window in volume elements $(R/R_{AV})^3 \geq 0.03$.

Figure 5.6 shows the misreconstruction rate as calculated by equation 5.18 as a function



Figure 5.6: A histogram showing the misrecontruction rate, as defined by equation 5.18, of all combined external background events with an energy between 2.2 and 2.8 MeV as a function of volume elements defined by the cube of reconstructed radius, R, normalised to the radius of the AV, R_{AV} . The entire external background non-normalised RAT 4.5 production data set was used.

of volume element. This shows energy misreconstruction has a clear radial dependence. If the parametrised Monte Carlo is to be used for external background events then this radial dependence must be taken into account when applying fiducial volume cuts or if events were binned by volume element in a two dimensional sensitivity study. This dependence is in addition to the isotope and background type dependencies shown in Table 5.2. The combination of these effects would lead to large uncertainties on the misreconstruction scale factors which would need to be applied to the individual backgrounds.

Figure 5.7 shows the energy distributions for reconstructed, true, true with an additional energy cut between 2.2 and 2.8 MeV in reconstructed energy and smeared energy distribution. The smeared energy distribution was constructed using the true energy distribution and echidna's energy resolution function as described in 5.2.1. The smeared and true energy distributions have been weighted by the misreconstruction rates described above. The ²¹⁴Bi and ²⁰⁸Tl misreconstruction rates were combined for each background source location due to the large uncertainties for the ²¹⁴Bi misreconstruction rates. The peaks on the left and right of Figure 5.7 in the true energies are due to gammas from ²¹⁴Bi



Figure 5.7: The number of expected external background counts per year as a function of energy in 0.3% Te-loaded scintinllator with a light yield of 200 NHit/MeV. A fiducial volume cut was made on reconstructed radius R normalised to av radius R_{AV} at $(R/R_{AV})^3 = 0.2$. All events were scaled to the background rates in [110] and an external background event rejection of 50% was assumed. The red points are reconstructed energies, blue and yellow points are true quenched energy deposited in the scintillator. The yellow points have a further energy cut so their reconstructed energy is between 2.2 and 2.8 MeV. The green points are smeared from the blue points by echidnas energy resolution function as decribed in Section 5.2.1. A further weighting was applied to the blue and green points to account for misreconstruction rates described above. The misreconstruction rates for ²¹⁴Bi and ²⁰⁸Tl background types were combined due to the large uncertainty associated with ²¹⁴Bi. The error bars represent statistical and systematic uncertainty.

and ²⁰⁸Tl respectively. The smeared energy distribution is supposed to mimic the reconstructed energy distribution. There are a number of bins which are in poor agreement between the smeared and reconstructed distributions due to artefacts in the reconstruction. It is unclear why these artefacts appear. However, the majority of reconstructed and smeared energy bins are within error.

Taking all of the above findings into account it was decided that it would be more accurate to use the reconstructed energy and position in constructing the external energy PDFs rather than introducing uncertainties through misreconstruction rates due to the large number of variables that could affect them. As a result, the external PDFs the energy systematics are kept fixed in the following sensitivity studies. Only the external background event rate is allowed to vary.

5.2.5 Background rates

Event rates for each background component were modelled using a simple scale factor applied uniformly to the entire data set. Each PDF is built from a certain number of decays and a scale factor is applied to generate a second PDF which represents the number of decays for the required exposure. For the RAT 4.5 production data set the original number of decays is simply the number of entries in the data set as all interactions are stored regardless of whether they led to a triggered event. For the RAT 5.0.2 production data set, only the triggered events are stored so the original number of decays has to be reverse engineered. The macro used to generate each file contains an event rate, R, in Hz and is run for a year. Therefore the number of decays, D, for N files can be calculated by

$$D = N \cdot R \cdot 365 \cdot 24 \cdot 60 \cdot 60 \text{ seconds.}$$
(5.19)

5.3 Minimisation

Two methods of minimisation were used in this thesis, a grid search and MINUIT [111]. The grid search was used when examining two floating parameters and MINUIT was used for a greater number.

5.3.1 Grid Search

The grid search calculates the value of χ^2 for each combination of floating parameters between a range in a given step size. It then fills an N-dimensional grid for N parameters with these values. Once the grid has been filled the minimum value is found and the best fit for each parameter is the value used to calculate that particular χ^2 with the associated error being half the step size. The grid search is useful for examining the correlations between two floating parameters as the χ^2 is mapped to a two dimensional grid. The grid search is also useful for verification of other more complex minimisers as the true minimum is reliably found by using a grid search. The problem with the grid search is the computing time as when floating parameters with equal number of steps the number of calculations required goes to the power N dimensions. Due to this, the grid search is not used for minimising scenarios which have more than two floating parameters.

5.3.2 MINUIT

For three or more floating parameters the MINUIT minimiser is used. MINUIT was developed by Fred James in the 1970s and is a common piece of software used in the particle physics community [111]. It searches for a minimum in a user defined function by varying several parameters. The algorithms used by MINUIT in this thesis are the MIGRAD and MINOS algorithms. The MIGRAD algorithm is used to find the minimum χ^2 and the best fit of all floating parameters. The MINOS algorithm is then run to evaluate the errors of the parameters.

5.4 Conclusion

A python based analysis software called echidna has been developed to extract sensitivity limits of neutrinoless double-beta decay for SNO+. This is achieved by calculating a Poisson likelihood chi-square from the number of events predicted by a model and the number which is observed in a data set. The number of events observed in the data set is assumed to be at the background only level and this is compared to models which include a signal contribution and the expected background. Systematic uncertainties are included in the analysis by assuming they are Gaussian and adding a penalty term to the total chi-square.

A parametrised Monte Carlo was designed to mimic the effects of different energy resolutions, scales and shifts. To imitate different light yields and energy resolutions the true quenched energy deposited in the scintillator is smeared with a Poisson distribution. The energy scale represents a miscalibration of the light yield of the scintillator cocktail and is represented numerically by multiplying the energy by a coefficient. The energy shift can be viewed as the average noise hits in an event and is expressed numerically by the addition of a value in MeV to the energy. The event rates for each background is modelled using a scale factor which is applied uniformly to the data set.

Backgrounds originating from outside the scintillator cocktail are treated differently to those that originate from within. This is due to a large number of events being misreconstructed into the centre of the detector when they actually occur elsewhere in the the scintillator cocktail. The rate of misreconstruction is background dependent and differs for the various isotope origins within the detector. The misreconstruction rate is also radial dependent. For these reasons the reconstructed energy and position is used when constructed the PDFs for backgrounds originating from outside the scintillator cocktail. Two different techniques of minimisation are used in echidna. A grid search calculates the chi-square for each combination of floating parameters between a range of values in a given step size. The minimum chi-square is then extracted from this grid of values. The second method used is MINUIT which finds the minimum in a user defined function by varying the floating parameters.

Chapter 6

SNO+ Sensitivity to Neutrinoless Double Beta Decay

The main goal of SNO+ is to discover neutrinoless double-beta decay or, if no signal is detected, set the best possible limit on the half-life of the process and the effective Majorana mass. Therefore it is important that the analysis of the data collected during the double beta phase will maximise the sensitivity to this process. A python based data analysis software called echidna, described in the previous chapter, has been written to develop the analysis. Section 6.1 describes the current methods used by the SNO+ experiment to extract the sensitivity to neutrinoless double-beta decay and states the value obtained using the same methodology and the echidna software. Section 6.2 describes how the energy range and binning was optimised to extract the best limit for neutrinoless double-beta decay for the fixed background case. Section 6.3 introduces systematic uncertainties to the fit. In this section the sensitivity is examined as a function of the uncertainty on the individual backgrounds and energy resolution, scale and shift. Correlations between systematic uncertainties are also probed. All systematic uncertainties are then included and the sensitivity to neutrinoless double-beta decay is extracted. Backgrounds are then removed individually. This allows us to view the effect of a systematic uncertainty on the sensitivity to neutrinoless double-beta decay whilst accounting for correlations between other systematic uncertainties. The fit is extended to two dimensions in Section 6.4 so both energy and radius information is included in the fit.



Figure 6.1: The energy spectra within the fiducial volume of the different categories of background expected to be observed with the SNO+ detector with the assumptions stated in Table 6.1. The signal in black is the 90% confidence limit corresponding to a half-life of $T_{1/2} = 9.38 \times 10^{25}$ yr.

6.1 Current State of the Art Sensitivity Calculations

To get the best possible sensitivity to neutrinoless double-beta decay, SNO+ aims to maximise the number of signal counts to background in a given parameter space. Figure 6.1 shows the energy spectra of the backgrounds and potential signal between 1.7 and 3.7 MeV. Below 2.48 MeV the ¹³⁰Te two neutrino double-beta decay dominates. The region above 2.86 MeV is dominated by the thorium-chain backgrounds. There is also a relatively flat background across the energy range considered from solar ⁸B neutrinos elastic scattering in the scintillator and this dominates between 2.48 and 2.86 MeV.

Previously, an independent SNO+ analysis code was used to extract the 90% confidence limit (CL) sensitivity for neutrinoless double-beta decay [114]. This analysis was a simple counting analysis in a single bin of the energy spectrum with an energy range defined as 2.42 to 2.64 MeV corresponding to $-1/2 \sigma$ to $3/2 \sigma$ from the extracted mean value of a Gaussian fit to the signal spectrum as shown in Figure 6.2. The limit calculated in this analysis was $T_{1/2} = 9.38 \times 10^{25}$ yr at 90% CL for the experimental parameters listed in Table 6.1.



Figure 6.2: The ¹³⁰Te $0\nu 2\beta$ decay energy spectrum in blue fitted with a Gaussian denoted by the dashed red line. The results of the fit was a mean of 2.476 ± 0.018 MeV and sigma of 0.111 ± 0.012 MeV. The vertical solid line corresponds to the fitted mean and the small dashed vertical lines correspond to \pm one sigma with respect to the mean. The larger dashed vertical lines from left to right correspond to -1/2 sigma and 3/2 sigma with respect to the mean.

Using the same data set, RAT 4.5 software release [115], and experimental parameters in Table 6.1, the same limit was computed with the echidna software. In both analyses the true quenched energy in the scintillator for each internal background was binned in 500 eV bins before convolving with Gaussian distributions with a light yield

Experimental Parameter	Value
Live time	5 years
Natural Te loading	0.3%
Light yield	200 NHit/MeV
Fiducial volume	$< 3.5 \mathrm{m}$
Detection efficiency	100%
Position reconstruction efficiency	100%
BiPo separate trigger rejection	100%
BiPo pileup rejection	80%
External background event rejection	50%
(α, n) rejection	90.265%
Background rates	See [110]

Table 6.1: A list of experimental parameters used in the previous SNO+ sensitivity analysis [112, 113]. The same experimental parameters were used for the verification of echidna with this analysis. The BiPo rejection factors are applied to both $^{214}Bi^{-214}Po$ and $^{212}Bi^{-212}Po$ decays.



Figure 6.3: A graph showing the extracted χ^2 as a function of the inverted half-life $(1/T_{1/2})$ of $0\nu 2\beta$ decay in ¹³⁰Te. The horizontal dashed blue line is at $\chi^2 = 2.71$ which corresponds to the 90% CL. Each category of background labelled was reduced to zero in the fit. The extracted χ^2 for these data points corresponds to a single binned analysis from 2.42 MeV to 2.64 MeV. The experimental parameters in Table 6.1 were used.

of 200 NHit/MeV to mimic energy resolution as described in Section 5.2.1. For reasons described in Section 5.2.4, the reconstructed energy and position was used for the external backgrounds. The limit obtained using a single binned analysis with the echidna software is 9.41×10^{25} yr at 90% CL. This corresponds to a 1.0% difference with respect to the previous analysis [114].

Background Type Removed	$T_{1/2}$ 90% CL
$^{8}\text{B} \nu \text{ES}$	1.16×10^{26}
130 Te $2\nu 2\beta$	1.06×10^{26}
External	1.04×10^{26}
Th Chain	1.01×10^{26}
U Chain	9.73×10^{25}
Cosmogenics	9.45×10^{25}
(α, n)	9.41×10^{25}
All Backgrounds	9.41×10^{25}

Table 6.2: A table showing the different background types ranked in ascending order by their extracted neutrinoless double-beta decay half-life at 90% confidence level, $T_{1/2}$ 90% CL using the single bin analysis. Each background type was reduced to zero in the fit and was fixed at its expected value. The experimental parameters listed in 6.1 were used.

Each background has a different number of counts in the single energy bin and will have a lesser or greater influence on the sensitivity to neutrinoless double-beta decay. The echidna software was used to determine which backgrounds have the greatest influence on sensitivity in the single bin fit between 2.42 and 2.64 MeV. For each background, the sensitivity was calculated in the case where that background was reduced to zero and the other backgrounds were fixed to their nominal rates. As before, the experimental parameters in Table 6.1 were used. Figure 6.3 shows the χ^2 as a function of half-life for each individual background which was reduced to zero. With this counting experiment approach, ⁸B neutrinos have the greatest impact on sensitivity and the 2223 keV gammas from neutron capture have the least effect. Table 6.2 ranks the impact of backgrounds in terms of the extracted 90% confidence limit for neutrinoless double-beta decay.

6.1.1 Transition to RAT 5.0.2

Since the independent SNO+ analysis was performed using the RAT version 4.5 software release, there was a new software release, version 5.0.2, with an associated production of backgrounds. A full list of changes between the software versions can be found here [116]. These include changes to the geometry, to the scintillator and loaded scintillator optical properties and multiple new production macros which simulate the data sets used in the analyses. These new production macros include simulations of the hold up rope backgrounds which were previously neglected. These changes affected the resultant sensitivity.

Also, approximately 10 times more statistics were generated with RAT 5.0.2. As an example, for 208 Tl hold down ropes, which is the dominant external background, only 10 years worth of data was simulated using RAT 4.5 whereas approximately 100 years was simulated with RAT 5.0.2. Similarly for 130 Te two neutrino double-beta decay, 10 years was simulated with RAT 4.5 and 100 years was simulated with RAT 5.0.2. Similar results are found with other backgrounds which suggests that there was statistical fluctuations in the RAT 4.5 data set may have been an issue as the sensitivity studies were examining 5 year livetimes where only 10 years worth of data had been simulated. Statistical fluctuations should not be an issue with the RAT 5.0.2 data set.

The sensitivity achieved using a single bin analysis and the data simulated by RAT 5.0.2 was 8.81×10^{25} years. This is a reduction of 6.4% in comparison to RAT 4.5. For the rest of this thesis RAT 5.0.2 data set and software release will be used.



Figure 6.4: The half-life at 90% confidence limit as a function of energy range of the region of interest centred at 2.5 MeV. The assumptions used are listed in Table 6.1. Seven different bin sizes were investigated listed in the legend.

6.2 Energy Range and Binning Optimisation

The previous analysis only used a single bin and hence a narrow energy window in its fit. Signal to background varies as a function of energy and radius and therefore an improved limit should be attainable by binning and fitting in these dimensions. To obtain a more realistic sensitivity, systematic uncertainties also need to be accounted for.

The advantages of using binned energy spectra in respect to a simple counting experiment can be demonstrated by binning the region used in the previous analysis with 20 keV bins. With a binned analysis a sensitivity of 9.05×10^{25} years at 90% CL is achieved. In comparison, the simple counting analysis, where one bin is used, reaches a sensitivity of 8.81×10^{25} years.

As information from the varying signal to background across the energy spectrum is included in the fit when the window is divided into bins, the energy window is no longer optimised to achieve the maximum sensitivity to neutrinoless double-beta decay. To optimise the energy window the energy range was extended symmetrically around a mid point of 2.5 MeV which is close to the Q-value of double beta decay in ¹³⁰Te, 2.527 MeV.

Figure 6.4 shows the extracted 90% confidence limit half-life for fits which have various



Figure 6.5: The reduction in half-life at 90% confidence limit, once a systematic uncertainty of 10% on light yield is introduced, as a function of energy range of the region of interest centred at 2.5 MeV. The assumptions used are listed in Table 6.1.

bin sizes and energy ranges. The smaller the bin size the more sensitive the fit is. However there is little to be gained for a bin size below 40 keV. In terms of energy window size, no significant sensitivity is gained above 0.36 MeV with 40 keV bins in the fixed background case. A sensitivity of 9.60×10^{25} years at 90% CL is achieved with this energy window and bin size an 8.97% improvement over the single bin analysis.

