

On the Measurement of Optical Scattering and Studies of Background Rejection in the SNO+ Detector

Krishanu Majumdar

Oriel College, Oxford

Thesis submitted in partial fulfilment of the requirements for the
degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2015

On the Measurement of Optical Scattering and Studies of Background Rejection in the SNO+ Detector

Krishanu Majumdar
Oriel College, Oxford

Thesis submitted in partial fulfilment of the requirements for
the degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2015

Abstract

SNO+ is a liquid scintillator experiment designed to study a wide range of neutrino-related physics goals, such as solar neutrinos, reactor and geo-neutrinos, and neutrinoless double beta decay ($0\nu\beta\beta$). The success of the experiment depends in part on the ability to accurately characterise the detector's optical properties, and also develop effective methods to suppress contributions from unwanted backgrounds.

This thesis presents one of the central calibration systems which will be used to make measurements of the detector's optical scattering properties. The hardware and its integration into the wider detector will be discussed. A simplified analysis of simulated water data is also presented, in order to show that the scattering properties can be accurately measured. The reconstruction of photons' scattering position, length and angles can be achieved to a high degree of accuracy, with a small reconstruction bias and resolution in comparison to the scale of the detector itself.

This thesis also presents a study into the rejection of two important radioactive backgrounds that will be encountered during the scintillator phases of SNO+: the ($^{212}\text{Bi} + ^{212}\text{Po}$) and ($^{214}\text{Bi} + ^{214}\text{Po}$) $\beta - \alpha$ chains, where both components of each chain occur within a single detector trigger window, creating what is known as a "BiPo pileup" event. Two methods have been considered: one using a comparison between the cumulative time residual distributions of BiPo pileup and $0\nu\beta\beta$ signal events, and the other using a log-likelihood difference method. Both methods perform extremely well, with the first being capable of rejecting $> 95\%$ of BiPo pileup events for a 1% loss of signal, and the second rejecting $> 97\%$ of backgrounds for the same signal loss.

Acknowledgements

My first thanks must go to my parents - although it has probably been quite a few years since they fully understood what exactly it is that I do, they've never stopped supporting and encouraging me, and I wouldn't be where I am today without them.

Of course, much of my work over the past four years has been helped in no small part by my colleagues in the Oxford SNO+ group.

My thanks go to my supervisor Steve Biller, who always has a few new ideas on hand to talk about, and a relentlessly infectious enthusiasm that then makes me want to try them out. Looking back, it's amazing how much I was able to pick up just from having occasional, unfocussed discussions about the experiment, the software and physics in general, and this has certainly helped me to become a more capable physicist.

Thank you also to Armin Reichold, who I've come to think of as my unofficial co-supervisor as we've worked together on the SMELLIE hardware, software and analysis. It has been an incredibly well-rounded and rewarding experience, and I hope that all of the work we've put in pays off after I'm gone.

The most illuminating discussions in the past four years have arguably come from my fellow SNO+ graduate students and postdocs - both past (Ian, Phil and Ken) and present (Luca, Jeff, Jack, Laura and yes - even Chris too), as well as the various Masters and Summer students who have come and gone during my time here. I think being one of the smaller experiments in the department (and having only a single student office) has made us a closer-knit group of people, and so thank you to each of them for giving me someone to bounce my ideas off - and subsequently correcting the ones that needed it - and for the impromptu discussions (and arguments!) about everything and anything we could think of - whether it was SNO+, or the merits of computer games, or the latest news headlines.

Even the people who I didn't interact with on a daily basis have also had a major impact on me in my time here.

Thank you to the support staff in the Department of Physics - in particular, to John Saunders and Rik Elliott, for keeping up with the ever-changing (and occasionally improbable!) demands of our electronics, and also to the department secretaries Kim Proudfoot and Sue Geddes for keeping things running smoothly behind the

scenes.

I'd also like give my thanks to Jeanne Wilson and Stefanie Langrock at Queen Mary University of London, for their invaluable contributions to many aspects of the SMELLIE analysis, and my best wishes to the wider SNO+ collaboration. Over the past four years there have been some ups and downs, both major and minor, in the progress of the experiment, and I admit that there have been some times when I really didn't like the situations that we found ourselves in. However, I feel that we are moving in the right direction ... just in time for me to leave ... and I look forward to seeing some interesting results in the scientific literature in the years to come.

Finally, and I think most importantly, I would like to express my gratitude to all of my non-SNO+ friends in the Department of Physics. In particular, I thank Neven for always being up for a conversation over tea, Nazim for his endless supply of entertaining anecdotes and some fantastic discussions about computer gaming, board gaming and Canada, and Kate for being the only other person I've ever met who wields the same brand of optimistic pessimism as I do.

Without them and many others not named here, these past four years would have been much less interesting and much less enjoyable, and in honour of that I have included some of their more insightful (or at least funny) thoughts on the various topics presented in this thesis.

Contents

1	The SNO+ Experiment	1
1.1	The Detector	1
1.1.1	Overview	1
1.1.2	Electronics and Detector Readout	4
1.1.3	Simulation Software	6
1.2	Scintillation and Scintillator	7
1.2.1	Experimental Phases	7
1.2.2	Scintillation	8
1.2.3	The SNO+ Scintillator: LABPPO	11
1.2.4	Effects of Scintillator Properties	14
2	Neutrinos	18
2.1	A Brief History of Neutrino Physics	18
2.1.1	The Need for a New Particle	18
2.1.2	Discovery of ν_e	20
2.1.3	Discovery of ν_μ	21
2.1.4	Subsequent Experimental Results	22
2.1.5	Discovery of ν_τ	24
2.2	Neutrinos in the SM	24
2.2.1	Fundamental Properties	24
2.2.2	Sources of Neutrinos	26
2.3	Beyond the SM: Neutrino Flavour Oscillations	30
2.3.1	Introduction	30
2.3.2	Evidence for ν Flavour Oscillations	31
2.3.3	Mathematical Framework and Implications	36
2.3.4	Current Measurements	38
2.3.5	Mass Hierarchies	39
2.4	Beyond the SM: Massive Neutrinos	39
2.4.1	Introduction and Current Measurements	39
2.4.2	Dirac Neutrinos	41
2.4.3	Majorana Neutrinos	42
2.5	Double β Decay	44
2.5.1	Overview	44
2.5.2	Neutrinoless Double β Decay	46
2.5.3	$0\nu\beta\beta$ at SNO+	50

3	Optical Scattering of Photons	52
3.1	General Principles	52
3.2	Elastic Scattering of Photons	54
3.2.1	Overview	54
3.2.2	Mie Scattering	57
3.2.3	Rayleigh Scattering I - Electrodynamical Approach	58
3.2.4	Rayleigh Scattering II - Density Fluctuation Approach	61
3.3	Current Measurements	62
3.4	Simulation of Optical Scattering in RAT	63
3.4.1	Pure Water	63
3.4.2	LABPPO	68
4	The SMELLIE Calibration System	72
4.1	The Scattering Module of ELLIE	72
4.1.1	Overview	72
4.1.2	Operating Modes	76
4.2	The SMELLIE Hardware	79
4.2.1	Laser Heads	79
4.2.2	Laser Head Control	83
4.2.3	Optical Paths within SMELLIE	87
4.2.4	Fibre Switch	92
4.2.5	In-Situ Components	98
4.2.6	Ex-Situ Monitoring	102
5	Scattering Measurements in Water	108
5.1	Simulation of SMELLIE in RAT	108
5.1.1	The ELLIE Monte-Carlo Generator	108
5.1.2	Comparisons between Real and Simulated Data	109
5.2	Isolation of Scattered Light	118
5.2.1	Comparison of Different Optical Processes	118
5.2.2	Time Residuals	130
5.2.3	Acceptance Cuts I: Varying n_γ	134
5.2.4	Acceptance Cuts II: Varying Fibre & Wavelength	136
5.2.5	Post-Acceptance Cuts PMTs	139
5.2.6	Extraction of l_{Ray}	142
5.3	General Analysis Technique	148
5.3.1	Introduction	148
5.3.2	Path Time Calculation	149
5.3.3	Scattering Position Reconstruction	153
5.4	Pencil Beam Analysis	154
5.4.1	Reconstruction Method	154
5.4.2	Reconstruction Performance I: s	156
5.4.3	Reconstruction Performance II: d	159
5.4.4	Reconstruction Performance III: θ	161
5.4.5	Reconstruction of Other Accepted Photons	163