6.2.1 Optimisation With Systematic Uncertainties

The choice of energy window affects the impact of systematic uncertainties on sensitivity to neutrinoless double-beta decay. Optimising the energy range once systematic uncertainties are included is complex. Many of the systematic uncertainties have degeneracies and correlations so an optimum energy range with one systematic uncertainty included is likely to be different once other systematic uncertainties are considered. An added complication is that the level of uncertainty that can be achieved on a systematic is dependent on the energy range used in external analysis.

To investigate the optimal energy range to use once systematic uncertainties are included, the impact of the energy resolution uncertainty was used. Figure 6.5 shows the reduction in half-life sensitivity at 90% CL as a function of energy window size when a 10% uncertainty in NHit/MeV has been introduced. This figure shows that the wider the energy window the less of an impact a systematic uncertainty on energy resolution has on sensitivity. This is due to the large number of background counts, particularly two neutrino double-beta decay, outside the range of the signal spectrum which act as a constraint on the NHit/MeV. Similar behaviour is observed for other systematic uncertainties and this is explored in the following section. The reduction in half-life plateaus at an energy range of around 1.36 MeV.

An energy range of 1.4 MeV (1.8 to 3.2 MeV) with 40 keV bins was chosen for the studies in this thesis to optimise the ability of the fit to reduce the impact of systematic uncertainties.

6.3 Application of systematic uncertainties

There are many systematic uncertainties which will effect the sensitivity to neutrinoless double-beta decay. This thesis focuses on background rate uncertainties and energy systematics: resolution, scale and shift. The techniques used to incorporate systematic uncertainties into the analysis are described in chapter 5. A two stage approach for evaluating systematic uncertainties is taken. Firstly, systematic uncertainties are investigated individually to examine the reductions in half-life sensitivity as a function of the size of the uncertainty. Secondly, they are combined with other systematic uncertainties to allow for the effect of correlations between the different systematic uncertainties on half-life sensitivity.

A background rate systematic uncertainty was viewed as important if the reduction in half-life was greater than the reduction in half-life associated with the 4% uncertainty, determined by the SNO experiment, on elastic scattering ⁸B solar neutrinos. The most important rate systematics are included with the energy systematic uncertainties to give the final systematic uncertainty. Systematics are removed individually to obtain the reduction in sensitivity associated with that systematic with all correlations included.

Four different analyses were used in the sensitivity studies as defined and motivated in Table 6.4. These analyses allow the effect of systematic uncertainties and correlations to be studied as a function of the size of their energy window. It also allows the impact of systematic uncertainties on the previous single binned analysis technique to be examined.

It is also important to examine these different energy ranges as there may be external analysis reasons why a wide 1.4 MeV window is not used. This may be due to the use of side regions outside the energy range applied in the fit which are used to extract the prior

	Counts		
Background	$2.2 \le E \le 2.42$	$2.42 \le E \le 2.64$	$2.64 \le E \le 2.8$
130 Te $(2\nu 2\beta)$	1192	30.53	0.1125
^{234m} Pa	214.8	1.379	2.827×10^{-4}
⁸ B (ES ν)	35.72	34.21	23.75
²¹² BiPo	15.42	12.67	8.116
²²⁸ Ac	13.83	7.456	0.2841
$2223 \text{ keV } \gamma$	18.04	0.2893	6.559×10^{-5}
²¹⁰ Tl	6.109	5.962	4.555
²⁰⁸ Tl Hold Down Ropes	0.8500	4.325	3.075
²⁰⁸ Tl AV	0.4986	4.238	2.992
²⁰⁸ Tl AV Outerdust	0.3497	1.524	1.249
²¹⁴ Bi External Water	1.909	0.7953	0.1591
PMT β - γ	0.3823	1.784	0.5098
²⁰⁸ Tl External Water	0.3094	1.309	0.9757
²¹⁴ BiPo	0.4158	0.7404	1.073
²¹⁴ Bi AV	2.072	0.	0.
²⁰⁸ Tl	0.	1.745×10^{-2}	1.5189
¹²⁴ Sb	0.5602	0.4552	0.1517
²⁰⁸ Tl Hold Up Ropes	0.1000	0.3501	0.3251
⁸⁸ Y	1.364×10^{-2}	0.2332	0.2100
²¹⁴ Bi AV Outerdust	0.2509	0.1254	0.7994
²¹⁴ Bi Hold Down Ropes	0.1750	0.1500	5.001×10^{-2}
^{110m} Ag	1.125×10^{-2}	2.805×10^{-2}	4.228×10^{-2}
102mRh	7.311×10^{-2}	4.084×10^{-4}	3.857×10^{-10}
²⁰⁸ Tl AV Innerdust	0.	2.505×10^{-2}	2.505×10^{-2}
²² Na	2.570×10^{-2}	1.855×10^{-2}	2.706×10^{-3}
⁶⁰ Co	8.380×10^{-3}	1.764×10^{-2}	3.739×10^{-3}
⁶⁸ Ga	5.512×10^{-3}	2.259×10^{-3}	4.201×10^{-3}
102Rh	4.407×10^{-3}	5.777×10^{-5}	1.079×10^{-8}
46Sc	3.450×10^{-3}	4.919×10^{-5}	3.031×10^{-8}
42K	3.934×10^{-4}	3.202×10^{-4}	1.788×10^{-4}
106Rh	3.504×10^{-4}	2.689×10^{-4}	1.431×10^{-4}
⁸² Rb	1.470×10^{-4}	1.488×10^{-4}	1.049×10^{-4}
44 Sc	7.307×10^{-5}	8.875×10^{-5}	6.138×10^{-5}
$^{126}\mathrm{Sb}$	2.317×10^{-5}	3.723×10^{-5}	3.878×10^{-5}
$^{110}\mathrm{Ag}$	6.014×10^{-5}	2.503×10^{-5}	4.771×10^{-6}
⁸⁴ Rb	6.925×10^{-5}	1.295×10^{-5}	4.246×10^{-7}
⁵⁸ Co	3.178×10^{-5}	1.354×10^{-7}	7.409×10^{-12}
⁵⁶ Co	1.478×10^{-6}	1.026×10^{-6}	1.350×10^{-5}
$^{126m}\mathrm{Sb}$	2.102×10^{-6}	2.057×10^{-6}	1.323×10^{-6}
⁹⁰ Y	1.988×10^{-6}	1.578×10^{-8}	4.954×10^{-12}
²⁶ Al	9.857×10^{-11}	7.961×10^{-8}	3.799×10^{-8}

Table 6.3: A list of all backgrounds in descending order of their expected count rate for three energy ranges with the assumptions in Table 6.1.

Lower Bound (MeV)	Upper Bound (MeV)	Bin Size	Motivation
2.42	2.64	Single Bin	Used in previous independent analysis
2.42	2.64	20 keV	Introduces binning
2.2	2.8	20 keV	Introduces binning and entire signal region
1.8	3.2	40 keV	Optimised analysis, see Section 6.2.1

Table 6.4: The different energy ranges and bin sizes for the four analysis approaches. A motivation is provided for each selected range and bin size.

values and associated uncertainties for each of the systematics. This side region method allows the extracted values for priors and uncertainties not to be biased by any fluctuations in the signal region as the signal counts are not included in these fits. It also maintains independence between the analyses. These obtained prior values can then be used in the fit used to extract an associated confidence level for the neutrinoless double-beta decay half-life.

Only the energy spectrum is fit for the results presented in this section. The experimental parameters used are as given in Table 6.1. RAT 5.0.2 software and simulation data sets were used in the following studies. Section 6.4 extends the fit so energy and radius parameters are both included in the fit.

6.3.1 Individual Systematics

Each systematic uncertainty is investigated individually, over a range of uncertainties, using a grid search to ensure the global minimum is found in each case. The important rate systematics are defined as those that reduce the half-life sensitivity by more than the reduction observed for the ⁸B neutrino elastic scattering rate systematic.

Energy Resolution

As previously explained in Section 5.2.1, SNO+ is a single photon counting experiment and the number of photons produced by the scintillator is directly proportional to the quenched energy deposited in the scintillator. The energy resolution is therefore directly related to the light yield, L, as measured in NHit/MeV. Assuming the energy response is Gaussian the energy resolution is related to the L for a given quenched energy E by $\sigma = \sqrt{E/L}$. To keep the energy dependence of energy resolution separate from the uncertainty in energy resolution, it is investigated in terms of an uncertainty on L.

A range of 180 to 220 NHit/MeV in steps of 1 NHit/MeV was used to probe the uncertainty in energy resolution. The prior value of L was 200 NHit/MeV. The uncertainty



Figure 6.6: Top: the half-life at 90% confidence limit as a function of uncertainty in light yield, L, measured in NHit/MeV. Bottom: the reduction in half-life sensitivity as a function of uncertainty in L. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line in the bottom plot corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.



Figure 6.7: A graph showing the best fits at the 90% confidence limit as a function of uncertainty in energy resolution. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins.

on energy resolution was defined in terms of NHit/MeV and converted into a percentage. For example an uncertainty of 20 NHit/MeV relates to a 10% uncertainty, σ_L .

Figure 6.6 shows the value and reduction in neutrinoless double-beta decay half-life sensitivity for the four different analyses described in Table 6.4. KamLAND-Zen quoted their uncertainty in L to be 4.54% [117] so assuming SNO+ achieves the same level of uncertainty, the reduction in half-life sensitivity is 3.86%, 3.0%, 3.89% and 1.09% for the four analyses in Table 6.4 (single bin through to wide energy window respectively). Figure 6.6 shows that the 1.8 to 3.2 MeV is the best of the four analyses to use for the entire uncertainty range in energy resolution. The benefit of using a wide energy window is clearly demonstrated at 10% uncertainty on L where the single bin analysis has its half life sensitivity reduced by 13.6% and the wide energy window is reduced only by 1.44%. This is due to the increased amount of shape information in the wide energy window acting to constrain the fit.

Above 5.5% there is clear separation of reduction in sensitivity for the four analyses. Figure 6.7 shows that this is due to the wider energy range allowing the energy resolution to be most strongly constrained to lower light yields than the smaller energy range analyses. A higher light yield will lead to events from the two neutrino double-beta spectrum shifting towards lower energies. Therefore, the increased signal counts observed at the 90% CL can be compensated for by shifting two neutrino double-beta decay events outside the signal energy spectrum. However, there will also be an observed increase in two neutrino events at energies below the Q-value of ¹³⁰Te as the energy spectrum becomes steeper. Therefore, for the wider energy range analyses, the differing number of counts per energy bin in the background-only case at the prior light yield and the background-and-signal case at higher light yields becomes significant at energies outside the signal energy spectrum peak. The difference in counts in these energy bins outside the peak of the signal region constrain the fit to lower light yields when the wider energy range analyses, higher light yields are preferred in order to compensate for the signal counts. The penalty term stops the analysis fitting to too high values of light yield in this case through the addition of χ^2 as stated in Equation 5.7.

Energy Scale

As stated in Section 5.2.2, an energy scale uncertainty is incorporated into the analysis using a scale factor which is multiplied with the quenched energy. This uncertainty can arise from variations in calibration and detector stability issues which include potential changes to the number of online PMTs, scintillator degradation and changes to trigger thresholds.

A range of 0.95 to 1.05 in steps of 0.0001 was used to investigate the energy scale with a prior value of 1. Figure 6.8 shows the half-life and reduction in neutrinoless double-beta decay half-life sensitivity for the four analyses as described in Table 6.4. It can be seen in these plots that energy scale is a dominant systematic for the smaller, 2.42 - 2.64 MeV, analyses. Figure 6.9 shows that, for analyses with this energy range, that the energy scale is pulled by the fit to lower values and constrained towards the prior by the penalty term. Lower energy scales behave in a similar way to higher light yields. With lower energy scales, events from the two neutrino double-beta decay spectrum are shifted to lower energies. Hence the number of two neutrino double-beta decay background events between 2.42 and 2.64 MeV will decrease to compensate for any increase in the number of signal counts in this region. The energy spectra in the fit also becomes steeper with lower values of energy scale. Due to the added shape information in analyses with an energy



Figure 6.8: Top: the half-life at 90% confidence limit as a function of uncertainty in energy scale. Bottom: the reduction in half-life sensitivity as a function of uncertainty in energy scale. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line in the bottom plot corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.



Figure 6.9: A graph showing the best fits at the 90% confidence limit as a function of uncertainty in energy scale. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins.

range of 2.2 to 2.8 MeV or above, the fits are constrained to values close to the prior as the difference in counts in the energy bins outside the 2.42 to 2.64 MeV range is too large. Figure 6.8 shows that the reduction in sensitivity is controlled and less important for the analyses with wider energy windows.

KamLAND-Zen quoted their systematic uncertainty in energy scale to be 0.3% [117] and assuming the same level of systematic uncertainty can be achieved for SNO+ the reduction in half-life sensitivity in this study are 5.08%, 3.64%, 1.09% and 0.727% for the four analyses in Table 6.4 (single bin through to wide energy window respectively).

Energy Shift

As described in Section 5.2.3, energy shift is where a constant offset is added to the quenched energy deposited in the scintillator. Its value is dependent on the average noise hits included in an event being too high or low in the Monte Carlo simulation.

A range of -0.1 to 0.1 MeV in steps of 0.0001 MeV was used to investigate the energy scale. A prior value of zero shift was used.



Figure 6.10: Top: the half-life at 90% confidence limit as a function of uncertainty in energy shift. Bottom: the reduction in half-life sensitivity as a function of uncertainty in energy shift. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line in the bottom plot corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.



Figure 6.11: A graph showing the best fits at the 90% confidence limit as a function of uncertainty in energy shift. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins.

Figure 6.10 shows the half-life and reduction in neutrinoless double-beta decay half-life sensitivity for the four analyses stated in Table 6.4. Clearly this systematic has a large impact on sensitivity for the analyses with the smaller energy range, particularly in the single bin case. Figure 6.11 shows that, for the analyses with the smaller energy ranges, negative values of energy shift are preferred at the 90% CL of neutrinoless double-beta decay. Negative values of energy shift have the effect of shifting events to lower energies. This is preferred as the largest number of events are at energies lower than the Q-value of ¹³⁰Te. In the single bin case there is no shape information so any increase in signal counts at the 90% CL can be compensated by shifting an amount of the two neutrino double-beta double-beta and signal spectrum out of the energy range. The energy shift is therefore only constrained by the penalty term in this case. A similar situation is found in the binned 2.42 to 2.64 MeV fit although the effect is not as great as the shape information provides some constraint on the fit. It is clear from these plots that an analysis with an energy range between 2.2 and 2.8 MeV or larger is required in order to constrain the energy shift. This occurs because more shape information from the two neutrino double-beta spectrum

is contained in the fit which constrains it. It can be seen in Figure 6.11 that the analysis with the 1.8 to 3.2 MeV energy range is constrained to -0.2 keV for the entire range of uncertainty.

KamLAND-Zen quoted their systematic uncertainty in energy shift at 2.5 MeV to be 0.05 MeV [118] and assuming the same level of uncertainty can be achieved for SNO+ the reduction in half-life sensitivity in this study are 69.92%, 48.9%, 0.73% and 0.36% for the four analyses in Table 6.4 (single bin through to wide energy window respectively).

¹³⁰Te Two Neutrino Double-Beta Decay Rate

An uncertainty in the two neutrino double-beta decay rate occurs from uncertainties in the level of loading and uncertainties in the decay rate itself. A range of 1.5σ to -1.5σ in 30000 steps was used to investigate the ¹³⁰Te double beta decay rate. A prior value of 18621266 decays was used which is the expected number of decays after five years livetime with 0.3% Te loading.