5.5	Angular Beam Analysis	168
5.5.1	Introduction	168
5.5.2	Application of Pencil Beam Analysis	168
5.5.3	Possible Methods for Improvement	171
5.6	Systematic Uncertainties	172
5.6.1	Detector-level Systematics	172
5.6.2	SMELLIE Systematics	174
5.7	Summary	175
6	BiPo Pileup Rejection in the Loaded Phase	180
6.1	$^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ Backgrounds	180
6.1.1	Backgrounds at SNO+	180
6.1.2	“BiPo Pileup” Events	186
6.1.3	Simulation Considerations	187
6.1.4	Time Residuals	188
6.2	Cumulative Time Residuals Method	194
6.2.1	Kolmogorov-Smirnov Γ Parameter	194
6.2.2	Γ Cut and Remaining Signal vs. Background	195
6.2.3	Post- Γ Cut BiPo Pileup Events	199
6.3	Log-Likelihood Difference Method	201
6.3.1	Log-Likelihood Parameters	201
6.3.2	$\Delta\mathcal{L}$ Cut and Remaining Signal vs. Background	206
6.3.3	Post- $\Delta\mathcal{L}$ Cut BiPo Pileup Events	208
6.4	Systematic Uncertainties	208
6.5	Conclusion	210
A	Rayleigh Scattering Formulation	212
B	Extra Results: Section 5.2.4	216
C	Extra Results: Section 5.4.2 - 5.4.4	225
	Bibliography	232

List of Figures

1.1	Schematic drawing of the SNO+ detector	2
1.2	Overview of the SNO+ electronics flow	4
1.3	Re-emission probability distribution of PPO	12
1.4	Quantum Efficiency of the SNO+ PMTs	12
1.5	Absorption lengths of bisMSB and Perylene	13
1.6	Re-emission probability distributions of bisMSB and Perylene	13
1.7	Nhits in LABPPO as a function of varying scattering length	15
2.1	Neutrino mass hierarchies	40
2.2	Feynman diagram of $2\nu\beta\beta$	45
2.3	Feynman diagram of $0\nu\beta\beta$	46
2.4	Effective Majorana ν mass vs. lightest ν mass	48
3.1	Diagram of the notation used to describe Mie scattering	55
3.2	Scattering length of water as a function of wavelength	64
3.3	Distance travelled in water by photons before scattering	65
3.4	Scattering angle ϕ in water	67
3.5	Cosine of scattering angle θ in water	67
3.6	Scattering length of LABPPO as a function of wavelength	69
3.7	Distance travelled in LABPPO by photons before scattering	69
3.8	Scattering angle ϕ in LABPPO	70
3.9	Cosine of scattering angle θ in LABPPO	70
4.1	Schematic of the ex-situ SMELLIE hardware	74
4.2	Timing delays in master mode	78
4.3	Timing delays in slave mode	79
4.4	The four SMELLIE laser heads	80
4.5	Wavelength distributions of the laser heads	82
4.6	Single-pulse time profiles of the laser heads	84
4.7	The Laser Switch and SEPIA II Laser Driver	85
4.8	Schematic of the internal design of a fused fibre taper beamsplitter	90
4.9	The SMELLIE Fibre Switch	92
4.10	Schematic of the Fibre Switch's internal design	94
4.11	Fibre Switch output attenuation - 375nm and 405nm laser heads	96
4.12	Fibre Switch output attenuation - 440nm and 500nm laser heads	97

4.13	Internal design of the SMELLIE collimators	99
4.14	Intensity profiles of a SMELLIE optical fibre beamspot	101
4.15	Schematic of a SMELLIE mounting plate	102
4.16	The SMELLIE Monitoring PMT Unit	104
4.17	PMT and sampled pulses produced by the Monitoring PMT Unit	106
5.1	PMT Map - Data vs. original Simulation	110
5.3	Schematic of the PMTs at PSUP node 37	112
5.2	PMT Map - Data vs. direction-corrected Simulation	113
5.4	PMT Map - Data vs. direction/position-corrected Simulation	114
5.5	PMT Map - Data vs. current Simulation	117
5.6	Defintion of the angle α	119
5.7	Trigger time vs. angle α for beamspot PMTs	120
5.8	Trigger time distribution for beamspot PMTs	121
5.9	Photon transit time distribution used by RAT	122
5.10	Trigger time vs. angle α for AV-reflected PMTs	123
5.11	Trigger time vs. angle α for concentrator-reflected PMTs	124
5.12	Trigger time vs. angle α for cavity-scattered PMTs	125
5.13	Trigger time vs. angle α for innerAV-scattered PMTs	125
5.14	Trigger time vs. angle α for multiple-effect PMTs	127
5.15	Trigger time vs. angle α for noise PMTs	128
5.16	Trigger time vs. angle α for all PMTs	129
5.17	Time residual vs. angle α for all PMTs	132
5.18	Time residual vs. angle α for all PMTs in the acceptance region	140
5.19	True d for all accepted-scattered photons	141
5.20	True $\cos(\theta)$ for all accepted-scattered photons	142
5.21	Acceptance Fraction f_{acc} as a function of s_l	145
5.22	The times and positions used in this analysis	149
5.23	Time residual distribution for beamspot PMTs	151
5.24	True and reconstructed d for all accepted-scattered photons (Scen. 1)	160
5.25	True and reconstructed d for all accepted-scattered photons (Scen. 4)	160
5.26	True and reconstructed θ for all accepted-scattered photons (Scen. 1)	164
5.27	True and reconstructed θ for all accepted-scattered photons (Scen. 4)	164
5.28	Reconstructed d for all accepted photons	165
5.29	Reconstructed θ for all accepted photons	165
5.30	Examples of possible scattered photon paths in an angular beam	170
6.1	The Uranium chain of radioactive decays	181
6.2	The Thorium chain of radioactive decays	182
6.3	The energy spectra of ^{130}Te and BiPo pileup decays	185
6.4	Distributions of the 1st, 2nd, 5th and 10th earliest time residuals	190
6.5	Time residual distributions of ^{212}Bi β and ^{212}Po events	192
6.6	Time residual distributions of ^{214}Bi β and ^{214}Po events	192
6.7	Time residual distributions of $^{212}\text{BiPo}$ pileup events at differing δt	193
6.8	Cumulative time residual distribution of a $^{212}\text{BiPo}$ pileup event	194

6.9	Distribution of Γ for $^{212}\text{BiPo}$ pileup events	196
6.10	Distribution of Γ for $^{214}\text{BiPo}$ pileup events	196
6.11	Remaining $^{212}\text{BiPo}$ pileup percentages after Γ cuts	198
6.12	Remaining $^{214}\text{BiPo}$ pileup percentages after Γ cuts	198
6.13	Distribution of δt for remaining $^{212}\text{BiPo}$ pileup events after Γ cut . . .	200
6.14	Distribution of δt for remaining $^{214}\text{BiPo}$ pileup events after Γ cut . . .	200
6.15	Optimiser results for $^{212}\text{BiPo}$ pileup events	203
6.16	Optimiser results for $^{214}\text{BiPo}$ pileup events	203
6.17	Distribution of $\Delta\mathcal{L}$ for $^{212}\text{BiPo}$ pileup events	205
6.18	Distribution of $\Delta\mathcal{L}$ for $^{214}\text{BiPo}$ pileup events	205
6.19	Remaining $^{212}\text{BiPo}$ pileup percentages after $\Delta\mathcal{L}$ cuts	207
6.20	Remaining $^{214}\text{BiPo}$ pileup percentages after $\Delta\mathcal{L}$ cuts	207
6.21	Distribution of δt for remaining $^{212}\text{BiPo}$ pileup events after $\Delta\mathcal{L}$ cut .	209
6.22	Distribution of δt for remaining $^{214}\text{BiPo}$ pileup events after $\Delta\mathcal{L}$ cut .	209

List of Tables

2.1	Current best measurements of ν oscillation parameters	39
2.2	Current limits on the $0\nu\beta\beta$ half-life in different isotopes	51
4.1	Inter-component ex-situ SMELLIE hardware connections	75
4.2	Specifications of the four SMELLIE laser heads	80
4.3	Measured radioactivity of the SMELLIE optical fibres	88
4.4	Nominal and measured splitting ratios of the SMELLIE beamsplitters	91
4.5	Internal components of the SMELLIE collimators	99
4.6	Measured radioactivity of the SMELLIE mounting plates	103
4.7	Connections to and from the SMELLIE Monitoring PMT Unit	105
5.1	Mapping between simulated and real SMELLIE fibres	111
5.2	Acceptance region - FS025 / 407nm / varying n_γ	135
5.3	Acceptance region - FS025 / 407nm / $n_\gamma = 1500$ / varying s_l	143
5.4	Coefficient A of the relationship between s_l and f_{acc}	146
5.5	Coefficient B of the relationship between s_l and f_{acc}	147
5.6	Coefficient C of the relationship between s_l and f_{acc}	147
5.7	Rules for calculating the line search step size and direction	156
5.8	Pencil beam: Bias and resolution of the scattering position	157
5.9	Pencil beam: Bias and resolution of d	159
5.10	Pencil beam: Bias and resolution of θ	162
5.11	Angular beam: Bias and resolution of the scattering position and d	169
5.12	Angular beam: Bias and resolution of θ	169
6.1	Backgrounds in the detector before and after RoI and FV cuts	186
6.2	Remaining $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ pileup percentages after Γ cuts	197
6.3	Remaining $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ pileup percentages after $\Delta\mathcal{L}$ cuts	206
B.1	Acceptance region for nodes 07 and 25 at 375nm, with $n_\gamma = 1500$	217
B.2	Acceptance region for nodes 37 and 55 at 375nm, with $n_\gamma = 1500$	218
B.3	Acceptance region for nodes 07 and 25 at 405nm, with $n_\gamma = 1500$	219
B.4	Acceptance region for nodes 37 and 55 at 405nm, with $n_\gamma = 1500$	220
B.5	Acceptance region for nodes 07 and 25 at 440nm, with $n_\gamma = 1500$	221
B.6	Acceptance region for nodes 37 and 55 at 440nm, with $n_\gamma = 1500$	222
B.7	Acceptance region for nodes 07 and 25 at 500nm, with $n_\gamma = 1500$	223
B.8	Acceptance region for nodes 37 and 55 at 500nm, with $n_\gamma = 1500$	224