Figure 6.12 shows the half-life and reduction in neutrinoless double-beta decay halflife sensitivity for the four analyses as stated in Table 6.4. From these figures it is clear that the smaller the energy range the more of an impact on sensitivity there is due to an uncertainty of two neutrino double-beta decay. Figure 6.13 shows the best fits of ¹³⁰Te double beta decay rate divided by its prior value as a function of its uncertainty at the 90% confidence limit. From this it can be seen that lower values of rate are preferred by the fit for the analyses with the smaller energy ranges. These lower rates compensate for the increase in counts from the signal contribution. The 2.2 to 2.8 MeV is constrained to close to the prior value and the 1.8 to 3.2 MeV is constrained to the prior value as more of the shape information is contained in the fit. The analyses with the smaller energy range, by contrast, fits to lower rates because of the absence of this shape information and the fit is being constrained by the penalty term instead.

The leading result in 130 Te double beta decay rate is by the NEMO3 experiment and has an associated uncertainty of 20.3% [119] so assuming that SNO+ achieves the same level of uncertainty the reduction in half-life sensitivity in this study are 14.3%, 9.35%, 1.81% and 0.0% for the four analyses in Table 6.4 (single bin through to wide energy window respectively).


Figure 6.12: Top: the half-life at 90% confidence limit as a function of uncertainty in the 130 Te double beta decay rate. Bottom: the reduction in half-life sensitivity as a function of uncertainty in the 130 Te double beta decay rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line in the bottom plot corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.



Figure 6.13: A graph showing the best fits at the 90% confidence limit as a function of uncertainty in the 130 Te double beta decay rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins.

⁸B Neutrino Elastic Scattering Rate

The ⁸B Neutrino Elastic Scattering Rate is a relatively flat background across the energy ranges considered. A range of 1.5σ to -1.5σ in 30000 steps was used to investigate the ⁸B neutrino elastic scattering rate. A prior value of 5105 decays was used which is the expected number of decays after five years livetime with 0.3% Te loading.

Figure 6.14 shows the half-life and reduction in neutrinoless double-beta decay half-life sensitivity for the four analyses. Similar reductions in half-life are observed for the four analyses considered in Table 6.4. As ⁸B Neutrino Elastic Scattering Rate is a relatively flat background there is little to be gained from binning the energy ranges. It can be seen in Figure 6.15 that lower values of the rate are preferred which compensate for the increase in number of counts in the signal region at the 90 % CL. It can also be seen that approximately the same rates are preferred by the binned and single bin fit in the analysis with the 2.42 to 2.64 MeV energy range due to the flat nature of the background. A lower, approximately the same rate is preferred for the analyses with the 2.2 to 2.8 MeV and 1.8



Figure 6.14: Top: the half-life at 90% confidence limit as a function of uncertainty in the ⁸B neutrino elastic scattering rate. Bottom: the reduction in half-life sensitivity as a function of uncertainty in the ⁸B neutrino elastic scattering rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line in the bottom plot corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.



Figure 6.15: A graph showing the best fits at the 90% confidence limit as a function of uncertainty in the ⁸B neutrino elastic scattering rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins.

to 3.2 MeV energy ranges due to the entire signal region being included in these fits.

The ⁸B neutrino elastic scattering rate was measured by the SNO experiment and has an associated uncertainty of 4% [120] so assuming that SNO+ achieves the same level of uncertainty, the reduction in half-life sensitivity in this study are 0.993%, 0.683%, 1.09% and 1.09% for the four analyses in Table 6.4 (single bin through to wide energy window respectively).

Other Rate Systematics

The other systematics which were considered in this study are ²¹²Bi²¹²Po, ²¹⁴Bi²¹⁴Po, thorium-chain (no ²¹²Bi²¹²Po), uranium-chain (no ²¹⁴Bi²¹⁴Po), external backgrounds and 2223 keV gammas, from neutron capture on ¹H nuclides, rates. The Bi-Pos were considered separately due to the added uncertainty from analysis cuts. A systematic was considered in further studies if its reduction in half-life sensitivity was greater than the reduction due to a 4% uncertainty in ⁸B neutrino elastic scattering rate, 1.09%.

Figure 6.16 shows the reduction in half-life due to uncertainty in the $^{212}Bi^{212}Po$ rate.



Figure 6.16: A graph showing the reduction in half-life sensitivity as a function of uncertainty in the $^{212}\text{Bi}^{212}\text{Po}$ pileup rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.

A 12% uncertainty is required in the rate before it is dominant over the ⁸B neutrino elastic scattering rate. A greater reduction in half-life is observed for the analyses with a larger energy range due to the larger number of counts per bin between 2.2 and 2.42 MeV, see Table 6.3. This will compensate for the increase in signal counts below 2.42 MeV at the 90% CL sensitivity to neutrinoless double-beta decay.

No reduction in sensitivity is observed for the binned regions of interest up to the 25% uncertainty investigated for the ²¹⁴Bi²¹⁴Po rate systematic uncertainty. This is due to the low number of counts expected from this background in the regions of interest, see Table 6.3. As a result this rate will be fixed in further studies.

Figure 6.17 shows the reduction in half-life due to uncertainty in the thorium-chain rate without the ²¹²Bi²¹²Po background. A 22% uncertainty is required in the rate before it is dominant over the ⁸B neutrino elastic scattering rate for the analysis with the 2.2 to 2.8 MeV energy range. A greater reduction in half-life is observed for the 2.2 to 2.8 MeV



Figure 6.17: A graph showing the reduction in half-life sensitivity as a function of uncertainty in the Th chain rate with the 214 Bi²¹⁴Po pileup fixed. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.

due to the larger number of counts per bin between 2.2 and 2.42 MeV, see Figure 6.1 which compensates for the increase in signal counts in this energy range. The thorium-chain rate is constrained for all uncertainties considered for the analysis with the 1.8 to 3.2 MeV energy range. This is due to the thorium-chain becoming dominant above 2.86 MeV and an increase in the rate below approximately 2.2 MeV. Both of which contribute to the fit being constrained in the analysis with the 1.8 to 3.2 MeV energy window.

Figure 6.18 shows the reduction in half-life due to uncertainty in the uranium-chain rate without the ²¹⁴Bi²¹⁴Po background. A 24% uncertainty is required in the rate before it is dominant over the ⁸B neutrino elastic scattering rate for the analysis with the 2.2 to 2.8 MeV energy range. A similar reduction is observed in all analyses apart from the analysis with the 1.8 to 3.2 MeV energy range where the uranium-chain rate is constrained for all uncertainties considered. The fit is constrained for the 1.8 to 3.2 MeV analysis and not the others considered due to the significant number of counts from the uranium-chain



Figure 6.18: A graph showing the reduction in half-life sensitivity as a function of uncertainty in the U chain rate with $^{214}\text{Bi}^{214}\text{Po}$ pileup fixed. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.

for energies between 1.8 and 2.2 MeV, see Figure 6.1.

Figure 6.19 shows the reduction in half-life due to uncertainty in the external background rate. An 11% uncertainty is required in the rate before it is dominant over the ⁸B neutrino elastic scattering rate. A similar reduction is observed for the analyses with the 2.2 to 2.8 MeV and 1.8 to 3.2 MeV energy ranges. This is because the external background spectrum peaks in the energy range between 2.2 to 2.8 MeV where the signal counts occur and there is a less significant contribution for energies outside of this window which is why the analysis with the 1.8 to 3.2 MeV energy range performs marginally better at greater uncertainties. the analyses with the smaller energy range performs marginally better than the other analyses as less of the external energy spectrum is included in the fit, see Figure 6.1.

No reduction in half-life is observed up to a 25% uncertainty in 2223 keV gamma rate for all regions of interest. This is due to the low number of counts expected in the signal



Figure 6.19: A graph showing the reduction in half-life sensitivity as a function of uncertainty in the external backgrounds rate. Four analyses are shown, the black spots show a singled binned region between 2.42 and 2.64 MeV; the red diagonal crosses show a binned 2.42 and 2.64 MeV region, the vertical blue crosses show a binned 2.2 and 2.8 MeV region both with 20 keV bins and the magenta diamonds shows a region of 1.8 to 3.2 MeV with 40 keV bins. The dashed red line corresponds to a 1.09% reduction in half-life sensitivity due to a 4% uncertainty in the ⁸B neutrino elastic scattering rate in the 1.8 to 3.2 MeV analysis.

region, see Figure 6.1 and Table 6.3. This background will be fixed in further studies.

6.3.2 Correlations Between Systematic Uncertainties

The correlations between systematic uncertainties are probed using a grid search minimiser which allows the χ^2 space to be mapped out between a pair of systematic uncertainties. This is achieved by obtaining a χ^2 at each combination of values of the systematics. The signal is included in the fit at a fixed half-life and uncertainties on the systematics are fixed. The distributions presented are with the signal half-life fixed at the 90% CL from the fit with the pair of systematic uncertainties under investigation. Relatively large systematic uncertainties were used to investigate the correlations such that the χ^2 contribution of the penalty term is reduced and the parameters are more free to change and so allow the underlying correlations to be observed. The full correlation matrix will be examined in Section 6.3.3.

Energy Resolution and Energy Scale

Energy resolution and energy scale systematics behave similarly. Both systematics move two neutrino double-beta events in or out of the signal region. Due to this it is likely that these parameters are correlated, to some extent, in the fit.

The limit obtained for the analysis with 2.2 to 2.8 MeV energy range when light yield and energy scale are both floated at a 10% uncertainty is 9.08×10^{25} years at 90% CL corresponding to a reduction in neutrinoless double-beta decay half-life sensitivity of 5.2%. The top plot in Figure 6.20 shows the χ^2 space for light yield and energy scale with a signal present at this 90% CL half-life. The figure shows a positive correlation between energy resolution and scale where at higher light yields a higher energy scale value is favoured. The two parameters are therefore competing with each other with the energy resolution parameter moving two neutrino events outside the signal region whilst the scale parameter moves the two neutrino events back into the signal region. The best fits, marked with a white cross in the figure, are a light yield of 206 NHit/MeV and $E_{scale} = 1.0005$. In comparison the best fits in the case where each individual parameter is floated one at a time at 10% uncertainty are a light yield of 203 NHit/MeV and $E_{scale} = 0.9994$. The increase in best fit to the energy resolution is due to the positive correlation and the two parameters effectively working against each other. The χ^2 space is smooth and continuous.

The limit obtained for the analysis with 1.8 to 3.2 MeV energy range when light yield and energy scale are both floated at a 10% uncertainty is 9.55×10^{25} years at 90% CL corresponding to a reduction in neutrinoless double-beta decay half-life sensitivity of 1.4%. The bottom plot in Figure 6.20 shows the χ^2 space for light yield and energy scale with a signal present at this 90% CL half-life. The figure shows there is a reduction in the positive correlation with this energy window due to the added shape information included in the fit, particularly the large number of two neutrino events at energies below 2.2 MeV, acting to constrain the two parameters. The best fits, marked with a white cross in the figure, are a light yield of 201 NHit/MeV and $E_{scale} = 1$. The best fits obtained when both parameters are floated are consistent with what was observed when the systematics were investigated individually which also demonstrates that the correlation has much less of an effect when the 1.8 to 3.2 MeV energy range is used. The reduction in correlation with respect to widening the energy range is observed in other pairs of systematic uncertainties. The analysis with the 2.2 to 2.8 MeV window will be used from now on to show the correlation in the entire signal region.



Figure 6.20: A two dimensional histogram showing the χ^2 space of energy scale on the y-axis and energy resolution on the x-axis at the 90% confidence limit. Top: The analysis with an energy range of 2.2 to 2.8 MeV and 20 keV bins were used in the fit. Bottom: The analysis with an energy range of 1.8 to 3.2 MeV and 40 keV bins were used in the fit. There is a 10% uncertainty in energy scale and energy resolution. The best fit of the two parameters is marked with a white cross.



Figure 6.21: A two dimensional histogram showing the χ^2 space of energy shift on the y-axis and energy resolution on the x-axis at the 90% confidence limit. There is a 10% uncertainty in energy resolution and a 0.1 MeV uncertainty in energy shift. The best fit of the two parameters is marked with a white cross.

Energy Resolution and Energy Shift

The energy resolution parameter, when investigated individually, prefers to fit to higher values which reduces the smearing of the energy. This effectively moves two neutrino double-beta decay events out of the signal region. Clearly the energy shift will exhibit similar behaviour by moving two neutrino double-beta events into or out of the signal region and so there is likely to be a correlation.

The limit obtained for the analysis with 2.2 to 2.8 MeV energy range when light yield and energy shift are floated at a 10% and 0.1 MeV uncertainty respectively is 9.08×10^{25} years at 90% CL corresponding to a reduction in neutrinoless double-beta decay halflife sensitivity of 5.2%. Figure 6.21 shows the χ^2 space for light yield and energy shift with a signal present at this 90% CL half-life. The figure shows a positive correlation between energy resolution and shift where at higher light yields a higher value in energy shift is favoured. The best fits, marked with a white cross in the figure, are a light yield of 209 NHit/MeV and $E_{shift} = 2$ keV. Here the resolution parameter is shifting two neutrino double-beta decay events out of the signal region whereas the shift parameter is attempting to compensate for the energy resolution parameter by shifting two neutrino double-beta



Figure 6.22: A two dimensional histogram showing the χ^2 space of ¹³⁰Te double beta decay rate normalised to its prior on the y-axis and energy resolution on the x-axis at the 90% confidence limit. There is a 10% uncertainty in energy resolution and a 20.3% uncertainty in ¹³⁰Te double beta decay rate. The best fit of the two parameters is marked with a white cross.

events back into the signal region. In comparison the best fits in the case where each individual parameter is floated at 10% and 0.1 MeV uncertainty respectively are a light yield of 203 NHit/MeV and $E_{shift} = -1$ keV. The χ^2 space does not appear smooth and continuous at lower energy shifts. This could pose problems for other minimisers.

Energy Resolution and ¹³⁰Te Two Neutrino Double-Beta Decay Rate

The energy resolution parameter compensates for an increasing signal by reducing the smearing on the two neutrino double-beta energy spectrum and shifting background events out of the signal region. Due to this it is likely that the energy resolution and the ¹³⁰Te two neutrino double-beta decay rate are correlated.

The limit obtained for the analysis with 2.2 to 2.8 MeV energy range when light yield and two neutrino double-beta decay rate are floated at a 10% and 20.3% uncertainty respectively is 8.84×10^{25} years at 90% CL corresponding to a reduction in neutrinoless double-beta decay half-life sensitivity of 7.8%. Figure 6.22 shows the χ^2 space for light yield and two neutrino double-beta decay rate with a signal present at this 90% CL half-



Figure 6.23: A two dimensional histogram showing the χ^2 space of energy shift on the y-axis and energy scale on the x-axis at the 90% confidence limit. There is a 10% uncertainty in energy scale and a 0.1 MeV uncertainty in energy shift. The best fit of the two parameters is marked with a white cross.

life. The figure shows a positive correlation between energy resolution and rate where at higher light yields, a higher value in two neutrino double-beta decay rate is favoured. The χ^2 space is smooth and continuous. The best fits, marked with a white cross in the figure, are a light yield of 208 NHit/MeV and (¹³⁰Te $2\nu\beta\beta$ rate)/prior=1.02842. Here the two neutrino double-beta decay rate is compensating for the increase in light yield by increasing the number of two neutrino double-beta decay events in the signal region whilst the light yield parameter is attempting to shift the events out of the signal region. In comparison the best fits in the case where each individual parameter is floated at 10% and 20.3% rate uncertainty respectively are a light yield of 203 NHit/MeV and (¹³⁰Te $2\nu\beta\beta$ rate)/prior=0.9903778.

Energy Scale and Energy Shift

Not only does the energy scale change the shape of the energy spectra, it also shifts the background events into or out of the signal region. For this reason it is likely that energy scale and shift are correlated.