C.1	Pencil beam: Bias of the scattering position	226
C.2	Pencil beam: Resolution of the scattering position	227
C.3	Pencil beam: Bias of d	228
C.4	Pencil beam: Resolution of d	229
C.5	Pencil beam: Bias of θ	230
C.6	Pencil beam: Resolution of θ	231

Chapter 1

The SNO+ Experiment

‘Is it that one in Japan?’

— A disappointingly common response when I say that I work on a neutrino experiment

1.1 The Detector

1.1.1 Overview

The SNO+ experiment is located at SNOLAB, in the Inco-Vale Creighton Mine #9 near the city of Sudbury in Ontario, Canada. Designed to detect neutrinos originating from a variety of sources and processes, it has inherited much of its infrastructure from the original Sudbury Neutrino Observatory (SNO), which previously operated between 1999 and 2006. However, unlike its heavy-water predecessor, SNO+ will operate using a liquid scintillator detection medium once the initial water-based commissioning phase is complete.

A schematic of the SNO+ detector is shown in Figure 1.1.

The centre-piece of the experiment is the 12m-diameter spherical *acrylic vessel*




Figure 1.1: A schematic drawing of the SNO+ detector, showing the acrylic vessel and neck (blue), the hold-up (pink) and hold-down (red) ropes, the PMT support structure (green), and the location of the deck (indicated by the figure of a person), all enclosed by the cavity (black). All components are shown approximately to scale. For clarity, the individual PMTs and the cables supporting the PMT support structure have not been shown.

(AV), which is constructed from 122 56mm-thick transparent panels [1] and will contain the detection medium. The AV is supported and secured in place by 10 *hold-up ropes* and 20 *hold-down ropes*, all manufactured from Tensylon. The hold-up ropes, of diameter 19mm, are anchored to the deck above the detector, and the 38.1mm diameter hold-down ropes are fixed to the floor of the cavity in which the experiment is located [2]. The AV is capped by the *neck* - a chimney-shaped structure 6.8m tall and 1.5m in diameter manufactured from the same material as the AV. The neck allows access into the AV for personnel, deployment of the various in-situ radioactive calibration sources, and filling of the AV with the detection medium. The top of the neck at deck-level is sealed by the *cover gas system* - comprising 3 nitrogen-filled bags and associated pipework that all together are designed to reduce the amount of radon entering the AV via the mine air, as well as compensating for pressure fluctuations in the mine that can cause the ambient radon levels to fluctuate [3].

A geodesic framework called the *PMT Support Structure* (PSUP) surrounds the AV at a diameter of 17.6m, constructed from stainless steel struts and supported from the deck by 15 stainless steel cables. There are 91 *nodes* on the PSUP where the struts' endpoints come together. The PSUP houses the ~ 9500 Hamamatsu R1408 *photomultiplier tubes* (PMTs) - 91 facing outwards away from the AV (one at each node) and the rest facing in towards the detection medium - these are used to detect the light created by events passing through the detector. Each PMT is individually housed within a hex-cell, and all inward-facing PMTs are accompanied by a *light concentrator* - a ring of mirrors designed to focus light onto its PMT's cathode, so as to increase the effective PMT area coverage of the detector to 54% (up from the 31% that can be achieved with the PMTs alone) [1][4]. The region between the PSUP and the AV is filled with ultra-pure light water, which provides the detection medium with shielding from any radioactive emanations originating in the PMTs and PSUP, and *vica versa*.

The entire detector is located in a cavern - the *cavity* - 34m high and 22m in diameter, which is also filled with ultra-pure light water acting as a shield between the PMTs and the radioactivity present in the surrounding rock of the mine. At the top of the cavity is the deck - this is where the detector's electronics are located, as well as the clean room that encloses the neck access. The entire cavity is located under 2039m of granitic rock which provides an overburden of 6010m water equivalent, reducing the muon flux through the AV to approximately 70 per day [1].

1.1.2 Electronics and Detector Readout

The PMTs provide all of the observable quantities from the detector, and so the readout of the PMTs is central to data acquisition. A basic overview of the PMT readout electronics is shown in Figure 1.2.

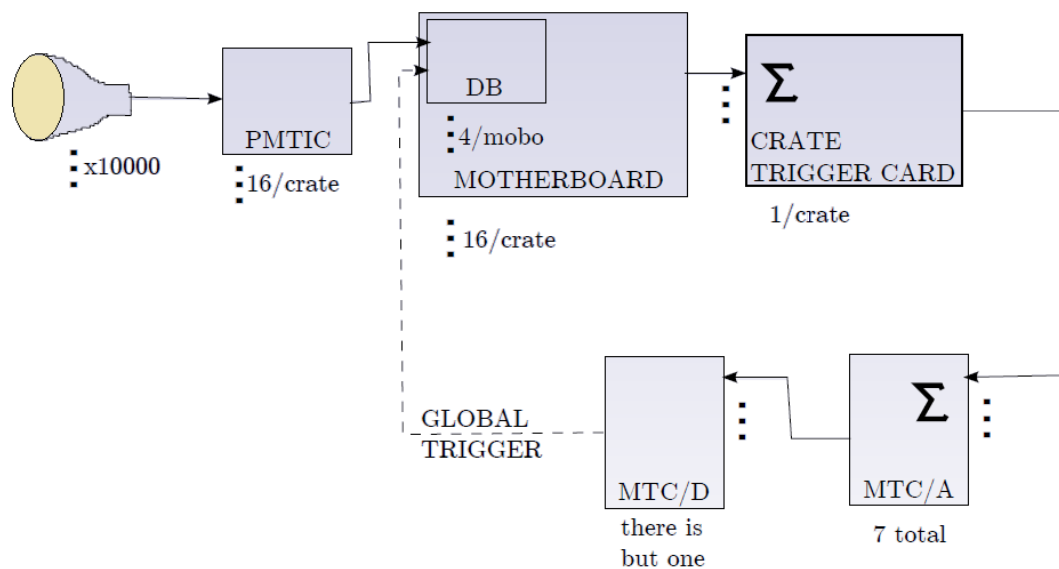


Figure 1.2: A basic overview of the SNO+ electronics flow, from a single PMT through the PMT Interface Card (PMTIC), the daughter board (DB) that is located on the mother board, through the Crate Trigger Card and on to the Analog and Digital Master Trigger cards (MTC/A and MTC/D respectively). Not shown are the XL3 boards that read out the non-trigger information from each mother board. This image has been adapted from [5].

The signal pulse produced by a single active PMT (referred to as a *channel*) is read out by a PMT Interface Card (PMTIC), with each PMTIC able to read up to 32 channels. The PMTIC passes this signal to a daughter board (DB) which performs discrimination, integration and triggering across 8 channels (therefore requiring 4 DBs per PMTIC). The DB first checks if the channel's signal has crossed the discriminator threshold required to "trigger" that particular channel, and if this check is successful, it goes on to create the per-channel trigger pulse for each available trigger type. Some example trigger types are the NHIT100 (the number of PMTs triggered in a 100ns-wide time window), NHIT20 (the number of triggered PMTs in a 20ns-wide window) and ESUM (the total charge across all triggered PMTs). The DB also measures the time window over which the signal was above the discriminator threshold, and calculates the integrals of the signal over various time windows. This results in a set of four channel-level quantities in addition to the trigger pulses: the TAC (timing amplitude conversion) and 3 charges: QHS (charge at high gain over a short integration time), QHL (charge at high gain over a long integration time) and QLX (charge at low gain over a long integration time). All of this information is then sent to a Front End Card (FEC), also known as a mother board, which receives and digitises signals from up to 4 DBs (therefore mirroring the same 32 channels that are read by each PMTIC) [5][6].

The FECs are grouped into *crates*, with 16 FECs per crate, and therefore requiring a total of 19 crates to read out all of the PMTs. The per-channel trigger signals from each FEC are sent to a Crate Trigger Card (CTC), which sums them together by type to produce a set of per-crate trigger signals - one signal per trigger type. The non-trigger information from each channel - the TAC and charges - are read out to an XL3 card (one per crate), which then waits for further instructions based on the global triggering of the detector [5].

The MTC/A (Master Trigger Card - Analog) receives the per-crate trigger signals

from the CTCs, and performs an analog summation over each trigger type to produce a single detector-wide trigger signal for each type, all of which are then passed to the MTC/D (Master Trigger Card - Digital). The MTC/D compares these overall trigger signals to preset trigger thresholds, issuing a detector-wide global trigger if the trigger signals exceed their respective thresholds within a *trigger window* of approximately 400ns starting from the earliest triggered channel [5][6]. Once a global trigger has been issued, the XL3 cards flush their stored time and charge information to the Data Acquisition (DAQ) and Event Builder systems. If a global trigger is not issued within the given time window (i.e. the detector-wide trigger signals do not exceed threshold quickly enough), the XL3 cards are reset and any channels that did in fact individually exceed their own trigger thresholds will not be associated with a global triggered event in the detector.

1.1.3 Simulation Software

SNO+ uses the RAT software framework for Monte-Carlo simulation. This C++ package is based on software originally developed by Stan Seibert of the Braidwood collaboration for simulating a generic liquid scintillator detector, consisting of a central volume of scintillator held in a vessel that is surrounded by a non-scintillating buffer region and PMTs. The original framework is the basis for software now used in a number of experiments around the world, such as DEAP/CLEAN, CLEAR and of course SNO+, but each one has its own distinct version of RAT specifically designed for purpose. RAT interfaces with standard GEANT4 libraries to provide well-tested, reliable and up-to-date simulation of physical processes, and also makes use of ROOT for storage of and access to simulated data.

SNO+ RAT versions 5.2.1 and 5.0.3 have been used for the work in Chapters 5 and 6 of this thesis respectively, and any additions and changes made specifically for the various simulations and analyses have been described in the appropriate sections.

1.2 Scintillation and Scintillator

1.2.1 Experimental Phases

The SNO+ experiment's lifetime will be divided into several phases, depending on (and named for) the detection medium contained in the AV.

The first three phases are grouped together into Phase I of SNO+.

The first of these - the *water phase* - will run from the end of 2015 to early 2016. The presence of ultra-pure light water in the detector will allow many of the instrumental systems to be tested and calibrated prior to their use in subsequent phases of the experiment, as the optical and physical properties of water are very well understood and therefore provide a known set of benchmarks with which to check the performance of the detector. In addition, the *external* radioactive backgrounds - those originating in the materials outside the detection medium itself, such as the surrounding rocks and the materials used to construct the PMTs and AV - can be measured cleanly during the water phase, without any additional contributions from the water itself. (A more detailed discussion on the various radioactive backgrounds present at SNO+ can be found in Section 6.1.1.) Finally, some physics analyses can be performed even in this early phase of the experiment, such as invisible nucleon decay and supernova neutrino searches.

The second phase of SNO+ - the *scintillator phase* - will use the liquid scintillator LABPPO as the detection medium (see Section 1.2.3), and is expected to run during the second half of 2016. During this phase, the radioactive background contribution from the LABPPO will be measured, along with other scintillator properties such as the light yield (the number of photons produced per unit of energy deposited in the scintillator) and the optical scattering and attenuation (using instrumental systems that will have been tested and calibrated in the water phase). Physics goals during

this phase include reactor and geo-neutrino measurements, low energy solar neutrinos and further work on supernova neutrinos.

The final part of phase I, beginning in 2017, will be the *loaded phase*, during which the LABPPO will be loaded with a small quantity (0.3% by mass) of the isotope ^{130}Te - a material that is able to undergo double beta decay (see Section 2.5.2). Therefore, the main focus of this phase will be the search for neutrinoless double beta decay in ^{130}Te .

Phase II of SNO+ will be designated by percent-level loadings of ^{130}Te in the LABPPO, as well as further upgrades to the SNO+ infrastructure that will be required for such a change in the detection medium.

1.2.2 Scintillation

Scintillation is just one of the many forms of *luminescence*: the emission of light from a molecule following the absorption of energy by that same molecule, with the intensity spectrum of the emitted light being determined by the specifics of the molecule in question. Luminescence in general only occurs in molecules that contain either one or two un-hybridized 2p valence electrons, i.e. 2p electrons that have *not* mixed with 2s electrons (with the mixing process being known as *hybridization*) [7]. Un-hybridized 2p electrons are referred to as π -electrons, and if these are first excited through the absorption of energy and then subsequently de-excite, their extra energy is released as luminescence.

The direct excitation of molecules containing π -electrons can only proceed into excited singlet states - moving directly from the ground state into excited triplet states is forbidden by spin. However, excited triplet states can still be reached indirectly, via inter-system crossing from singlet to triplet states [7]. Excited π -electron singlet and triplet states can also be produced indirectly if the molecules are first ionised and then

undergo ion-electron recombination. A simple statistical argument shows that 75% of all ion-electron recombinations will result in excited triplet states, with 12% going to excited singlet states [7][8], and the rest proceeding directly to the ground state. Therefore, if molecules undergo ionisation more readily than direct excitation, there will be more excited π -electron triplet states produced than excited singlet states.

There are three fundamental types of luminescence, which differ according to their emission spectra and lifetimes [7].

The first type is *fluorescence* - the emission of light via the de-excitation of a π -electron from the first excited singlet state down to ground, occurring with a lifetime of $\sim 10^{-9}$ s and an exponentially decaying intensity spectrum (as a function of time). Although π -electrons can be excited into second and higher-order singlet states, these states de-excite in a non-radiative manner down to the first excited state very quickly - with a lifetime of $\sim 10^{-11}$ s, thus leading to standard fluorescence.

The second type of luminescence - *phosphorescence* - is also emitted with an exponential time spectrum, but has a longer lifetime of 10^{-4} s and is caused by de-excitation from excited triplet states down to ground. Phosphorescence is also emitted at longer wavelengths (i.e. lower energies) than fluorescence.

The final type of luminescence is known as *delayed fluorescence*, and is caused by an excited π -electron triplet state acquiring enough thermal energy to jump to the corresponding excited singlet state, and from there undergoing standard fluorescence. As a result, delayed fluorescence is emitted at the same wavelengths as standard fluorescence, but with a non-exponential spectrum.

The energy transferred to a medium from a particle passing through it is dependent on both the particle's and medium's properties:

$$\frac{dE}{dr} = \frac{4\pi z^2 e^4}{m_e v^2} N Z \left[\ln \left(\frac{2m_e v^2}{I} \right) - \ln(1 - \beta^2) - \beta^2 \right] \quad (1.1)$$

where $\frac{dE}{dr}$ is the energy transferred from particle to medium along an element dr of the particle's path, m_e and e are the electron mass and charge respectively, v , E and ze are the particle's velocity ($= \beta c$), energy and charge respectively, N and Z are the medium's number density and atomic number respectively (and the medium is assumed to be isotropic and homogeneous), and I is a medium-specific constant that describes the mean excitation and ionisation potential [7]. Due to the dependence of $\frac{dE}{dr}$ on the particle energy and velocity (and therefore mass), heavy and light particles lose energy at different rates and via different processes.

Heavy, non-relativistic particles are more likely to ionise the medium's molecules than directly excite them, resulting in more excited π -electron triplet states than singlet states, and therefore a larger fraction of luminescence emitted via phosphorescence than through fluorescence. Heavy particles also have a large $\frac{dE}{dr}$ and therefore lose a large fraction of their energy over a short distance, creating many closely-spaced molecular ionisations and excitations. The ionisations can be close enough to interact with each other, and this has the effect of de-exciting some of the excited molecules non-radiatively, thus reducing the total energy that is released radiatively via luminescence - a process known as *ionisation quenching*. (Other forms of quenching which are not related to the particle itself - such as *chemical quenching* - can additionally be present, and serve to further reduce the light output of the scintillator.)

In contrast, light, relativistic particles interact with the medium primarily via elastic collisions with molecular electrons, producing more molecular excitations than ionisations and therefore more excited π -electron singlet states than triplet states, leading to a greater contribution to the emitted light from fluorescence than from phosphorescence. Light particles are easily subjected to large-angle deflections, and

have a smaller $\frac{dE}{dr}$, allowing them to travel large distances through the medium and only produce a few excitations along their path, resulting in (almost) no ionisation quenching.