The limit obtained for the analysis with 2.2 to 2.8 MeV energy range when energy

scale and energy shift are floated at a 10% and 0.1 MeV uncertainty respectively is $9.51 \times$ 10^{25} years at 90% CL corresponding to a reduction in neutrinoless double-beta decay halflife sensitivity of 0.73%. Figure 6.23 shows the χ^2 space for energy scale and energy shift with a signal present at this 90% CL half-life. The figure shows a negative correlation between energy scale and shift where at higher values of scale, lower values of energy shift are favoured and vice versa. The χ^2 space is smooth and continuous but also asymmetric around an energy shift of zero. This asymmetry arises from the anisotropy in background counts across the energy range. For negative shifts the two neutrino double-beta decay events are being moved out and thorium chain events are being moved in to the energy range and vice versa for positive shifts. Due to this the amount of scaling needed to compensate for the shift changes. The best fits, marked with a white cross in the figure, are $E_{scale} = 0.998$ and $E_{shift} = 2$ keV. Here the energy scale is reducing the smearing on the energy spectra and shifting events out of the signal region whilst the energy shift is trying to compensate for this by shifting the two neutrino double-beta decay events back into the signal region. In comparison the best fits in the case where each individual parameter is floated at 10% and 0.1 MeV uncertainty respectively are $E_{scale} = 0.9994$ and $E_{shift} = -1$ keV.

Energy Scale and ¹³⁰Te Two neutrino Double-Beta Decay Rate

As the energy scale compensates for the increased signal counts at the 90% CL by shifting two neutrino double-beta decay events out of the signal region it is likely that the energy scale and the 130 Te two neutrino double-beta decay rate are correlated to some extent.

The limit obtained for the analysis with 2.2 to 2.8 MeV energy range when energy scale and two neutrino double-beta decay rate are floated at a 10% and 20.3% uncertainty respectively is 9.02×10^{25} years at 90% CL corresponding to a reduction in neutrinoless double-beta decay half-life sensitivity of 5.9%. Figure 6.24 shows the χ^2 space for energy scale and two neutrino double-beta decay rate with a signal present at this 90% CL half-life. The figure shows a negative correlation between energy scale and rate where at higher scale, a lower value in rate is favoured and vice versa. The χ^2 space appears smooth and continuous. The best fits, marked with a white cross in the figure, are $E_{scale} = 0.9972$ and $(^{130}\text{Te } 2\nu\beta\beta \text{ rate})/\text{prior}=1.0609$. Here the energy scale is making the two neutrino double-beta spectrum less smeared which effectively shifts the two neutrino events outside the signal region, the two neutrino double-beta decay rate by contrast is trying to compensate for this effect by increasing the number of two neutrino double-beta decay counts in the



Figure 6.24: A two dimensional histogram showing the χ^2 space of ¹³⁰Te double beta decay rate normalised to its prior on the y-axis and energy scale on the x-axis at the 90% confidence limit. There is a 10% uncertainty in energy scale and a 20.3% uncertainty in ¹³⁰Te double beta decay rate. The best fit of the two parameters is marked with a white cross.

signal region. For comparison, the best fits in the case where each individual parameter is floated at 10% and 20.3% rate uncertainty respectively are $E_{scale} = 0.9994$ and (¹³⁰Te $2\nu\beta\beta$ rate)/prior= 0.9903778.

Energy Shift and ¹³⁰Te Two neutrino Double-Beta Decay Rate

The energy shift parameter compensates for the increase in counts from the signal at the 90% CL sensitivity by shifting two neutrino double-beta decay events outside the signal region. For this reason it is likely that energy shift and 130 Te two neutrino double-beta decay rate are correlated.

The limit obtained for the analysis with 2.2 to 2.8 MeV energy range when energy shift and two neutrino double-beta decay rate are floated at a 0.1 MeV and 20.3% uncertainty respectively is 9.34×10^{25} years at 90% CL corresponding to a reduction in neutrinoless double-beta decay half-life sensitivity of 2.5%. Figure 6.25 shows the χ^2 space for energy shift and two neutrino double-beta decay rate with a signal present at this 90% CL halflife. The figure shows a negative correlation between energy shift and rate where at higher



Figure 6.25: A two dimensional histogram showing the χ^2 space of ¹³⁰Te double beta decay rate normalised to its prior on the y-axis and energy shift on the x-axis at the 90% confidence limit. There is a 0.1 MeV uncertainty in energy shift and a 20.3% uncertainty in ¹³⁰Te double beta decay rate. The best fit of the two parameters is marked with a white cross.

scale, a lower value in rate is favoured and vice versa. The χ^2 space appears smooth and continuous. The best fits, marked with a white cross in the figure, are $E_{shift} = -4$ keV and (¹³⁰Te $2\nu\beta\beta$ rate)/prior=1.02842. Here the energy shift is compensating for the signal counts at the 90% CL sensitivity by shifting two neutrino double-beta decay counts outside the signal region, the two neutrino double-beta decay rate then compensates for the effect of the energy shift by increasing the number of two neutrino double-beta decay counts. For comparison, the best fits in the case where each individual parameter is floated at 10% and 20.3% rate uncertainty respectively are $E_{shift} = -1$ keV and (¹³⁰Te $2\nu\beta\beta$ rate)/prior= 0.9903778.

⁸B Neutrino Elastic Scattering Rate

As the ⁸B spectrum is a relatively flat background and the energy parameters effect the shape of the energy spectra it is unlikely a priori that there is very much, if at all, correlation between this systematic uncertainty and the energy systematic uncertainties. The correlations between all systematics are quantified in the following section.

Systematic Uncertainty	Optimistic Value	Conservative Value
Light Yield	4.54%	10%
Energy Scale	0.3%	2%
Energy Shift	50 keV	125 keV
Two neutrino double-beta Decay Rate	10%	20.3%
⁸ B Neutrino Elastic Scattering Rate	4%	4%
$^{212}\text{Bi}^{212}\text{Po}$ Rate	10%	50%
Thorium-Chain Rate	10%	50%
Uranium-Chain Rate	10%	50%
External Backgrounds Rate	10%	50%

Table 6.5: Values for each systematic uncertainty for the two scenarios investigated, an optimistic scenario and a conservative systematics scenario.

6.3.3 Combined Systematic Uncertainties Study

To assess the impact on sensitivity to neutrinoless double-beta decay half-life for each systematic uncertainty when the correlation between other systematics is accounted for two cases of systematic uncertainties are compared. The first case is the sensitivity when all systematic uncertainties are included in the fit. The second case is the sensitivity when one systematic uncertainty has been removed.

The difference, in quadrature, between the reduction in sensitivity, r, with all systematic uncertainties included and the reduction where one systematic uncertainty has been removed gives the reduction in half-life sensitivity for that systematic

$$r_{syst} = \sqrt{r_n^2 - r_{n-1}^2}.$$
 (6.1)

Due to the number of systematic uncertainties included in the study, MINUIT is used to minimise the fit. MINUIT also outputs a correlation matrix which allows the correlations between the different systematics to be quantified.

The effect of systematics on two analyses are examined in this study: firstly, the analysis with energy range 2.2 to 2.8 MeV and 20 keV bins; and secondly the analysis with an energy range of 1.8 to 3.2 MeV and 40 keV bins. For each analysis two case studies are examined with different systematic uncertainties: one with conservative values on systematic uncertainties and one with a more optimistic scenario. The systematic uncertainties are given in Table 6.5. The systematic uncertainties for the optimistic uncertainties are taken as those that KamLAND-Zen achieved where available.

Table 6.6: The correlation coefficient matrix for the analysis with a 2.2 to 2.8 MeV energy window with 20 keV bins in the optimistic systematics scenario. Global is the overall correlation of the parameter with all other parameters. LY is light yield which is directly related to energy resolution. Scale and shift are the energy scale and shift parameters. $2\nu 2\beta$, ⁸B ν , ²¹²BiPo, U, Th and Ext correspond to the ¹³⁰Te two neutrino double-beta decay, ⁸B neutrino elastic scattering, ²¹²Bi²¹²Po, uranium, thorium-chain and external background rate uncertainties respectively.

	Global	LY	Scale	Shift	$2\nu 2\beta$	$^{8}\mathrm{B} \nu$	²¹² BiPo	U	Th	Ext
LY	0.53560	1.000	-0.311	-0.085	0.319	-0.050	-0.047	-0.276	-0.029	-0.046
Scale	0.31180	-0.311	1.000	0.012	-0.091	0.011	0.010	0.083	0.006	0.009
Shift	0.92364	-0.085	0.012	1.000	-0.846	-0.064	-0.063	-0.238	-0.048	-0.054
$2\nu 2\beta$	0.92410	0.319	-0.091	-0.846	1.000	0.032	0.030	-0.122	0.020	0.034
$^{8}\mathrm{B} \nu$	0.15748	-0.050	0.011	-0.064	0.032	1.000	-0.030	-0.041	-0.016	-0.029
²¹² BiPo	0.15328	-0.047	0.010	-0.063	0.030	-0.030	1.000	-0.036	-0.015	-0.029
U	0.66949	-0.276	0.083	-0.238	-0.122	-0.041	-0.036	1.000	-0.018	-0.039
Th	0.10810	-0.029	0.006	-0.048	0.020	-0.016	-0.015	-0.018	1.000	-0.015
Ext	0.13490	-0.046	0.009	-0.054	0.034	-0.029	-0.029	-0.039	-0.015	1.000

Correlations Between Systematic Uncertainties

MINUIT outputs a correlation coefficient matrix which can be used to examine the correlations between the various parameters. A value of 1 corresponds to a complete positive correlation, -1 is a complete negative correlation and 0 corresponds to uncorrelated parameters. The correlation coefficient between the i^{th} and j^{th} parameters is

$$\rho_{ij} = \frac{V_{ij}}{\sqrt{V_{ii}V_{jj}}}.$$
(6.2)

where V is the covariance matrix. For details on how MINUIT calculates the covariance matrix see [111]. MINUIT also provides a global correlation coefficient which is the correlation between a parameter and the linear combination of all other parameters. It takes a value between zero and one for uncorrelated and highly correlated parameters respectively.

Table 6.6 shows the correlation coefficients for the 2.2 to 2.8 MeV analysis with the optimistic systematic scenario. One can see that the two neutrino double-beta decay rate is the most correlated with all other parameters closely followed by the energy shift. The parameters which have the next highest correlations are the uranium chain rate and light yield. This is followed by energy shift and the other background rate systematics which are uncorrelated. The highest correlation between two different systematic uncertainties is a negative correlation of coefficient value -0.846 between energy shift and the two neutrino double-beta decay rate.

Table 6.7 shows the correlation coefficients for the 2.2 to 2.8 MeV analysis with the

Table 6.7: The correlation coefficient matrix for the analysis with a 2.2 to 2.8 MeV energy window with 20 keV bins in the conservative systematics scenario. Global is the overall correlation of the parameter with all other parameters. LY is light yield which is directly related to energy resolution. Scale and shift are the energy scale and shift parameters. $2\nu 2\beta$, ⁸B ν , ²¹²BiPo, U, Th and Ext correspond to the ¹³⁰Te two neutrino double-beta decay, ⁸B neutrino elastic scattering, ²¹²Bi²¹²Po, uranium, thorium-chain and external background rate uncertainties respectively.

	Global	LY	Scale	Shift	$2\nu 2\beta$	$^{8}\mathrm{B} \nu$	²¹² BiPo	U	Th	Ext
LY	0.81025	1.000	-0.713	0.175	0.177	0.012	0.067	-0.105	0.042	0.026
Scale	0.81011	-0.713	1.000	-0.472	0.163	0.035	0.079	-0.020	-0.002	-0.027
Shift	0.88963	0.175	-0.472	1.000	-0.572	-0.022	-0.026	-0.017	-0.064	-0.032
$2\nu 2\beta$	0.94266	0.177	0.163	-0.572	1.000	0.033	0.075	-0.705	-0.023	0.089
$^{8}\mathrm{B} \nu$	0.18364	0.012	0.035	-0.022	0.033	1.000	-0.042	-0.036	-0.055	-0.091
²¹² BiPo	0.49567	0.067	0.079	-0.026	0.075	-0.042	1.000	-0.124	-0.096	-0.376
U	0.90388	-0.105	-0.020	-0.017	-0.705	-0.036	-0.124	1.000	0.026	-0.098
Th	0.26260	0.042	-0.002	-0.064	-0.023	-0.055	-0.096	0.026	1.000	-0.055
Ext	0.44564	0.026	-0.027	-0.032	0.089	-0.091	-0.376	-0.098	-0.055	1.000

conservative systematics scenario. All of the global correlation parameters for the systematic uncertainties are larger than the optimistic systematics scenario apart from the energy shift which decreases very slightly. It is expected that there will be more correlation in the conservative systematics scenario as the penalty term can vary more which weakens the constraints on the other systematic uncertainties. The energy shift best fit was -3.79 keV and -3.90 keV for the optimistic and conservative systematics scenarios respectively which is close to the prior value with respect to their uncertainties and the pull of penalty term will be negligible in both cases as it is a fraction of a χ^2 -unit away from the prior which may be why the increase in correlation is not observed for energy shift. The systematic uncertainties with the highest correlation are light yield and energy scale with a correlation of -0.713 and is due to the similar behaviour of the systematics on the energy distributions. The ⁸B neutrino elastic scattering rate has the least correlation with other parameters with a global correlation coefficient of 0.18364 which is due to the flat nature of this background across the entire energy range.

Table 6.8 shows the correlation coefficients for the 1.8 to 3.2 MeV analysis with the optimistic systematics systematic scenario. With respect to the 2.2 to 2.8 MeV optimistic systematics analysis, the global correlation coefficients have increased for the light yield, ⁸B neutrino elastic scattering rate, ²¹²Bi²¹²Po rate and thorium-chain rate. The global correlation coefficients of the other systematic uncertainties have decreased. This shows that extending the energy range is having an effect on the correlations between parameters. The systematic uncertainty which is most correlated with the other systematics for this scenario is energy shift with a global correlation coefficient of 0.83275. The external

Table 6.8: The correlation coefficient matrix for the analysis with a 1.8 to 3.2 MeV energy window with 40 keV bins in the optimistic systematics scenario. Global is the overall correlation of the parameter with all other parameters. LY is light yield which is directly related to energy resolution. Scale and shift are the energy scale and shift parameters. $2\nu 2\beta$, ⁸B ν , ²¹²BiPo, U, Th and Ext correspond to the ¹³⁰Te two neutrino double-beta decay, ⁸B neutrino elastic scattering, ²¹²Bi²¹²Po, uranium, thorium-chain and external background rate uncertainties respectively.

	Global	LY	Scale	Shift	$2\nu 2\beta$	$^{8}\mathrm{B} \nu$	²¹² BiPo	U	Th	Ext
LY	0.78472	1.000	0.178	0.450	0.090	-0.274	0.236	-0.124	-0.314	0.036
Scale	0.19752	0.178	1.000	0.043	0.072	-0.069	0.044	-0.066	-0.093	0.005
Shift	0.83275	0.450	0.043	1.000	-0.532	-0.116	0.040	-0.394	0.347	-0.004
$2\nu 2\beta$	0.75011	0.090	0.072	-0.532	1.000	-0.038	0.032	-0.050	-0.519	0.005
$^{8}\mathrm{B} \nu$	0.46418	-0.274	-0.069	-0.116	-0.038	1.000	-0.013	0.100	-0.176	-0.025
²¹² BiPo	0.32854	0.236	0.044	0.040	0.032	-0.013	1.000	-0.018	0.021	-0.020
U	0.57680	-0.124	-0.066	-0.394	-0.050	0.100	-0.018	1.000	-0.242	0.004
Th	0.74138	-0.314	-0.093	0.347	-0.519	-0.176	0.021	-0.242	1.000	-0.017
Ext	0.05832	0.036	0.005	-0.004	0.005	-0.025	-0.020	0.004	-0.017	1.000

Table 6.9: The correlation coefficient matrix for the analysis with a 1.8 to 3.2 MeV energy window with 40 keV bins in the conservative systematics scenario. Global is the overall correlation of the parameter with all other parameters. LY is light yield which is directly related to energy resolution. Scale and shift are the energy scale and shift parameters. $2\nu 2\beta$, ⁸B ν , ²¹²BiPo, U, Th and Ext correspond to the ¹³⁰Te two neutrino double-beta decay, ⁸B neutrino elastic scattering, ²¹²Bi²¹²Po, uranium, thorium-chain and external background rate uncertainties respectively.