1.2.3 The SNO+ Scintillator: LABPPO

The scintillator that will be used in the SNO+ experiment is comprised of two parts. The first component is a solvent called *linear alkyl benzene* (abbreviated to LAB) - an organic compound with the chemical formula $C_6H_5C_{12}H_{25}$ [9]. LAB is excited by physical processes occurring within the detector, but does not itself emit light very easily - its emission and absorption wavelengths are very similar, and so any light that might be emitted is likely to be immediately reabsorbed. Therefore, a second component is added to the LAB: a primary fluor *2,5 Diphenyloxazole* (commonly known as PPO) at a concentration of 2g/l of LAB. Any excitations of the LAB molecules are non-radiatively transferred to the PPO molecules much faster than the LAB absorption timescale, and the PPO then emits this energy as light at wavelengths predominantly between 350 and 400nm, as shown in Figure 1.3. The combination of LAB and PPO will hereafter be referred to simply as *LABPPO*.

Figure 1.4 indicates that, as it is, the PPO re-emission wavelength range is already very close to that at which the SNO+ PMTs achieve their highest quantum efficiency. However, the addition of a *wavelength shifter* may improve this matching even further - that is, a compound that will absorb the light emitted from the PPO, and re-emit it at a different wavelength. Two compounds under consideration at the time of writing are *bisMSB* and *Perylene*, both of which absorb strongly at wavelengths below 400nm, as shown in Figure 1.5 (i.e. the same wavelengths at which PPO predominately re-emits), and then re-emit light either between 380 and 450nm (*bisMSB*) or between 440 and 500nm (*Perylene*), as shown in Figure 1.6 - matching very well with the wavelengths at which the PMTs achieve their highest efficiency.




Figure 1.3: The normalised probability of light re-emission as a function of wavelength for the primary fluor 2,5 Diphenyloxazole (PPO). This plot has been adapted from [10].




Figure 1.4: The normalised quantum (photocathode) efficiency of the SNO+ PMTs, i.e. the probability of a photoelectron being liberated from the cathode by an incident photon of a given wavelength [11]. The highest efficiency ($\sim 22\%$) lies at a wavelength of $\sim 440\text{nm}$, and between 350 and 400nm (the wavelengths which PPO is most likely to re-emit at) the efficiency is $\sim 20\%$.




Figure 1.5: The absorption length of the wavelength shifters bisMSB (green) and Perylene (blue) as a function of wavelength. A smaller absorption length indicates that more light is absorbed. This plot has been adapted from [10].




Figure 1.6: The normalised probability of light re-emission as a function of wavelength for the wavelength shifters bisMSB (green) and Perylene (blue). This plot has been adapted from [10].

As previously noted in Section 1.2.1, during both Phase I's loaded phase and Phase II, the LABPPO (and potential wavelength shifter) will be combined with a mixture containing the isotope ^{130}Te . The exact contents of the overall cocktail are still being investigated as of writing, but several potential candidates have been added to RAT for simulation purposes. The particular one selected for use in Chapter 6 of this thesis contains LABPPO, bisMSB, various compounds that allow Tellurium to actually be loaded into the scintillator, and finally natural Tellurium itself at a level of 0.3% of the total mass of the entire cocktail (corresponding to Phase I of the experiment). Of this natural Tellurium, 34% (approximately 800kg) is ^{130}Te [12]. For convenience, this overall cocktail will be referred to simply as *Te_LABPPO*.

1.2.4 Effects of Scintillator Properties

The detector's performance will obviously be affected by the properties of the detection medium, whatever it may be, but given that most of the new physics that SNO+ is sensitive to will potentially be encountered during the scintillator and loaded phases, it is of particular importance that the properties of LABPPO are well understood. Two of the more important properties for the purposes of SNO+ (but by no means the only ones) are the scintillator's *light yield* (the number of photons emitted per MeV of deposited energy) and its *optical scattering*.

The light yield is directly related to the number of PMTs that will be triggered within a given trigger window - a quantity referred to as the *Nhits* of the triggered event, and therefore sets a lower limit on the energy of events that SNO+ can observe: a scintillator with too low a light yield will result in only a few PMTs being triggered by a single physics event - and this small Nhits may not be high enough to trigger the entire detector and create an observed event. However, having a very high light yield can also potentially create problems, as a high-energy event occurring in a very high light yield scintillator could saturate the detector - triggering every PMT and

potentially causing damage.

The scattering of light in the scintillator is also important to understand, as it will have a major impact on the total time that a photon spends in the scintillator before entering and triggering a PMT, as illustrated by Figure 1.7.




Figure 1.7: The number of PMTs that were triggered within a single trigger window (Nhits) for 2.526MeV electrons located at the centre of a detector filled with LABPPO, where the scattering length l_{Ray} of the LABPPO has been scaled by various factors. No scaling represents the default value of l_{Ray} that is currently used in RAT.

When using the default scattering length, the vast majority of electrons produce a single triggered event with Nhits between 1200 and 1450. However, a small number of electrons produce two consecutive triggered events - the first of which has Nhits of ~ 1300 , and the second with Nhits < 50 . These low-Nhits events are known as *retriggers*, and are caused by a photon taking longer than a single trigger window to enter and trigger a PMT, instead doing so during the next window. (It should be noted that these retriggers cannot be reliably reconstructed in position, time or energy, since they do not actually represent physical particles depositing energy in the scintillator.) This delay between photon emission and PMT triggering - which

only occurs rarely, hence why only few PMTs are found in a single retrigger - can be caused by a number of effects, including scattering of photons in the scintillator.

It can be seen that when scattering length is increased from the default value, there is very little change in the overall Nhits distribution. This is because the default scattering length of photons in LABPPO is already much larger than the size of the detector itself (as will be detailed further in Section 3.4.2), and so increasing it further does not significantly impact how photons traverse the scintillator or how much they scatter.

However, there is a very noticeable trend in the distribution of Nhits as the scattering length is progressively shortened: the average Nhits of the first triggered event per electron decreases, and both the number of retriggers per electron and the Nhits per retrigger increase, as indicated by the low-Nhits peak both growing in height and also shifting to slightly higher Nhits. However, the total number of triggered PMTs produced by a single electron (i.e. summed across all triggered events - the first triggered event and all subsequent retriggers) remains the same at ~ 1350 . This change in the distribution is due to the fact that when the scattering length is decreased, it is more likely that a particular photon will scatter, and so it is more likely to spend more than a single trigger window within the scintillator before entering a PMT than if the scattering length is long. Therefore, the Nhits of the first event will decrease, since fewer photons successfully make it to PMTs within the first trigger window after emission from the scintillator, and the number and Nhits of any retriggers will increase due to more late-arriving photons - thus producing the very clear two-peak distributions that can be seen in the $(l_{Ray} \times 2)$ and $(l_{Ray} \times 10)$ scenarios.

The Nhits of an event is the most important quantity for energy reconstruction, and therefore it can be seen that the absolute scattering properties of the scintillator have a significant knock-on effect on the overall reconstruction performance of the detector.

For example, if the Nhits of an event is smaller than expected due to more scattering occurring, the reconstructed energy will also be lower than it should be, and the energy resolution of the detector will become larger (i.e. worse) - in part simply due to the larger uncertainty in a smaller Nhits. Higher levels of scattering also mean that the photons' paths through the scintillator will deviate further and further away from a simple straight line between emission and PMT, and this could potentially affect the position reconstruction - which relies on certain assumptions being made about the photon trajectories within the scintillator volume.

The end result of a potential worsening of the detector's performance is that the eventual extraction of a ^{130}Te neutrinoless double beta decay signal during the loaded phase will be made considerably more difficult: the energy and position reconstruction are directly related to how many non-signal (background) events overlap in energy with the signal, and so poor position and energy resolution will affect the selection of signal events and rejection of backgrounds.

Chapter 2

Neutrinos

‘It’s reinvigorated my interest in Physics, which was stamped out by seven years of secondary school.’

— Scott Stevenson, on writing the theory chapter of his own thesis

2.1 A Brief History of Neutrino Physics

2.1.1 The Need for a New Particle

Many scientists in the first two decades of the twentieth century were keen to study the properties of radioactivity - the α , β and γ decays of certain atomic nuclei, which had only recently been discovered at the turn of the century. Experiments had clearly shown that α and γ decays occur in discrete emission lines - that is, the respective daughter particles are emitted with specific energies that are only dependent on the particular nucleus that decayed. Therefore, it seemed entirely reasonable to assume that the same result would be observed for β particles as well.

However, in 1914 it was discovered that the e^- emitted by the β decay of ^{214}Pb has a continuous energy spectrum [13]. Physicists were split between two different interpretations of this result - either the e^- did have a single energy that was subsequently broadened into a spectrum through other atomic processes, or it was in

fact emitted with a non-discrete energy and the assumption that β decays occurred completely analogously to α and γ decays was incorrect. Further experiments, such as [14] and [15], revealed that the latter approach was correct: the e^- produced by a β decay has a continuous energy spectrum at emission, rather than one that is acquired through the smearing of an originally discrete energy.