	Global	LY	Scale	Shift	$2\nu 2\beta$	$^{8}\mathrm{B} \nu$	²¹² BiPo	U	Th	Ext
LY	0.51929	1.000	0.415	0.007	-0.053	0.009	0.034	0.024	0.204	0.003
Scale	0.85519	0.415	1.000	-0.386	-0.073	-0.008	-0.083	-0.028	0.470	-0.006
Shift	0.60752	0.007	-0.386	1.000	0.014	0.002	0.003	-0.016	0.011	0.004
$2\nu 2\beta$	0.99660	-0.053	-0.073	0.014	1.000	0.012	0.020	-0.990	0.008	0.052
$^{8}\mathrm{B} \nu$	0.18925	0.009	-0.008	0.002	0.012	1.000	-0.093	-0.009	-0.055	-0.083
²¹² BiPo	0.49645	0.034	-0.083	0.003	0.020	-0.093	1.000	-0.010	-0.132	-0.443
U	0.99661	0.024	-0.028	-0.016	-0.990	-0.009	-0.010	1.000	-0.094	-0.051
Th	0.62418	0.204	0.470	0.011	0.008	-0.055	-0.132	-0.094	1.000	-0.001
Ext	0.47302	0.003	-0.006	0.004	0.052	-0.083	-0.443	-0.051	-0.001	1.000

background rate is the least correlated with a global correlation coefficient of 0.05832. The two most correlated systematics are energy shift and two neutrino double-beta decay rate with a correlation of -0.532 which is less strongly correlated than the 2.2 to 2.8 MeV optimistic systematics analysis.

Table 6.9 shows the correlation coefficients for the 1.8 to 3.2 MeV analysis with the conservative systematics scenario. The uranium-chain rate and the two neutrino double-beta decay rate are the most correlated systematic uncertainties with the other parameters. These two systematics have the highest correlation between pairs of systematic uncertainties as well with a correlation coefficient of -0.99. It is likely that these two systematics will be highly correlated due to the similar shape the energy spectra have for energies be-

Systematia	Reduction in half-life sensitivity $(\%)$				
Systematic	Optimistic	Conservative			
Light Yield	3.09	4.67			
Energy Scale	2.21	4.67			
Energy Shift	5.59	5.72			
Two Neutrino Double-Beta Decay Rate	5.52	14.4			
⁸ B Neutrino Elastic Scattering Rate	4.30	3.30			
$^{212}\text{Bi}^{212}\text{Po}$ Rate	3.75	17.7			
Thorium-Chain Rate	2.21	10.9			
Uranium-Chain Rate	2.21	4.67			
External Backgrounds Rate	4.30	21.3			
All Systematics	7.93	33.3			
Sum In Quadrature	11.7	34.7			

Table 6.10: The reductions in half-life for each systematic uncertainty for the optimistic and conservative systematics scenarios using the analysis with the 2.2 to 2.8 energy range and 20 keV bins.

low approximately 2.4 MeV. The ⁸B neutrino elastic scattering rate is the least correlated with the other systematic uncertainties with a global correlation coefficient of 0.18925. All systematic uncertainties have higher global correlation coefficients than the 2.2 to 2.8 MeV conservative systematics analysis apart from energy shift and light yield. This shows that, in the conservative systematics scenario, increasing the energy range is having less impact on reducing the correlations between parameters once all systematic uncertainties are included.

Reductions in half-life

Table 6.10 shows the reductions in half-life for each systematic uncertainty using the 2.2 to 2.8 MeV analysis. In the optimistic systematics scenario the resultant sensitivity extracted at the 90% CL for neutrinoless double-beta decay is 8.83×10^{25} years. This sensitivity drops to 6.40×10^{25} years in the conservative systematics scenario. This means that the optimistic systematics scenario has a 38% increase in sensitivity with respect to the conservative systematics scenario are energy shift and two neutrino double-beta decay rate with reductions of 5.6% and 5.5% respectively. Two neutrino double-beta decay rate is still dominant in the conservative systematics scenario with a reduction of 14% but it is not as dominant as the external backgrounds and ²¹²Bi²¹²Po rates which have reductions in half-life sensitivity of 21% and 18%.

Table 6.11 shows the reductions in half-life for each systematic uncertainty using the 1.8 to 3.2 MeV analysis. The 90% CL sensitivity to neutrinoless double-beta decay with

Systematic	Reduction in half-life sensitivity $(\%)$				
Systematic	Optimistic	Conservative			
Light Yield	1.02	0.			
Energy Scale	0.835	2.81			
Energy Shift	0.	1.40			
Two Neutrino Double-Beta Decay Rate	0.	5.24			
⁸ B Neutrino Elastic Scattering Rate	1.96	2.43			
$^{212}\mathrm{Bi}^{212}\mathrm{Po}\ \mathrm{Rate}$	1.84	13.6			
Thorium-Chain Rate	0.835	2.43			
Uranium-Chain Rate	0.	5.77			
External Backgrounds Rate	2.08	17.3			
All Systematics	2.57	22.4			
Sum In Quadrature	3.74	23.8			

Table 6.11: The reductions in half-life for each systematic uncertainty for the optimistic and conservative systematics scenarios using the analysis with the 1.8 to 3.2 energy range and 40 keV bins

the optimistic systematics scenario is 9.30×10^{25} years and 7.41×10^{25} years with the conservative systematics scenario. This is an improvement of 5.3% and 16% with respect to the 2.2 to 2.8 MeV analysis and the optimistic systematics scenario has an improvement on sensitivity of 26% with respect to the conservative systematics scenario. Apart from the reduction due to the conservative uranium-chain rate systematic uncertainty, all other reductions are smaller than the 2.2 to 2.8 MeV analysis. This shows the benefit of using a wide energy range to constrain systematic uncertainties in the fit. The conservative uranium-chain rate systematic uncertainty may be higher than the 2.2 to 2.8 MeV analysis due to the large correlation between the other systematics in the analysis as the global correlation coefficient is almost one with a value of 0.99661. The global correlation coefficient in the 2.2 to 2.8 MeV analysis is slightly smaller with a value of 0.90388. The dominant systematic uncertainty for both the optimistic and conservative systematics scenarios is the external backgrounds rate with values of 2.1% and 17% respectively. This systematic may be constrained further by fitting as a function of volume as the event rate of external backgrounds increases as the radius from the centre of the detector increases. Therefore there will be more events at higher radial values which will constrain the fit. This is explored further in the following section.

6.4 Multidimensional Limit Setting

The analyses to this point have only been fitting in energy however there is the potential to gain more sensitivity by binning as a function of distance from the detector centre which



Figure 6.26: The half-life sensitivity at 90% CL as a function of range of volume, $(R/R_{AV})^3$ where R is the radius from the centre of the detector and R_{AV} is the radius of the Acrylic Vessel. Several bin sizes where investigated, 0.01 in the black spots, 0.05 in the vertical blue crosses, 0.1 in the diagonal red crosses and 0.2 in the magenta diamonds. The analysis with an energy range of 2.2 to 2.8 MeV was used to encapsulate the entire signal region.

benefits from self shielding of backgrounds originating from outside the liquid scintillator. The external background rate increases as a function of volume from the centre of the detector, see Section 5.2.4, whereas the internal backgrounds and signal events are constant as a function of volume. A fiducial volume radius cut of 3.5 m, which is approximately 20% of the AV volume, was applied thus far to maximise the signal to background counts. The issue with just using a fiducial volume cut is that it puts the external background events in the same energy spectrum as signal events from the detector centre even though it is known that most of these events are at the edge of the volume. If the volume is divided up in the analysis then the signal to background counts will increase at the centre of the detector. By binning the detector and R_{AV} is the radius of the Acrylic Vessel (AV), it should be possible to widen the fiducial volume and gain in sensitivity.



Figure 6.27: The χ^2 space at the 90% CL sensitivity in the parameters of volume element, $(R/R_{AV})^3$, and energy.

6.4.1 Optimisation With Fixed Backgrounds

To optimise the bin size and examine the point at which no more sensitivity is gained the fixed background case was used as a first step. The current optimisation of 1.8 to 3.2 MeV and 40 keV bins is used for the energy parameter.

Figure 6.26 shows the half-life sensitivity at 90% CL achieved for each range in volume from the centre of the detector for different bin sizes. Little sensitivity is gained above 40% of AV volume. There is a slight advantage in using finer binning but the effect is negligible. A bin size of 0.1 AV volume was chosen as this has some flexibility in choosing the maximum range in volume considered whilst also maintaining enough Monte Carlo statistics in each bin. With the previous single volume bin analysis (one 0.2 AV volume bin and hence a 20% AV volume) the sensitivity achieved is 9.60×10^{25} years. With 0.1 AV volume bins and a 40% AV volume a sensitivity of 1.10×10^{26} years is reached. This is an additional improvement of 15% with respect to the analysis presented in the previous section (1.8 to 3.2 MeV with 40 keV bins and one volume bin) and a total improvement of 25% over the single volume and energy bin analysis.

Figure 6.27 shows the χ^2 space with the signal set at the 90% CL sensitivity for the two dimensions that are being fit in the sensitivity study. Here you can see that there is



Figure 6.28: The reduction in half life sensitivity as a function of volume element $(R/R_{AV})^3$, where R is the radius from the centre of the detector and R_{AV} is the radius of the Acrylic Vessel, with a 50% uncertainty on external background rate. Two bin sizes where investigated: 0.05 in the vertical blue crosses and 0.1 in the diagonal red crosses.

little signal contribution above 2.7 MeV and no signal contribution below 2.3 MeV. There is significant signal contribution in volume elements above the current fiducial volume cut at 20% AV volume. As also shown in Figure 6.27, little sensitivity is gained above 40% of the AV volume.

6.4.2 Optimisation With Systematic Uncertainties

As external background counts vary with volume element it make sense to optimise the size of the volume with the external rate systematic uncertainty included. An analysis is used with an energy range 1.8 to 3.2 MeV, 40 keV bins, as optimised in Section 6.2.1, and a systematic uncertainty of 50% on the external backgrounds rate. Figure 6.28 shows the reduction in half-life sensitivity due to the systematic uncertainty in external background rate as the size of the fiducial volume is increased. From this it appears that a maximum bin at 70% of the AV volume is sufficient for reducing the impact of a systematic uncertainty in the external backgrounds rate. With a bin size of 0.1 $(R/R_{AV})^3$ and a 70% total volume the reduction in sensitivity to neutrinoless double-beta decay half-life due to the

uncertainty in the external backgrounds rate is 0.42%. This is an improvement from a 10% reduction in half-life sensitivity at a fiducial volume of 20%.

The estimated 0.42% for the reduction in sensitivity assumes we have the correct model for the external backgrounds. If the way we reconstruct external backgrounds behaves differently in Monte Carlo to data then the above conclusion may not be true. For instance, the bias in energy and position reconstruction may not be simulated correctly, or the attenuation and interaction length of gammas in the scintillator cocktail may not be modelled accurately enough. The proportion of events detected from the different external source locations may also not be correctly modelled. These variables will affect the external background event rate as a function of volume element. The external background rate systematic uncertainty assumes perfect knowledge of these other factors.

Clearly it is important that the model of external backgrounds is developed to be accurate. There are a variety of ways this could be achieved. The external background model could be tuned using pure scintillator data (before the tellurium is loaded in). A gamma emitting source could also be deployed in the external water near the AV to calibrate the Monte Carlo model of external backgrounds. During the tellurium loaded phase, a volume region outside that used in the neutrinoless double-beta decay analysis could also be used to constrain the model of external backgrounds.

6.4.3 Combined Systematic Uncertainties Study

The same methodology is used here as described in section 6.3.3 to extract the reductions in half life for each systematic uncertainty with the correlations of each systematic uncertainty accounted for. The correlation coefficient matrices are also extracted from the MINUIT minimiser. The energy range used in this analysis is 1.8 to 3.2 MeV with 40 keV bins and 0.1 AV volume bins up to 70% AV volume.

Correlations

Table 6.12 shows the correlation coefficients for the multidimensional analysis with the optimistic systematics scenario. As has been seen previously in the energy fit analyses, the uranium-chain and two neutrino double-beta decay rates are highly correlated with the other systematic uncertainties with global correlation coefficients of 0.99171 and 0.99151 respectively. They are also the two systematic uncertainties which are most correlated with each other with a correlation coefficient of -0.975. The energy shift, ²¹²Bi²¹²Po rate and ⁸B neutrino elastic scattering rate all appear to be relatively uncorrelated with the

Table 6.12: The correlation coefficient matrix for the multidimensional analysis in the optimistic systematics scenario. Global is the overall correlation of the parameter with all other parameters. LY is light yield which is directly related to energy resolution. Scale and shift are the energy scale and shift parameters. $2\nu 2\beta$, ⁸B ν , ²¹²BiPo, U, Th and Ext correspond to the ¹³⁰Te two neutrino double-beta decay, ⁸B neutrino elastic scattering, ²¹²Bi²¹²Po, uranium, thorium-chain and external background rates respectively.

	Global	LY	Scale	Shift	$2\nu 2\beta$	$^{8}\mathrm{B} \nu$	²¹² BiPo	U	Th	Ext
LY	0.75297	1.000	0.639	-0.005	-0.116	0.036	0.027	0.075	0.278	0.020
Scale	0.87186	0.639	1.000	-0.010	-0.047	-0.001	-0.012	-0.107	0.530	-0.074
Shift	0.01115	-0.005	-0.010	1.000	0.000	-0.000	-0.000	0.002	-0.005	0.000
$2\nu 2\beta$	0.99151	-0.116	-0.047	0.000	1.000	0.005	-0.002	-0.975	-0.089	-0.034
${}^{8}\mathrm{B} \nu$	0.14096	0.036	-0.001	-0.000	0.005	1.000	-0.049	0.003	-0.098	-0.026
²¹² BiPo	0.09089	0.027	-0.012	-0.000	-0.002	-0.049	1.000	0.006	-0.040	-0.021
U	0.99171	0.075	-0.107	0.002	-0.975	0.003	0.006	1.000	-0.049	0.030
Th	0.64973	0.278	0.530	-0.005	-0.089	-0.098	-0.040	-0.049	1.000	-0.070
Ext	0.21734	0.020	-0.074	0.000	-0.034	-0.026	-0.021	0.030	-0.070	1.000

Table 6.13: The correlation coefficient matrix for the multidimensional analysis in the conservative systematics scenario. Global is the overall correlation of the parameter with all other parameters. LY is light yield which is directly related to energy resolution. Scale and shift are the energy scale and shift parameters. $2\nu 2\beta$, ⁸B ν , ²¹²BiPo, U, Th and Ext correspond to the ¹³⁰Te two neutrino double-beta decay, ⁸B neutrino elastic scattering, ²¹²Bi²¹²Po, uranium, thorium-chain and external background rates respectively.

	Global	LY	Scale	Shift	$2\nu 2\beta$	$^{8}\mathrm{B} \nu$	²¹² BiPo	U	Th	Ext
LY	0.55662	1.000	0.004	-0.025	-0.065	0.024	0.169	0.107	-0.062	0.077
Scale	0.00995	0.004	1.000	-0.000	0.001	0.000	-0.000	-0.001	0.004	-0.000
Shift	0.02477	-0.025	-0.000	1.000	0.002	-0.001	-0.004	-0.003	0.002	-0.002
$2\nu 2\beta$	0.99768	-0.065	0.001	0.002	1.000	0.012	0.007	-0.996	0.052	-0.042
$^{8}\mathrm{B} \nu$	0.21826	0.024	0.000	-0.001	0.012	1.000	-0.176	-0.009	-0.085	-0.013
²¹² BiPo	0.30984	0.169	-0.000	-0.004	0.007	-0.176	1.000	-0.001	-0.129	-0.074
U	0.99771	0.107	-0.001	-0.003	-0.996	-0.009	-0.001	1.000	-0.083	0.035
Th	0.44027	-0.062	0.004	0.002	0.052	-0.085	-0.129	-0.083	1.000	-0.026
Ext	0.21533	0.077	-0.000	-0.002	-0.042	-0.013	-0.074	0.035	-0.026	1.000

other systematics.

Table 6.13 shows the correlation coefficients for the multidimensional analysis with the conservative systematics scenario. Again the uranium-chain and two neutrino double-beta rates are the systematic uncertainties which are most correlated with the other systematics. They have global correlation coefficients of 0.99771 and 0.99768 respectively. They are also the most correlated pair of systematics with a correlation coefficient of -0.996.