Assuming that energy conservation still holds, the only possible explanation for the continuous β energy spectrum is the undetected presence of a second daughter particle - one which carries away a certain fraction of the total decay energy such that the e^- is then left with an energy anywhere between 0 and the maximum Q -value of the decay. This solution was first proposed by Wolfgang Pauli in 1930, in his famous letter addressed to the “*Dear Radioactive Ladies and Gentlemen ...*” who were meeting at Tübingen to discuss issues in radioactivity [16]. Pauli’s new particle, which he actually called the “neutron”, had to be electromagnetically uncharged and spin- $\frac{1}{2}$, with a mass of not more than 1% of the proton mass, in order to satisfy the constraints given by the results of β decay experiments. (This in fact makes the ν_e only the third Standard Model particle to be theorised, being preceded only by the e^- and the γ , and itself preceding the second and third lepton generations and the entire quark sector.)

Pauli himself acknowledged that at the time his new particle was a “desperate remedy”, and in fact encouraged the readers of his letter to explore other solutions to the problem of the continuous β energy spectrum. However, only two years later in 1932 the true neutron was discovered [17], allowing Enrico Fermi to develop his theory of β decay [18]. Fermi included Pauli’s particle - now called the *electron neutrino* (Italian for “little neutral one”) and denoted by ν_e - in his framework:

$$n \rightarrow p + e^- + \bar{\nu}_e \tag{2.1}$$

but it would take many more years for experimental evidence to prove that Pauli's original idea was correct.

2.1.2 Discovery of ν_e

Twenty-two years after Pauli's original proposal, experiments using the *electron capture* (EC):



of ^{37}Ar showed the first evidence of the existence of neutrinos [19]. Due to the two-body final state of EC, the recoil energy of the daughter ^{37}Cl nucleus is constant and can be calculated - along with the nucleus' corresponding velocity - using simple energy-momentum conservation. The experiments proceeded to measure this recoil velocity, and found it to be in good agreement with the theoretical prediction, indicating that EC must include the release of a ν_e alongside the recoiling nucleus.

The direct discovery of the $\bar{\nu}_e$ occurred a year later in 1953, in an experiment performed by Clyde Cowan Jr. and Frederick Reines at the Hanford nuclear reactor (Washington State, USA) [20] - the same reactor that produced the plutonium that was used in the Trinity nuclear test in 1945.

Reactor neutrinos (discussed further in Section 2.2.2 below) were used to stimulate the *inverse β decay* (IBD):



of CdCl_2 dissolved in water. The daughter e^+ would very quickly annihilate with

an e^- in the water, producing a first light pulse that could be observed by the surrounding 300l of liquid scintillator, and at some later time (due to its heavier mass and therefore lower velocity), the daughter n would be captured by another CdCl_2 molecule, releasing a γ to give a second scintillation signal. The delayed coincidence between the e^+ annihilation and n capture signals would be proof of the occurrence of IBD, and therefore of the existence of $\bar{\nu}_e$ emissions from the reactor.

Unfortunately, due to experimental limitations, the results of this first attempt were somewhat inconclusive, but Cowan and Reines repeated the experiment three years later in 1956, now at the Savannah River reactor (South Carolina, USA) [21]. This new experiment used 4200l of liquid scintillator and much more effective shielding against other radioactive backgrounds, and conclusively observed the delayed coincidence between the e^+ annihilation and n capture signals, confirming the existence of the $\bar{\nu}_e$. (This also made the $\nu_e/\bar{\nu}_e$ only the fourth Standard Model particle to be conclusively discovered, after the e^\pm , γ and μ^\pm .) Reines was awarded the 1995 Nobel Prize in Physics for this work, with Cowan unfortunately having died in 1974.

2.1.3 Discovery of ν_μ

In the intervening years between Pauli's proposal and Cowan and Reines' studies, experiments using cosmic rays had resulted in the discoveries of muons μ^\pm (1936), charged pions π^\pm (1947) and the neutral pion π^0 (1950). Shortly after their discovery, it was found that π^\pm decay into a μ^\pm along with a second, neutral daughter particle of undetermined mass [22], but at the time it was unclear as to whether this second particle - tentatively named the *muon neutrino* and denoted by ν_μ - was the same as the ν_e emitted in β decays. (Interestingly, the original discoverers of π^\pm decay concluded that the second daughter particle cannot in fact be a low-mass neutrino, due to an incorrect measurement of the ratio between the π^\pm and μ^\pm masses.)

In 1962, and using an almost pure beam of accelerator ν_μ (discussed in Section

2.2.2) produced at the Brookhaven National Laboratory (New York, USA), physicists sought to test the idea that the ν_μ and ν_e were the same particle by measuring the rates of four processes:

$$n + \nu_\mu \rightarrow p + \mu^- \quad (2.4)$$

$$p + \bar{\nu}_\mu \rightarrow n + \mu^+ \quad (2.5)$$

$$n + \nu_\mu \rightarrow p + e^- \quad (2.6)$$

$$p + \bar{\nu}_\mu \rightarrow n + e^+ \quad (2.7)$$

If $\nu_\mu = \nu_e$, then processes 2.6 and 2.7 would be possible - the final-state e^\pm would be seen in the experiment's tracker as electromagnetic showers. In fact, although many final-state μ^\pm were observed - proving that processes 2.4 and 2.5 did indeed occur, almost no e^\pm showers were seen (and those that were were deemed to be a result of kaons decaying to e^\pm , producing ν_e contamination of the ν_μ beam and resulting in IBD, as in Equation 2.3) [23]. The experiment was repeated two years later at CERN [24], where this result was verified - confirming that a second type of neutrino existed.

2.1.4 Subsequent Experimental Results


Helicity Measurements

The *helicity*, \mathcal{H} , of any particle is defined by its spin $\underline{\sigma}$ and momentum $\underline{\rho}$:

$$\mathcal{H} = \frac{\underline{\sigma} \cdot \underline{\rho}}{|\underline{\rho}|} \quad (2.8)$$

In 1958, only two years after its discovery, the helicity of the ν_e was measured [25]

using the EC of ^{152}Eu :



As ^{152}Eu undergoes EC at rest, the ν_e and γ - which have opposite spin - are emitted back to back, and therefore equation 2.8 dictates that $\mathcal{H}(\nu_e)$ and $\mathcal{H}(\gamma)$ must be equal. The latter of these is relatively easy to measure, as it is simply the γ 's circular polarisation. This measurement showed that $\mathcal{H}(\gamma) = \mathcal{H}(\nu_e) = -1$, and together with the established electroweak theory implied that neutrinos are massless, as will be discussed in Section 2.2.1 below.

Number of Light Neutrino Flavours

With the discovery of the τ^- at SLAC in 1975 [26] and given that both of the neutrinos known at the time were associated with a charged lepton, it would be easy to assume that there had to be at least one more neutrino waiting to be discovered. This was tested by the ALEPH detector at LEP [27], where the total decay width of the Z^0 boson was calculated using very precise measurements of its resonance peak. From this, and the calculable hadronic and charged-leptonic decay widths, it is then possible to find the neutral-leptonic decay width, and therefore (assuming equal decay widths for each $\nu_l\bar{\nu}_l$ final state, where l represents the neutrino flavour) the maximum number of possible $\nu_l\bar{\nu}_l$ final states with total mass $< m_Z$ - equivalent to the number of neutrino flavours with mass $< \frac{m_Z}{2}$ (referred to as *light* neutrinos).

This number was found to be 3.01 ± 0.15 , completely consistent with there being exactly three known neutrino flavours - one for each known charged lepton - but

it should be noted that this does not preclude the existence of one or more *heavy* neutrinos with a mass $> \frac{m_Z}{2}$.

2.1.5 Discovery of ν_τ

Even with the results of the ALEPH experiment, it still took another eleven years for the ν_τ to be formally discovered. This was finally achieved at the DONUT experiment [28], which first used the decays of charmed mesons D_S^- and the subsequent decay of the daughter τ^- to produce a beam of ν_τ :

$$\begin{aligned} D_S^- &\rightarrow \tau^- + \bar{\nu}_\tau \\ \tau^- &\rightarrow X^- + \nu_\tau \end{aligned} \tag{2.10}$$

The ν_τ beam was then made incident on nuclear emulsion plates, such that the ν_τ would interact in the charged-current interaction (i.e. involving a mediating W boson):

$$A + \nu_\tau \rightarrow B + \tau^- \tag{2.11}$$

and the daughter τ^- could be then identified via their subsequent decays.