Reductions in half-life

Table 6.14 shows the reductions in half-life sensitivity at 90% CL to neutrinoless doublebeta decay for each systematic uncertainty with all correlations accounted for. The benefits in using volume elements to divide up the data can be clearly seen in the improvement on

Systematic	Reduction in half-life sensitivity $(\%)$				
Systematic	Optimistic	Conservative			
Light Yield	0.813	1.40			
Energy Scale	1.14	2.43			
Energy Shift	0.813	1.99			
Two Neutrino Double-Beta Decay Rate	0.813	4.63			
⁸ B Neutrino Elastic Scattering Rate	2.78	3.13			
$^{212}\text{Bi}^{212}\text{Po}$ Rate	2.68	16.4			
Thorium-Chain Rate	1.14	0.0			
Uranium-Chain Rate	1.14	5.40			
External Backgrounds Rate	1.40	1.99			
All Systematics	4.29	16.9			
Sum In Quadrature	4.77	18.6			

Table 6.14: The results presented are for the multidimensional analysis.

the reductions in half-life for the external background rates when comparing to Table 6.11. For the energy-only analysis the reductions observed due to the external background rate are 2.08% and 17.3% for the optimistic and conservative systematics scenarios respectively. This improves to 1.40% and 1.99% for the optimistic and conservative systematics scenarios respectively using the multidimensional analysis technique. This shows that the external background rate is being constrained by volume elements outside the fiducial volume where the rate of those events is much higher.

The systematic uncertainty which is having the greatest effect on sensitivity, particularly in the conservative systematics scenario is the $^{212}Bi^{212}Po$ rate which has a 16.4% reduction in half-life sensitivity in the conservative systematics scenario. This background should therefore be a priority for SNO+ to be reduced and constrained as much as possible.

6.5 Summary of Results

The previously developed analysis which was a simple counting experiment performed on an energy range of 2.42 to 2.64 MeV achieved a neutrinoless double-beta decay halflife sensitivity of 8.81×10^{25} yr. By binning the data set with 20 keV in this energy range the sensitivity improved by 2.7% to 9.05×10^{25} yr. A further 6.1% improvement was made by extending the energy range to 2.32 to 2.68 MeV to achieve a sensitivity of 9.60×10^{25} yr corresponding to nearly a 10% improvement over the previously developed analysis technique.

The analyses to this point had not yet included systematic uncertainties. Four analyses were considered to examine the effects of individual systematics on the sensitivity to neutrinoless double-beta decay: a single bin counting-type analysis between 2.42 and 2.64 MeV, which was used in the previous independent analysis; an energy range of 2.42 and 2.64 MeV with 20 keV bins which shows the advantages of simply binning the data; an energy range of 2.2 to 2.8 MeV with 20 keV bins which encapsulates the entire signal region; and the optimised analysis of 1.8 to 3.2 MeV with 40 keV bins which minimises the effect of systematics. The optimised range was found by investigating the energy range which would minimise the reduction in neutrinoless double-beta decay sensitivity due to the light yield systematic uncertainty. Using these analyses it was found that the ²¹⁴Bi²¹⁴Po and the 2223 keV gamma background rate systematic uncertainties were unimportant. The original single binned counting analysis was found to be very sensitive to systematic uncertainties. As an example, an uncertainty of 0.05 MeV on energy shift led to a 70% reduction in neutrinoless double-beta decay sensitivity.

The 2.2 to 2.8 MeV and 1.8 to 3.2 MeV analyses were used to evaluate the reductions in neutrinoless double-beta decay half-life sensitivities with two different scenarios of systematic uncertainties: one where the systematic uncertainties were set to optimistic values and a second where the systematic uncertainties were conservative. By comparing the result where all systematics were included and the result where one systematic had been removed the correlations between systematics were accounted for. Table 6.15 shows the results of the analysis with correlations accounted for with the wide 1.8 to 3.2 MeV energy window and compares it to the wide and narrow 2.2 to 2.8 MeV results from the studies where systematics were investigated individually. The reductions in half-life generally improve for each systematic when moving from the narrow to the wide energy range. The table also shows that correlations are having some effect on the reductions whereby the reductions generally slightly increase for each systematic when the correlations are accounted for. The neutrinoless double-beta decay half-life sensitivities for the optimistic and conservative systematics scenario where calculated to be 9.30×10^{25} vr and 7.41×10^{25} yr at 90% CL respectively for the 1.8 to 3.2 MeV analysis corresponding to 5.3% and 16% improvement over the 2.2 to 2.8 MeV analysis.

Further sensitivity to neutrinoless double-beta decay was achieved by binning the data set in fractional AV-volume bins which acts as a self-shield from backgrounds originating from outside the liquid scintillator as well as fitting in energy using the optimised analysis (1.8 to 3.2 MeV with 40 keV bins). It was found that by extending the fiducial volume to 40% and using 0.1 AV volume bins the sensitivity improved to 1.10×10^{26} yr at 90% CL. This is an improvement of 25% over the previously developed single bin counting experiment technique. An optimised analysis for including systematic uncertainties was Table 6.15: The reductions in half-life sensitivity for each systematics when the systematics are investigated individually and when correlations are accounted for. The wide 1.8 to 3.2 MeV, 40 keV bins energy fit analysis was used with the optimistic systematics scenario. For comparison, the results where systematics are investigated individually for the narrow 2.2 to 2.8 MeV, 20 keV bins energy fit analysis are included. For the individual systematics studies a coarse binning was used so the error on each reduction is $\pm 0.18\%$.

Systematic	Reduction in half-life sensitivity $(\%)$				
Systematic	Indivi	dual	With Correlations		
	(Narrow)	(Wide)	(Wide)		
Light Yield	3.89	1.08	1.02		
Energy Scale	1.09	0.727	0.835		
Energy Shift	0.730	0.365	0.0		
Two Neutrino Double-Beta Decay Rate	1.45	0.0	0.0		
⁸ B Neutrino Elastic Scattering Rate	1.09	1.09	1.96		
$^{212}\mathrm{Bi}^{212}\mathrm{Po}\ \mathrm{Rate}$	0.730	1.09	1.84		
Thorium-Chain Rate	0.366	0.0	0.835		
Uranium-Chain Rate	0.366	0.0	0.0		
External Backgrounds Rate	1.09	1.09	2.08		

found by incorporating a 50% uncertainty on external backgrounds rate. It was found that a fiducial volume of 70% and 0.1 AV volume bins was necessary to constrain the external background systematic.

The optimistic and conservative systematics scenarios for the combined study was repeated for the optimised multidimensional analysis. The sensitivity for half-life in the multidimensional optimistic systematics scenario is 1.06×10^{26} years and 9.2×10^{25} years at 90% CL for the conservative systematics scenario. This is an improvement of 14% and 24% with respect to the optimistic and conservative systematics scenarios with the 1.8 to 3.2 MeV energy dimension only analysis. In comparison to the previous independent, single binned analysis where systematic uncertainties were neglected an improvement in sensitivity is still observed in both the optimistic, 20% improvement, and conservative, 4.1% improvement, multidimensional analysis which shows the benefit in utilising these two dimensions to extract a limit on half-life.

A final result of 1.06×10^{26} years in the optimistic systematics scenario for ¹³⁰Te is comparable to KamLAND-Zen's result of 1.07×10^{26} years for ¹³⁶Xe [66]. A half-life of 1.06×10^{26} years corresponds to a limit on the effective Majorana neutrino mass, $m_{\beta\beta}$, of 51.9 to 125 meV using a phase space factor of $G = 3.69 \times 10^{-14}$ yr⁻¹ [121] and $g_A = 1.269$; with the range of the values due to the various nuclear matrix element calculation methods [122, 121, 123, 124, 125, 126]. This is an improvement on the limit that KamLAND-Zen set on $m_{\beta\beta}$ of 61 to 165 meV. With improved rejection of external backgrounds and Bi-Po events it may be possible for SNO+ to improve the mass sensitivity further.

176

6.6 Potential Further Studies

There have been a number of systematic uncertainties that have not been considered in the above studies. One such systematic is the uncertainty on the position of the events which would lead to events moving into and out of the fiducial volume or between volume element bins.

There are a number of systematic uncertainties associated with the external backgrounds which have been assumed to be constant in the above studies including the energy and position systematics. Another systematic is the misreconstruction rate in the fiducial volume or volume elements. An additional systematic which has not been considered is the relative contribution of external backgrounds from different source locations to the overall external background rate. The external backgrounds originating from different source locations interact differently in the active volume as discussed in Section 5.2.4. For instance events which originate from the AV are more likely to be reconstructed outside the fiducial volume than events which originate in the external water.

A further systematic uncertainty which has not yet been studied is non-linearities in energy reconstruction. This causes events to move to different energies inside and outside of the expected energy PDF. This would lead to events being reconstructed into and out of the signal region.

There is also a potential to use other variables to extract more sensitivity to neutrinoless double-beta decay from the data. One such parameter could be the isotropy of the hits detected at the PMTs. Events occurring towards the centre of the detector are expected to be isotropic due to the nature of scintillator light emission. External background events which interact multiple times in the scintillator and interact near the AV are expected to be more anisotropic as the interactions near the AV occur closer to a selection of PMTs meaning the probability of generating a hit at those PMTs is higher than other PMTs in the detector. Cuts or fits to isotropy could therefore remove or help identify external background events in the data set and lead to an increase in sensitivity.

Another parameter which may be useful is the time profile of events in the detector. Background events with gammas that interact multiple times have a longer time profile than signal events and potentially have multiple peaks in their time distributions.

Chapter 7

Conclusions

The main aim of the SNO+ experiment is to discover neutrinoless double-beta decay which will tell us that neutrinos are Majorana. If neutrinos are Majorana then this could explain the lightness of neutrino mass through the see-saw mechanism and the process would violate lepton number. This is a requirement of some models to meet the Sakharov conditions and explain the observed matter-antimatter asymmetry in the universe.

In SNO+ neutrinoless double-beta decay will be searched for by loading a scintillator cocktail with tellurium whose isotope ¹³⁰Te decays via double beta decay. The first phase with 0.3% tellurium loading is expected to achieve a neutrinoless double-beta decay half-life sensitivity of about 9×10^{25} years at a 90% Confidence Limit (CL) and it is this phase which is the focus of the thesis. A second phase is planned with 3% loading which is expected to achieve a sensitivity around 7×10^{26} years [67].

The use of scintillator as the active mass has lowered the energy threshold of the detector and requires strict radiopurity control so the ratio of signal to background counts is as high as possible. This thesis has described a new external optical calibration system: the Timing component of the Embedded LED Light Injection Entity (TELLIE). This system is used to calibrate the timing and charge output of the photomultiplier tube (PMT) array.

This thesis has also demonstrated a new analysis which improves the sensitivity of SNO+ to neutrinoless double-beta decay. This was achieved by fitting in two binned dimensions: energy and fractional volume of the Acrylic Vessel (AV). Systematic uncertainties have been included in the studies when previously none were.

7.1 TELLIE

Due to the stringent radiopurity requirements on SNO+ it is important to reduce the use of deployed sources in the detector. To limit the use of these sources an external optical calibration source, TELLIE, has been developed to achieve continuous calibration of the PMT array. TELLIE generates light pulses from LED drivers on the deck above the detector. The light is transmitted down optical fibres that terminate at the PMT Support Structure (PSUP), which hosts the PMT array. The light detected by the PMTs opposite the fibre is used to calibrate the PMTs' timing and charge response. There are 91 injection points across the PSUP allowing calibration of the entire PMT array.

TELLIE has been successfully tested and the majority of fibres have been installed at the SNO+ detector. Prior to installation there were a series of quality assurance tests to ensure that TELLIE met the requirements for PMT calibration. These requirements were a peak wavelength of 500 to 520 nm so that absorption of light by the scintillator is minimised and the PMT quantum efficiency is maximised; a pulse intensity range of 10^3 to 10^5 photons per pulse so single photoelectron and multiple photoelectron hits at PMTs can be studied; and a time profile with a Full Width Half Maximum (FWHM) of less than 6 ns so that the time offsets between PMTs can be measured to a precision of 1 ns or less with a run time of up to 24 hours. The results I obtained showed that: the range of peak wavelengths measured were 500 to 510.1 nm with a mean of (505.5 ± 2.6) nm; the highest minimum output of a driver-LED combination was 510 photons per pulse; the lowest maximum output was 4.5×10^5 photons per pulse; and all 96 LED-driver board combinations met the FWHM requirements for PMT calibration.

There were two major commissioning periods since TELLIE was installed. One in February/March 2014 and one in December 2014. This allowed TELLIE to be tested with the SNO+ detector. There were four goals for TELLIE during these commissioning periods: investigate the stability of TELLIE by monitoring the intensity of light over the course of data taking; calibrate the PMT array; test the SNO+ DAQ system by examining any missed TELLIE triggers; and measure a true single photoelectron spectrum at a PMT. The intensity rise of TELLIE pulses over the course of data taking was found to be smaller than pulse to pulse fluctuations and did not significantly change the probability of multiphoton hits at PMTs. A total of 3224 PMT channels had enough triggered hits to perform PMT calibration. 81.7% of channels passed the PMT gain calibration checks and 87.6% of PMTs passed PMT timing calibration checks. It was found that 0.65% of TELLIE triggers were lost by the SNO+ DAQ system and there has since been significant work by

the SNO+ DAQ group to rectify this. Using the single photoelectron spectra of PMTs it was found that the thresholds could be reduced by an average of 13.6 ± 1.3 counts above pedestal (cap).

A series of data quality checks that I developed, designed to ensure any data taken with TELLIE is of a high enough standard for PMT calibration that, are presented. These include checks on the pulse intensity; efficiency of the trigger; the delay between pulses; the fibre which is injecting light; and the time profile of the events.

7.2 Sensitivity to Neutrinoless Double-Beta Decay

The main aim of SNO+ is to search for neutrinoless double-beta decay. This is achieved in two phases: a first phase with 0.3% tellurium loading; and a second phase with 3% loading. This thesis concentrates on the first phase. With a five year live time, 0.3% loading, a light yield of 200 NHit/MeV and 3.5 m fiducial volume cut, a previously developed analysis technique determined the sensitivity of SNO+ to neutrinoless double-beta decay to be 8.81×10^{25} yr at 90% CL. This analysis was a simple counting experiment performed in an energy range of 2.42 to 2.64 MeV.

This thesis developed an analysis which performs a fit on a binned data set which incorporates the shape information of the energy spectra rather than simply counting in an energy range. An analysis in this thesis determined that by fitting the energy spectra with 11 bins, this sensitivity improved to 9.05×10^{25} yr at 90% CL. Further improvements to the sensitivity of SNO+ to neutrinoless double-beta are now possible as the use of a binned energy spectrum allows the energy range fitted to be expanded. The analysis was optimised with fixed background rates and detector parameters to an energy range of 2.32 to 2.68 MeV and 40 keV bins to achieve a sensitivity of 9.60×10^{25} yr at 90% CL. This is a sensitivity improvement of nearly 10% over the previous counting analysis.

All of the SNO+ sensitivity calculations to this point have neglected systematic uncertainties and in this thesis I have developed a thorough treatment of them. The systematic uncertainties investigated in this thesis are light yield (and hence energy resolution), energy scale, energy shift and various background rates. The energy range used in analyses with systematic uncertainties was optimised using the energy resolution systematic uncertainy and was found to be 1.8 to 3.2 MeV. This demonstrated that a wider energy window minimises the reduction in half-life sensitivity for this systematic.

The reduction in neutrinoless double-beta decay half-life was investigated as a function of the level of uncertainty on a systematic. Four different analyses were used in the study: a single bin counting-type analysis between 2.42 and 2.64 MeV, which replicated the previous independent analysis; an energy range of 2.42 and 2.64 MeV with 20 keV bins which shows the advantages of simply binning the data; an energy range of 2.2 to 2.8 MeV with 20 keV bins which encapsulates the entire signal region; and the optimised analysis of 1.8 to 3.2 MeV with 40 keV bins which minimises the effect of systematics. It was found that the original single binned counting analysis was very sensitive to systematic uncertainties. For example, an uncertainty of 0.05 MeV on energy shift (the value achieved by KamLAND-Zen) led to a 70% reduction in neutrinoless double-beta decay sensitivity. These analyses allowed the systematics which are significant in reducing the sensitivity to neutrinoless double-beta decay to be identified. It was found through these studies that the uncertainties on the ²¹⁴Bi²¹⁴Po and the 2223 keV gamma background rates were unimportant.