2.2 Neutrinos in the SM

2.2.1 Fundamental Properties

Based on the previously discussed experiments, we now know that the three neutrinos - the electron neutrino ν_e , the muon neutrino ν_μ and the tau neutrino ν_τ - are some

of the most fundamental particles of the universe. They are uncharged and colourless spin- $\frac{1}{2}$ fermions (leptons), meaning that they do not interact via the electromagnetic or strong nuclear forces, but instead only through the weak nuclear force and (presumably) gravity, making them very difficult to detect due to the extremely short range of weak interactions. Indeed, a calculation of the neutrino interaction cross section performed very early in the history of Neutrino Physics by Hans Bethe and Rudolf Peierls gave a fantastically small value of $< 10^{-44}\text{cm}^2$ [29]. Being fermions, each neutrino *flavour* ν_l ($l = e, \mu$ or τ) has a corresponding antineutrino $\bar{\nu}_l$.

As their names suggest, the three neutrino flavours are associated with their charged lepton partners l^- , due to one of the fundamental neutrino production mechanisms being the decay of the weak W^- boson into $l^-\nu_l$ pairs. Neutrinos can also be produced by the decay of the other weak boson Z^0 into $\nu_l\bar{\nu}_l$ pairs. Neutrinos and antineutrinos have helicity \mathcal{H} of -1 and +1 respectively, and each of the three neutrino flavours ν_l has an individual *leptonic family number* $L_l = +1$, with the antineutrinos $\bar{\nu}_l$ possessing $L_l = -1$.

In the quantum theory that underlies the Standard Model (SM), the weak force only interacts with particles which are in a left-handed chiral state and antiparticles which are in a right-handed chiral state, and therefore neutrinos must be always be created left-handed in terms of chirality. On the other hand, as previously discussed in Section 2.1.4, experiments have shown that neutrinos are always observed with helicity $\mathcal{H} = -1$, i.e. left-handed helicity. Helicity and chirality would therefore appear to be equivalent for neutrinos (and antineutrinos), but this can only occur if they are massless particles.

Another way of looking at this is to consider the SM Higgs mechanism for generating particle masses. Unlike the weak force, the Higgs field couples to both left- and right-handed chiral states of a particle field in order to generate mass terms in

the Lagrangian that take the form [30]:

$$\mathcal{L}_{Yuk} = -\frac{c_p v}{\sqrt{2}}(\bar{p}_R p_L + \bar{p}_L p_R) \quad (2.12)$$

where p_L and p_R are the left- and right-handed chiral components of a general particle field, $\bar{p}_{L,R} = p_{L,R}^\dagger \gamma^0$ is the adjoint chiral component, c_p is a coupling constant and v is the Higgs field's vacuum expectation value. The mass of the particle is then given by:

$$m_p = \frac{c_p v}{\sqrt{2}} \quad (2.13)$$

However, in the electroweak sector of the SM, the neutrino fields do not have a right-handed component [30], and so such mass terms cannot be constructed - i.e. the Higgs mechanism cannot be used to give neutrinos mass.

The three leptonic family numbers L_l ($l = e, \mu$ or τ) and the *total lepton number* $L = \sum L_l$ are all conserved quantities in the SM. However, unlike some other conserved quantities, lepton number conservation is not fundamentally connected to the invariance of a particle's equation of motion under a particular symmetry - instead, it is assumed, simply because non-conservation had never been experimentally observed at the time when the various conserved quantities were being formulated.

2.2.2 Sources of Neutrinos

Neutrinos are produced in a number of processes, both naturally and artificially.

Accelerator

Particle accelerators can be used to produce very intense beams of high-energy neutrinos, by colliding a proton beam with a fixed target - producing large numbers of charged pions π^\pm and kaons K^\pm . These heavy particles have very short lifetimes, on the order of 10^{-8} s, and so will decay in a variety of ways - for example:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (2.14)$$

$$K^+ \rightarrow \mu^+ + \nu_\mu \quad (2.15)$$

$$K^+ \rightarrow \pi^0 + \pi^+ \rightarrow \pi^0 + \mu^+ + \nu_\mu \quad (2.16)$$

(and similarly for their antiparticles). By using a strong magnetic field to focus the π^\pm and K^\pm into a beam as they are produced, the resulting daughter particles will also be emitted in a beam, and the μ^\pm can be removed through the use of shielding in the forward beam direction, leaving only the ν_μ *accelerator neutrino* beam. This technique was originally proposed in 1960 [31], and was subsequently used in the experiments that discovered the ν_μ and ν_τ , as previously discussed in Sections 2.1.3 and 2.1.5.

Reactor

The other main source of artificially-created neutrinos is nuclear reactors. Very heavy isotopes - ^{235}U , ^{238}U , ^{239}Pu and ^{241}Pu - undergoing β decay produce huge numbers of $\bar{\nu}_e$ *reactor neutrinos* over a wide range of energies, the exact values of which depend on the specific isotope being used in the reactor. These $\bar{\nu}_e$ are produced isotropically rather than in a beam as with accelerator neutrinos, but this is offset somewhat by the much higher total neutrino flux from reactors compared to accelerators. The experiments that originally discovered the $\bar{\nu}_e$, noted above in Section 2.1.2, used

reactor neutrinos.

Geologic

Of course, radioactive isotopes do not undergo β decays only when they are found in nuclear reactors. Naturally-occurring radioactive materials in the Earth - for example: ^{238}U , ^{232}Th and their daughters, as well as unrelated isotopes such as ^{40}K - all produce $\bar{\nu}_e$ *geo-neutrinos* as they decay. Fundamentally there is therefore no difference between reactor and geo-neutrinos, apart from where they are produced. Studying geo-neutrinos can tell scientists about the geological properties of the Earth, such as isotope abundances and spatial densities, but only a few such experiments have been performed, such as at KamLAND [32] and Borexino [33]. SNO+ will be sensitive to geo-neutrinos, and there are also other proposals for future experiments that are specifically designed to study them.

Solar

By far the largest source of neutrinos - natural or artificial - that are detected on Earth is the Sun. The various nuclear fusion processes that occur in the Sun's core are well-described by the Standard Solar Model, in which the majority of the Sun's energy is produced via the fusion of hydrogen into helium through the *pp cycle* [30]:



(The numerous intermediate steps of the cycle have not been shown.) This overall process accounts for 98.4% of the Sun's energy output, with the rest originating from a second set of reactions known as the *CNO cycle*, which has the same overall reaction as the pp cycle but proceeds through the heavier carbon, nitrogen and oxygen atoms (hence its name), rather than through the pp cycle's deuterium and helium-3. From

the Sun's mean energy flux at Earth - known as the *solar constant* - the total flux of these ν_e *solar neutrinos* at Earth can be calculated as $\approx 6.5 \times 10^{10} \text{cm}^{-2}\text{s}^{-1}$, with energies ranging from 0.384MeV up to 18.77MeV.

Atmospheric

Extremely high energy neutrinos with energies of at least 1GeV are created when cosmic rays - consisting mainly of high energy protons - interact with the Earth's atmosphere and produce π^\pm , in a very similar process to that used to produce π^\pm in accelerators. Once again, π^\pm will themselves decay to μ^\pm as given previously in Equation 2.14, but now these μ^\pm can themselves go on to decay:

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \quad (2.18)$$

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu \quad (2.19)$$

Note that this decay only occurs if the μ^\pm energy is smaller than a few GeV - if this is not the case, relativistic effects result in the μ^\pm impacting the Earth's surface before it can decay. In the case when the μ^\pm does decay, it can be seen that the entire process from the decay of a single π^\pm to the production of the e^\pm results in three *atmospheric neutrinos*: two muon-flavoured and one electron-flavoured. This expected ratio of 2:1 will become important in the discussion of neutrino flavour oscillations in Section 2.3.2.

Supernovae

By far the rarest but most spectacular sources of neutrinos are core-collapse supernovae - the final stage of a massive star as it collapses under its own gravity when

nuclear fusion ceases. Neutrinos are an important component of supernovae - carrying $\sim 99\%$ of the total energy released in the event, spread across all three flavours of both neutrinos and antineutrinos [30]. These *supernovae neutrinos* are created in two distinct processes, the first of which occurs as the core of the star collapses. The density of matter increases to the point that free protons and heavy atoms undergo EC (Equation 2.2), creating a pure neutron core and emitting large numbers of ν_e and energy which heats the core. The second phase of neutrino creation occurs when the core subsequently cools, producing $\nu_l\bar{\nu}_l$ pairs in all three flavours. The exact flavour of the produced neutrino dictates its energy, which can range from 10MeV (ν_e) up to 30MeV (ν_μ and ν_τ).

Although supernovae neutrinos had been considered as early as 1966 [34], it was not until the 23rd of February, 1987 that the theory was tested. This was the famous “Supernova 1987A” - the first time that neutrinos originating outside our own Solar System were observed, at three detectors: Kamiokande II (Japan), Irvine-Michigan-Brookhaven (USA) and Baksan (Russia) [30]. Although these detectors could only observe $\bar{\nu}_e$ interactions via IBD (Equation 2.3), a total of 24 events were observed in a time window of just 12 seconds, and from these a wealth of information - about both supernovae and neutrinos - could be gathered.