The reduction in half-life for each systematic was determined for the 2.2 to 2.8 MeV and 1.8 to 3.2 MeV analyses with the uncertainties on the systematics fixed at either conservative or optimistic values. This was achieved through the comparison in sensitivities between the case where all systematics are included and a case where one systematic is removed. This allows the correlations between systematic uncertainties to be accounted for. It was found that, once the correlations have been included, the reductions in sensitivities for each systematic generally increased although this effect was found to be relatively minor. It was consistently found that the two neutrino double beta decay rate and the uranium chain rate were highly correlated with the other parameters. The resultant sensitivities in the optimistic and conservative systematics scenarios for the 1.8 to 3.2 MeV analyses were found to be 9.30×10^{25} yr and 7.41×10^{25} yr at 90% CL respectively corresponding to reductions in half-life sensitivity of 3% and 22%. The dominant systematic for this analysis in both scenarios was the external background rate with a reduction in half-life sensitivity of 2% and 17% in the optimistic and conservative systematics scenarios respectively.

By binning the data set as a function of fractional AV-volume, corresponding to the amount of self-shielding from external backgrounds, an improvement in sensitivity to neutrinoless double-beta decay was achieved. With fixed background rates and detector parameters a sensitivity of 1.10×10^{26} years is achieved with an energy range of 1.8 to 3.2 MeV and 40 keV bins; 0.1 AV-volume bins and a 40% total AV-volume. This is a 25% improvement in comparison with the previous independent single binned analysis. For comparison, the previous 3.5 m fiducial volume cut used approximately 20% of the AV volume. The

analysis was optimised using the external background rate systematic uncertainty and it was found that 0.1 AV volume bins and 70% of the AV volume was sufficient to substantially reduce the impact of a 50% uncertainty on the external backgrounds rate. Repeating the optimistic and conservative systematics scenarios with all systematics included reduced the impact of external backgrounds to be 1% and 2% respectively. The resultant sensitivities with all systematics included are 1.06×10^{26} yr and 9.17×10^{25} yr at 90% CL for the optimistic and conservative systematics scenarios respectively. This is a 20% and 4% improvement in sensitivity over the previous counting analysis even with systematics included.

It should be possible for SNO+ to reach the optimistic systematics scenario as there are lots of ways SNO+ can measure the prior values and reduce the uncertainties of systematics. For example the pure scintillator phase presents a good opportunity to measure the external background rate as the rate in this phase should be the same as the rate in the loaded tellurium phases. The use of a gamma emitting source deployed in the water outside the AV could also be used as a way to develop the simulation and quantify its accuracy. There are ongoing studies on whether internal backgrounds could be used to constrain the energy systematics of the SNO+ detector. The various background rates presented in this thesis could be constrained through the measurements of the rates outside the 1.8 to 3.2 MeV energy region used in the sensitivity studies. AV-volume regions outside those used in the analysis (> 0.7 AV volume) could also be used to constrain the backgrounds, particularly the internal backgrounds as these should have a constant event rate as a function of AV-volume. The multidimensional analysis in this thesis has shown that, even if the uncertainties are as a large as the conservative systematics scenario, the sensitivity to the half-life of neutrinoless double-beta decay should still reach the target of 9×10^{25} yr.

In the first phase of SNO+, if the optimistic systematics scenario for systematic uncertainties is achieved then the expected half-life sensitivity will be similar to the KamLAND-Zen result of 1.07×10^{26} yr obtained with a different isotope (¹³⁶Xe). Due to the favourable phase space factor and nuclear matrix elements of tellurium, SNO+ should achieve a limit on $m_{\beta\beta}$ of 52 to 125 meV in comparison to the limit KamLAND-Zen set corresponding to $m_{\beta\beta}$ of 61 to 165 meV. Future phases planned with a higher loading of tellurium and potential improvements to the light yield currently undergoing research and development should substantially increase the sensitivity to neutrinoless double-beta decay.
Appendix A

The radioactive decay chains of 238 U and 232 Th



Figure A.1: The decay chain of 238 U. The energy of the Q values of alpha and beta decays are in MeV and gammas in keV. The half life of each isotope is also shown [127].



Figure A.2: The decay chain of 232 Th. The energy of the Q values of alpha and beta decays are in MeV and gammas in keV. The half life of each isotope is also shown [127].

Bibliography

- Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998. Cited on 1
- [2] Q. Ahmad and et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the sudbury neutrino observatory. *Phys. Rev. Lett.*, 89(011301), 2002. Cited on 1, 7
- [3] N.D. Scielzo, S. Caldwell, G. Savard, J.A. Clark, C.M. Deibel, et al. Double-beta decay Q values of Te-130, Te-128, and Te-120. *Phys.Rev.*, C80:025501, 2009. Cited on 2, 27
- [4] W. Pauli. Letter, addressed to participants of the tubingen conference on radioactivity. Available from the CERN Document server: http://cdsweb.cern.ch/record/83282P., Dec 1930. Cited on 6
- [5] J. Chadwick. Possible existence of a neutron. Nature, 129(312), 1932. Cited on 6
- [6] F. Perrin. Possibilite d'emission de particules neutres de masse intrinseque nulle dans les radioactivites beta. *Comptes Rendus*, 197:1625, Dec 1933. Cited on 6
- [7] E. Fermi. Towards the theory of β -rays. Z. Phys, 88(161), 1934. Cited on 6
- [8] C. Cowan, F.Reines, F. Harrison, H. Kruse, and A. McGuire. Detection of the free neutrino: A confirmation. *Science*, 124, 1956. Cited on 6
- M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos. *Phys. Rev.*, 109(3):1015–1017, Feb 1958. Cited on 6
- [10] G. Danby, J. Gaillard, K.Goulianos, L. Lederman, N. Mistry, and M.Schwartz. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9(1), 1962. Cited on 7

- [11] R. Davis, D. Harmer, and K. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, 1968. Cited on 7
- [12] M. Perl and et al. Evidence for anomalous lepton production in electron positron annihilation. *Phys. Rev. Lett.*, 35(22):1489–1492, Dec 1975. Cited on 7
- [13] R. Bionta and et al. Contained neutrino interactions in an underground water detector. Phys. Rev. D, 38:768–775, 1988. Cited on 7
- [14] K. Hirata and et al. Experimental study of the atmospheric neutrino flux. *Phys. Lett. B*, 205(416), 1988. Cited on 7
- [15] LEP Collaboration. Precision electroweak measurements on the z resonance. Phys. Rept., 427:257–454, 2006. Cited on 7
- [16] Y. Fukuda and et al. Measurement of a small atmospheric nu/mu / nu/e ratio.
 Phys. Lett. B, 433:9–18, 1998. Cited on 7
- [17] Y. Fukuda and et al. Study of the atmospheric neutrino flux in the multi-gev energy range. *Phys. Lett. B*, 436:33–41, 1998. Cited on 7
- [18] K. Kodama and et al. Observation of tau neutrino interactions. Phys. Lett. B, 504:218–224, 2001. Cited on 7
- [19] Y. Abe and et al. Indication of reactor electron-antineutrino disappearance in the double chooz experiment. *Phys. Rev. Lett.*, 108(131801), March 2012. Cited on 8
- [20] F. An and et al. Observation of electron-antineutrino disappearance at daya bay. *Phys. Rev. Lett.*, 108, April 2012. Cited on 8
- [21] J. Ahn and et al. Observation of reactor electron antineutrinos disappearance in the reno experiment. *Phys. Rev. Lett.*, 108(191802), May 2012. Cited on 8
- [22] Y. Abe and et al. Reactor electron antineutrino disappearance in the double chooz experiment. ArXiv e-prints, (1207.6632), August 2012. Cited on 8
- [23] K. Abe et al. Indication of Electron Neutrino Appearance from an Acceleratorproduced Off-axis Muon Neutrino Beam. *Phys. Rev. Lett.*, 107:041801, 2011. Cited on 8

- [24] P. Adamson and et al. Measurement of neutrino and antineutrino oscillations using beam and atmospheric data in minos. *Phys. Rev. Lett.*, 110(251801), June 2013. Cited on 8
- [25] C. L. Bennett et al. Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. Astrophys. J. Suppl., 208:20, 2013. Cited on 11, 14
- [26] R. N. Cahn, D. A. Dwyer, S. J. Freedman, W. C. Haxton, R. W. Kadel, Yu. G. Kolomensky, K. B. Luk, P. McDonald, G. D. Orebi Gann, and A. W. P. Poon. White Paper: Measuring the Neutrino Mass Hierarchy. In *Proceedings, Community Summer Study 2013: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013, 2013.* Cited on xvii, 12
- [27] F. Capozzi, E. Lisi, A. Marrone, D. Montanino, and A. Palazzo. Neutrino masses and mixings: Status of known and unknown 3ν parameters. *Nucl. Phys.*, B908:218–234, 2016. Cited on xi, 12, 13
- [28] Ch. Kraus et al. Final results from phase II of the Mainz neutrino mass search in tritium beta decay. Eur. Phys. J., C40:447–468, 2005. Cited on 12
- [29] V. N. Aseev, A. I. Belesev, A. I. Berlev, E. V. Geraskin, A. A. Golubev, N. A. Likhovid, V. M. Lobashev, A. A. Nozik, V. S. Pantuev, V. I. Parfenov, A. K. Skasyrskaya, F. V. Tkachov, and S. V. Zadorozhny. Upper limit on the electron antineutrino mass from the troitsk experiment. *Phys. Rev. D*, 84:112003, Dec 2011. Cited on 12
- [30] Otokar Dragoun and Drahoslav Vénos. Constraints on the Active and Sterile Neutrino Masses from Beta-Ray Spectra: Past, Present and Future. 2015. Cited on xvii, 13, 14
- [31] P. A. R. Ade et al. Planck 2013 results. XVI. Cosmological parameters. Astron. Astrophys., 571:A16, 2014. Cited on 14
- [32] J. Dunkley et al. The Atacama Cosmology Telescope: likelihood for small-scale CMB data. JCAP, 1307:025, 2013. Cited on 14

- [33] Frank F. Deppisch, Martin Hirsch, and Heinrich Pas. Neutrinoless Double Beta Decay and Physics Beyond the Standard Model. J. Phys., G39:124007, 2012. Cited on 17
- [34] Werner Rodejohann. Neutrinoless Double Beta Decay in Particle Physics. Nucl. Phys. Proc. Suppl., 229-232:113–117, 2012. Cited on 18
- [35] S. M. Bilenky and C. Giunti. Neutrinoless Double-Beta Decay: a Probe of Physics Beyond the Standard Model. Int. J. Mod. Phys., A30(04n05):1530001, 2015. Cited on xvii, xviii, 18, 19, 21, 22
- [36] Tomas R. Rodriguez and G. Martinez-Pinedo. Energy density functional study of nuclear matrix elements for neutrinoless ββ decay. *Phys. Rev. Lett.*, 105:252503, 2010. Cited on 18
- [37] Michael Bender, Paul-Henri Heenen, and Paul-Gerhard Reinhard. Self-consistent mean-field models for nuclear structure. *Rev. Mod. Phys.*, 75:121–180, 2003. Cited on 19
- [38] F. Iachello and J. Barea. Fundamental processes in the interacting boson model: 0 nu beta beta decay. AIP Conf. Proc., 1355:7–14, 2011. Cited on 19
- [39] J. Menendez, A. Poves, E. Caurier, and F. Nowacki. Disassembling the Nuclear Matrix Elements of the Neutrinoless beta beta Decay. Nucl. Phys., A818:139–151, 2009. Cited on 19
- [40] P. K. Rath, R. Chandra, K. Chaturvedi, P. K. Raina, and J. G. Hirsch. Uncertainties in nuclear transition matrix elements for neutrinoless ββ decay within the PHFB model. *Phys. Rev.*, C82:064310, 2010. Cited on 19
- [41] F. Simkovic, G. Pantis, J. D. Vergados, and Amand Faessler. Additional nucleon current contributions to neutrinoless double beta decay. *Phys. Rev.*, C60:055502, 1999. Cited on 19
- [42] L. Winslow and R. Simpson. Characterizing Quantum-Dot-Doped Liquid Scintillator for Applications to Neutrino Detectors. JINST, 7:P07010, 2012. Cited on xvii, 19
- [43] A. Gando et al. Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen. 2016. Cited on xvii, 21

- [44] R. Arnold et al. Measurement of the double-beta decay half-life and search for the neutrinoless double-beta decay of ⁴⁸Ca with the NEMO-3 detector. *Phys. Rev.*, D93(11):112008, 2016. Cited on 23
- [45] R. Arnold et al. Study of 2b-decay of Mo-100 and Se-82 using the NEMO3 detector. JETP Lett., 80:377–381, 2004. [Pisma Zh. Eksp. Teor. Fiz.80,429(2004)]. Cited on 23
- [46] J. Argyriades et al. Measurement of the two neutrino double beta decay half-life of Zr-96 with the NEMO-3 detector. Nucl. Phys., A847:168–179, 2010. Cited on 23
- [47] R. Arnold, C. Augier, J. D. Baker, A. S. Barabash, A. Basharina-Freshville, S. Blondel, S. Blot, M. Bongrand, V. Brudanin, J. Busto, A. J. Caffrey, S. Calvez, C. Cerna, J. P. Cesar, A. Chapon, E. Chauveau, D. Duchesneau, D. Durand, V. Egorov, G. Eurin, J. J. Evans, L. Fajt, D. Filosofov, R. Flack, X. Garrido, H. Gómez, B. Guillon, P. Guzowski, R. Hodák, A. Huber, P. Hubert, C. Hugon, S. Jullian, A. Klimenko, O. Kochetov, S. I. Konovalov, V. Kovalenko, D. Lalanne, K. Lang, Y. Lemière, T. Le Noblet, Z. Liptak, P. Loaiza, G. Lutter, F. Mamedov, C. Marquet, F. Mauger, B. Morgan, J. Mott, I. Nemchenok, M. Nomachi, F. Nova, F. Nowacki, H. Ohsumi, R. B. Pahlka, F. Perrot, F. Piquemal, P. Povinec, P. Přidal, Y. A. Ramachers, A. Remoto, J. L. Reyss, B. Richards, C. L. Riddle, E. Rukhadze, R. Saakyan, X. Sarazin, Yu. Shitov, L. Simard, F. Šimkovic, A. Smetana, K. Smolek, A. Smolnikov, S. Söldner-Rembold, B. Soulé, I. Štekl, J. Suhonen, C. S. Sutton, G. Szklarz, J. Thomas, V. Timkin, S. Torre, Vl. I. Tretyak, V. I. Tretyak, V. I. Umatov, I. Vanushin, C. Vilela, V. Vorobel, D. Waters, and A. Žukauskas. Results of the search for neutrinoless double- β decay in ¹⁰⁰Mo with the nemo-3 experiment. Phys. Rev. D, 92:072011, Oct 2015. Cited on 23
- [48] Alexander S. Barabash. NEMO 3 double beta decay experiment: Latest results. J. Phys. Conf. Ser., 173:012008, 2009. Cited on 23
- [49] Roger Arnold et al. Measurement of the Double Beta Decay Half-life of ¹³⁰Te with the NEMO-3 Detector. Phys. Rev. Lett., 107:062504, 2011. Cited on 23
- [50] R. Arnold et al. Measurement of the $2\nu\beta\beta$ decay half-life of ¹⁵⁰Nd and a search

for $0\nu\beta\beta$ decay processes with the full exposure from the NEMO-3 detector. 2016. Cited on 23

- [51] R. Arnold et al. Results of the search for neutrinoless double- β decay in ¹⁰⁰Mo with the NEMO-3 experiment. *Phys. Rev.*, D92(7):072011, 2015. Cited on 23
- [52] S. Blot. Investigating ββ decay with the NEMO-3 and SuperNEMO experiments. J. Phys. Conf. Ser., 718(6):062006, 2016. Cited on 23
- [53] Paolo Bernardini, Gianluigi Fogli, Eligio Lisi, and Carlo Bucci. Proceedings of the neutrino oscillation workshop (now 2010) final results of cuoricino and status of cuore. Nuclear Physics B - Proceedings Supplements, 217(1):41 – 43, 2011. Cited on 23
- [54] C. Arnaboldi et al. Results from a search for the 0 neutrino beta beta-decay of Te-130. Phys. Rev., C78:035502, 2008. Cited on 23
- [55] K. Alfonso et al. Search for Neutrinoless Double-Beta Decay of ¹³⁰Te with CUORE0. Phys. Rev. Lett., 115(10):102502, 2015. Cited on 24
- [56] K. H. Ackermann et al. The GERDA experiment for the search of $0\nu\beta\beta$ decay in ⁷⁶Ge. *Eur. Phys. J.*, C73(3):2330, 2013. Cited on 24
- [57] W. Xu et al. The MAJORANA DEMONSTRATOR: A Search for Neutrinoless Double-beta Decay of ⁷⁶Ge. J. Phys. Conf. Ser., 606(1):012004, 2015. Cited on 24
- [58] M. Agostini et al. Results on Neutrinoless Double-β Decay of ⁷⁶Ge from Phase I of the GERDA Experiment. *Phys. Rev. Lett.*, 111(12):122503, 2013. Cited on 24
- [59] Valerio D'Andrea. Status Report of the GERDA Phase II Startup. 2016. Cited on 24
- [60] N. Abgrall et al. The Majorana Demonstrator Neutrinoless Double-Beta Decay Experiment. Adv. High Energy Phys., 2014:365432, 2014. Cited on 24
- [61] M. Auger et al. Search for Neutrinoless Double-Beta Decay in ¹³⁶Xe with EXO-200.
 Phys. Rev. Lett., 109:032505, 2012. Cited on 25
- [62] J. B. Albert et al. Search for Majorana neutrinos with the first two years of EXO-200 data. *Nature*, 510:229–234, 2014. Cited on 25

- [63] Reyco Henning. Current status of neutrinoless double-beta decay searches. Reviews in Physics, 1:29 – 35, 2016. Cited on 25
- [64] J. Martín-Albo et al. Sensitivity of NEXT-100 to neutrinoless double beta decay. JHEP, 05:159, 2016. Cited on 25
- [65] A. Gando et al. Measurement of the double- β decay half-life of ¹³⁶Xe with the KamLAND-Zen experiment. *Phys. Rev.*, C85:045504, 2012. Cited on 25
- [66] A. Gando et al. Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen. 2016. Cited on 25, 175
- [67] S. Andringa et al. Current Status and Future Prospects of the SNO+ Experiment. Adv. High Energy Phys., 2016:6194250, 2015. Cited on xviii, 26, 177
- [68] B. Aharmim et al. Measurement of the Cosmic Ray and Neutrino-Induced Muon Flux at the Sudbury Neutrino Observatory. *Phys.Rev.*, D80:012001, 2009. Cited on 27, 46
- [69] Philip G Jones. Background Rejection for the Neutrinoless Double Beta Decay Experiment SNO + Background Rejection for the Neutrinoless Double Beta Decay Experiment SNO+. PhD thesis, University of Oxford, 2011. Cited on xviii, 28, 29
- [70] J. Boger et al. The Sudbury neutrino observatory. Nucl.Instrum.Meth., A449:172–207, 2000. Cited on xix, 28, 29, 42, 44
- [71] L. Sibley. SNO+: Physics program and status update. AIP Conf.Proc., 1604:449–455, 2014. Cited on 29, 38
- [72] Ian T Coulter. Modelling and reconstruction of events in SNO + related to future searches for lepton and baryon number violation related to future searches for lepton and baryon number violation. PhD thesis, University of Oxford, 2013. Cited on xix, 29, 37, 48
- [73] Reliance Industry Limited. Linear Alkyl Benzene, Material Safety Data Sheet, 2011. Cited on 29, 38
- [74] D.F. Cowen, T. Ekenberg, J.R. Klein, F.M. Newcomer, R. Van Berg, et al. The

Sudbury Neutrino Observatory electronics chain. *IEEE Trans.Nucl.Sci.*, 42:925–932, 1995. Cited on 31

- [75] Robert Knapik, Richie Bonventre, and Tim Shokair. SNO+ Electronic Upgrades. Nucl.Phys.Proc.Suppl., 229-232:523, 2012. Cited on 31
- [76] G. D. Obrei Gann. TUBii Proposal, SNO+-doc-2167-v1, 2011. Cited on 32
- [77] G. D. Obrei Gann. ECA Calibration / Uncalibration in SNO+, SNO+-doc-1236-v2, 2011. Cited on xix, 34
- [78] B.E. Berger et al. The KamLAND Full-Volume Calibration System. JINST, 4:P04017, 2009. Cited on xix, 35
- [79] S. J. M. Peeters. Calibration With Radioactive Sources, SNO+-doc-2441-v2, 2014. Cited on 38
- [80] S. Biller. Progress towards a phase ii hybrid loading. SNO+ DocDB, (3798v1), Aug 2016. Cited on 38
- [81] M. Mingfang. Te-LS developments, SNO+-doc-2727-v1, 2014. Cited on 38
- [82] S. Grullon. Light Yield and Scintillation Timing Profiles of Te-loaded Liquid Scintillator for SNO+ (Neutrino 2014 Poster), SNO+-doc-2519-v2, 2014. Cited on 38
- [83] Steven Biller. SNO+ with Tellurium. 2014. Cited on 39
- [84] Boris Pritychenko. Systematics of Evaluated Half-Lives of Double-Beta Decay. Nucl.Data Sheets, 120:101–104, 2014. Cited on 39
- [85] M. Mingfang. SNO+ Te-LS updates, SNO+-doc-2426-v3, 2014. Cited on 40
- [86] P. G. Jones. The SNO+ Experiment, DBD14, 2014. Cited on xix, 41
- [87] C. Arpesella et al. Direct Measurement of the Be-7 Solar Neutrino Flux with 192 Days of Borexino Data. *Phys.Rev.Lett.*, 101:091302, 2008. Cited on 41
- [88] G. Alimonti et al. The liquid handling systems for the Borexino solar neutrino detector. Nucl.Instrum.Meth., A609:58–78, 2009. Cited on 41
- [89] S. Grullon and V. Lozza. A Summary of the Impact of EDTA During the Water Phase of SNO+, SNO+-doc-2780-v4, 2014. Cited on 46

- [90] M. Chen, V. Lozza, and H. O'Keeffe. Expected radioactive backgrounds in SNO+, SNO+-doc-507-v24, 2014. Cited on xi, 45
- [91] T. Hagner, R. von Hentig, B. Heisinger, L. Oberauer, Stefan. Schonert, et al. Muon induced production of radioactive isotopes in scintillation detectors. *Astropart.Phys.*, 14:33–47, 2000. Cited on xi, 46
- [92] P. G. Jones. Rat release 5.0.0 notes. SNO+ DocDB, (2798v1), Oct 2014. Cited on 47
- [93] T. Bolton. The Braidwood reactor antineutrino experiment. Nucl. Phys. Proc. Suppl., 149:166–169, 2005. [,166(2005)]. Cited on 47
- [94] S. Agostinelli et al. GEANT4: A Simulation toolkit. Nucl. Instrum. Meth., A506:250–303, 2003. Cited on 48
- [95] R. Alves et al. The calibration system for the photomultiplier array of the SNO+ experiment. JINST, 10(03):P03002, 2015. Cited on xii, xxi, xxii, xxiii, 50, 59, 61, 67, 70, 71, 73
- [96] S. Peeters. Tellie pulse fitting notes. SNO+ DocDB, (2880), 2014. Cited on 52
- [97] F. Descamps. Tellie hit probabilites revisited. SNO+ DocDB, (1703v2), Feb 2015.
 Cited on xii, 54
- [98] G. Lefeuvre. Sussex optical scan setup. SNO+ DocDB, (2882), 2012. Cited on xxi,
 62, 63, 64
- [99] Hamamatsu H10721P-110 PMT . Cited on xxiv, 63, 80
- [100] James R. Sinclair. PMT Calibration System of SNO+. PhD thesis, University of Sussex, 2015. Cited on xxi, 65, 86, 87
- [101] Eska polyethylene jacketed optical fiber cord: Sh4002, 2016. Cited on 67
- [102] Si pin photodiodes: S5971, s5972 and s5973 series., 2016. Cited on 71
- [103] Metal package pmt photosenso modules h10720/h10721, May 2015. Cited on xxiv, 79, 80

- [104] Bryce A. Moffat. The Optical Calibration of the SUdbury Neutrino Observatory.PhD thesis, Queen's University, 2001. Cited on 87
- [105] F. Descamps. Pca calibration. SNO+ DocDB, (1987v5), May 2016. Cited on 91
- [106] Steve Baker and Robert D. Cousins. Clarification of the Use of Chi Square and Likelihood Functions in Fits to Histograms. Nucl. Instrum. Meth., 221:437–442, 1984. Cited on 111
- [107] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J.*, C71:1554, 2011. [Erratum: Eur. Phys. J.C73,2501(2013)]. Cited on 112
- [108] K. A. Olive et al. Review of Particle Physics. Chin. Phys., C38:090001, 2014. Cited on xiii, 112
- [109] G. Cowan. Statistical techniques for incorporating systematic/theory uncertainties. Theory Experiment Interplay at the LHC Workshop, 2010. Cited on 113
- [110] V. Lozza M. Chen and H. O'Keeffe. Expected radioactive backgrounds in sno+. SNO+ DocDB, (507-v24), 2014. Cited on xxix, 124, 130
- [111] F. James and M. Roos. Minuit a system for function minimization and analysis of the parameter errors and correlations. *Computer Physics Communications*, 10(6):343 – 367, 1975. Cited on 125, 126, 163
- [112] A. Mastbaum. Double-beta phase plots. SNO+ DocDB, (2593-v8), 2014. Cited on xiii, 130
- [113] A. Mastbaum. Double-beta white paper sensitivity plots. SNO+ DocDB, (2593-v10),
 2014. Cited on xiii, 130
- [114] A. Mastbaum. Double-beta white paper sensitivity plots. SNO+ DocDB, (2593v10), Nov 2015. Cited on 129, 131
- [115] P. G. Jones. Rat release 4.5.0 notes. SNO+ DocDB, (2257v1), Jan 2014. Cited on 130
- [116] Releases of rat, 2016. Cited on 132

- [117] A. Gando, Y. Gando, H. Hanakago, H. Ikeda, K. Inoue, K. Ishidoshiro, R. Kato, M. Koga, S. Matsuda, T. Mitsui, D. Motoki, T. Nakada, K. Nakamura, A. Obata, A. Oki, Y. Ono, M. Otani, I. Shimizu, J. Shirai, A. Suzuki, Y. Takemoto, K. Tamae, K. Ueshima, H. Watanabe, B. D. Xu, S. Yamada, H. Yoshida, A. Kozlov, S. Yoshida, T. I. Banks, S. J. Freedman, B. K. Fujikawa, K. Han, T. O'Donnell, B. E. Berger, Y. Efremenko, H. J. Karwowski, D. M. Markoff, W. Tornow, J. A. Detwiler, S. Enomoto, and M. P. Decowski. Limit on neutrinoless ββ decay of ¹³⁶Xe from the first phase of kamland-zen and comparison with the positive claim in ⁷⁶Ge. *Phys. Rev. Lett.*, 110:062502, Feb 2013. Cited on 139, 142
- [118] G.L. Fogli, E. Lisi, A. Palazzo, and A.M. Rotunno. Kamland neutrino spectra in energy and time: Indications for reactor power variations and constraints on the georeactor. *Physics Letters B*, 623(1–2):80 – 92, 2005. Cited on 145
- [119] Roger Arnold et al. Measurement of the Double Beta Decay Half-life of ^{130}Te with the NEMO-3 Detector. *Phys. Rev. Lett.*, 107:062504, 2011. Cited on 145
- [120] B. Aharmin, S. N. Ahmed, A. E. Anthony, N. Barros, E. W. Beier, A. Bellerive, B. Beltran, M. Bergevin, S. D. Biller, K. Boudjemline, M. G. Boulay, B. Cai, Y. D. Chan, D. Chauhan, M. Chen, B. T. Cleveland, G. A. Cox, X. Dai, H. Deng, J. A. Detwiler, M. DiMarco, P. J. Doe, G. Doucas, P.-L. Drouin, F. A. Duncan, M. Dunford, E. D. Earle, S. R. Elliott, H. C. Evans, G. T. Ewan, J. Farine, H. Fergani, F. Fleurot, R. J. Ford, J. A. Formaggio, N. Gagnon, J. TM. Goon, K. Graham, E. Guillian, S. Habib, R. L. Hahn, A. L. Hallin, E. D. Hallman, P. J. Harvey, R. Hazama, W. J. Heintzelman, J. Heise, R. L. Helmer, A. Hime, C. Howard, M. Huang, P. Jagam, B. Jamieson, N. A. Jelley, M. Jerkins, K. J. Keeter, J. R. Klein, L. L. Kormos, M. Kos, C. Kraus, C. B. Krauss, A Kruger, T. Kutter, C. C. M. Kyba, R. Lange, J. Law, I. T. Lawson, K. T. Lesko, J. R. Leslie, J. C. Loach, R. MacLellan, S. Majerus, H. B. Mak, J. Maneira, R. Martin, N. McCauley, A. B. McDonald, S. R. McGee, M. L. Miller, B. Monreal, J. Monroe, B. G. Nickel, A. J. Noble, H. M. O'Keeffe, N. S. Oblath, R. W. Ollerhead, G. D. Orebi Gann, S. M. Oser, R. A. Ott, S. J. M. Peeters, A. W. P. Poon, G. Prior, S. D. Reitzner, K. Rielage, B. C. Robertson, R. G. H. Robertson, R. C. Rosten, M. H. Schwendener, J. A. Secrest, S. R. Seibert, O. Simard, J. J. Simpson, P. Skensved, T. J. Sonley, L. C. Stonehill,

G. Tešić, N. Tolich, T. Tsui, R. Van Berg, B. A. VanDevender, C. J. Virtue, H. Wan Chan Tseung, D. L. Wark, P. J. S. Watson, J. Wendland, N. West, J. F. Wilkerson, J. R. Wilson, J. M. Wouters, A. Wright, M. Yeh, F. Zhang, and K. Zuber. Combined analysis of all three phases of solar neutrino data from the sudbury neutrino observatory. *Phys. Rev. C*, 88:025501, Aug 2013. Cited on 149

- [121] J. Kotila and F. Iachello. Phase-space factors for double-β decay. Phys. Rev. C, 85:034316, Mar 2012. Cited on 175
- [122] J. Barea, J. Kotila, and F. Iachello. Nuclear matrix elements for double-β decay. Phys. Rev. C, 87:014315, Jan 2013. Cited on 175
- [123] Fedor Šimkovic, Vadim Rodin, Amand Faessler, and Petr Vogel. $0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements, quasiparticle random-phase approximation, and isospin symmetry restoration. *Phys. Rev. C*, 87:045501, Apr 2013. Cited on 175
- [124] J. Menendez, A. Poves, E. Caurier, and F. Nowacki. Disassembling the Nuclear Matrix Elements of the Neutrinoless beta beta Decay. Nucl. Phys., A818:139–151, 2009. Cited on 175
- [125] Juhani Hyvärinen and Jouni Suhonen. Nuclear matrix elements for $0\nu\beta\beta$ decays with light or heavy majorana-neutrino exchange. *Phys. Rev. C*, 91:024613, Feb 2015. Cited on 175
- [126] Tomás R. Rodríguez and Gabriel Martínez-Pinedo. Energy density functional study of nuclear matrix elements for neutrinoless ββ decay. Phys. Rev. Lett., 105:252503, Dec 2010. Cited on 175
- [127] Helen Mary O'Keeffe. Low Energy Background in the NCD Phase of the Sudbury Neutrino Observatory. PhD thesis, University of Oxford, 2008. Cited on xxxv, 182, 183