2.3 Beyond the SM: Neutrino Flavour Oscillations

2.3.1 Introduction

The concept of a particle oscillating between different flavours has been studied in the quark sector of the SM for a number of years. What we think of as the three down-type quarks - d , s and b - are in fact the *mass eigenstates*, and the down-type *flavour eigenstates* - that is, the particles that actually interact via the weak force, denoted by d' , s' and b' - are found by applying the Cabibbo-Kobayashi-Maskawa

(CKM) unitary mixing matrix to the mass eigenstates. Similarly, the neutral K^0 and its antiparticle \bar{K}^0 are the mass eigenstates of the kaon, and can be thought of as being superpositions of the flavour eigenstates K_S^0 (“K-Short”) and K_L^0 (“K-Long”).

In the SM electroweak theory, the left-handed fundamental particles are arranged into generational doublets, i.e. the first generations of quarks and leptons are denoted by $\begin{pmatrix} u \\ d' \end{pmatrix}_L$ and $\begin{pmatrix} e \\ \nu_e \end{pmatrix}_L$ respectively. Qualitatively, it is not unreasonable to expect a certain mirroring of aspects between the two sectors - for example, since the lower components of the quark doublets are known to experience flavour mixing, perhaps the lower components of the leptons doublets - the neutrinos - do so as well. If so, then the three neutrino flavours ν_e , ν_μ and ν_τ that we observe are in fact superpositions of some set of neutrino mass eigenstates, denoted by ν_1 , ν_2 and ν_3 , where the exact mixing is governed by some matrix analogous to the quark sector’s CKM matrix. Following from this, if the mass eigenstates propagate through space with even slightly different phases, the result is that the overall flavour superposition changes over time.

2.3.2 Evidence for ν Flavour Oscillations

Atmospheric Neutrinos

As noted previously in Section 2.2.2, the ratio of muon-type to electron-type atmospheric neutrinos, denoted by R :

$$R = \frac{\text{number of } \nu_\mu + \text{number of } \bar{\nu}_\mu}{\text{number of } \nu_e + \text{number of } \bar{\nu}_e} \quad (2.20)$$

is expected to have a value of 2, and this should be constant regardless of the neutrinos’ incoming direction, since cosmic rays impact the Earth’s atmosphere isotropically and so the subsequent daughter particles are also produced without any locational biases.

In 1998, the Super-Kamiokande experiment - a water Cherenkov detector in Japan - was able to measure R , in both the downwards (produced in the atmosphere directly above the detector) and upwards (produced in the atmosphere on the other side of the Earth and then travelled through the planet and into the detector) directions. The different types of neutrino were distinguished by their different final-state leptons:

$$e^- + \nu_e \rightarrow e^- + \nu_e \quad (2.21)$$

$$e^- + \nu_\mu \rightarrow \mu^- + \nu_e \quad (2.22)$$

$$e^- + \bar{\nu}_\mu \rightarrow \mu^- + \bar{\nu}_e \quad (2.23)$$

where the e^- produces a very diffuse Cherenkov ring, compared to the much sharper one resulting from the μ^- [30].


The initial results from Super-Kamiokande showed that there was a clear difference between the observed and expected values of R [35], consistent with the discrepancy that might be expected from neutrino flavour oscillations (among other explanations), and further measurements found that the major source of discrepancy is found in the upwards-travelling ν_μ - there were fewer such events than expected [36]. It should be noted that this observation on its own does not conclusively prove that neutrino flavour oscillation occurs - it is possible that the ν_μ could be decaying, or that they are undergoing one-way flavour changes rather than oscillations. However, all of these involve a change of the neutrino's state over time, implying that there is a sense of time passing in the neutrino's frame of reference, meaning that they cannot be travelling at the speed of light and therefore that they are massive - directly contradicting the SM, and showing that at least in some way, neutrino properties cannot be explained

by the existing framework of electroweak theory.

In hindsight along with later evidence from other sources (detailed below), the results from Super-Kamiokande point towards the idea that some of the upwards-travelling ν_μ have oscillated into at least one of the other flavours - this is not seen for downwards-travelling ν_μ since these particles only travel a few hundred kilometres before entering the detector, whereas upwards-travelling particles must travel over 13,000km through the Earth, allowed them enough time and distance to oscillate. Due to no corresponding increase in upwards-travelling ν_e being observed at Super-Kamiokande, the preferred solution to the ν_μ deficiency is that some of them have oscillated into ν_τ [37].

The Solar Neutrino Problem

At the same time as the Brookhaven experiments to understand the ν_μ were being performed in the early 1960s, much progress was being made in developing the Standard Solar Model of nuclear fusion processes that are occurring within the Sun. Due to the opacity of the Sun's outer layers to the photons that are produced deep within its core, the only way to study these processes is to detect the emitted solar neutrinos, which do not interact with the other matter in the Sun and so can travel to the Earth unhindered. In 1963 and 1964, John Bahcall and others calculated the total solar neutrino flux from all nuclear processes expected to occur in the Sun [38], and Raymond Davis proposed an experiment to detect these neutrinos on Earth [39] using a large volume of chlorine as a target in the IBD reaction:



The ${}^{37}\text{Ar}$ atoms would be extracted from the chlorine solution, and the number of these would then be a direct measure of the number of incident $\bar{\nu}_e$.

The experiment - located at the Homestake Mine (South Dakota, USA) - began in 1968, and even early indications showed that there was a significant difference between the theorised and measured solar neutrino fluxes. The experiment ran for almost twenty-five years, and this remained a consistent aspect of the results: the measured solar neutrino flux was approximately a third of the expected value [40] - a discrepancy known as the *Solar Neutrino Problem*.

By the end of the twentieth century, all indications were that the Standard Solar Model was indeed an excellent description of the nuclear reactions occurring inside the Sun, and therefore the Solar Neutrino Problem could not be explained by a simple lack of understanding of these processes. This left only one other alternative - that the neutrinos themselves were behaving in an unexpected way.

The SNO experiment finally provided a solution to the problem, with its simultaneous measurements of: 1) the solar neutrino flux at Earth across *all* neutrino flavours, and 2) the solar neutrino flux at Earth from only ν_e (the same as had been measured previously at the Homestake experiment). SNO used heavy water (D_2O , where the D nuclei are denoted by d in Equations 2.25 to 2.27 below) as its detection medium, and was therefore sensitive to a wider variety of neutrino interactions than other contemporary experiments that used standard H_2O . The three key interactions in SNO [41] were the charged-current (CC, mediated by the W^+ boson) interaction:



which can only occur through an incoming ν_e and not via ν_μ or ν_τ , the elastic scattering (ES) process:



which can occur via all neutrino flavours l , but with a higher sensitivity to ν_e than to ν_μ and ν_τ , and the neutral-current (NC, mediated by the Z^0 boson) interaction:

$$\nu_l + d \rightarrow p + n + \nu_l \quad (2.27)$$

which is equally sensitive to all neutrino flavours. The neutron from the NC interaction could be detected via a number of possible neutron capture mechanisms, which were varied throughout the lifetime of the experiment [30].

The ν_e flux calculated from the CC interaction was found to be significantly different from the total ν flux from the ES process [41]. However, the ES ν flux measured at SNO was found to be very consistent with the ν flux that had been previously measured at the Super-Kamiokande experiment using the same ES process for detection [42], indicating that the reduced ν_e flux seen at SNO was indeed real. The irrefutable evidence of neutrino flavour oscillations came from the total ν flux measured using the NC interaction [43] - this was shown to be consistent with the Standard Solar Model's predictions of the total (purely ν_e) neutrino output of the Sun. These two results therefore indicate that, although all solar neutrinos are indeed produced as ν_e , the majority of them ($\sim 65\%$) reach the Earth as either ν_μ or ν_τ , having changed flavour along the way.

The 2015 Nobel Prize in Physics was awarded to the leaders of the SuperK and SNO experiments - Takaaki Kajita and Arthur McDonald respectively - in recognition of the work performed by the two groups in showing that neutrinos undergo flavour oscillations.

2.3.3 Mathematical Framework and Implications

Consider that in quantum field theory, the neutrino flavour eigenstates $|\nu_\alpha\rangle$ ($\alpha = e, \mu$ or τ) can be written as superpositions of the neutrino mass eigenstates $|\nu_i\rangle$ ($i = 1, 2$ or 3) in the form:

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle \quad (2.28)$$

with a similar relation connecting the antineutrino flavour and mass eigenstates:

$$|\bar{\nu}_\alpha\rangle = \sum_i U_{\alpha i}^* |\bar{\nu}_i\rangle \quad (2.29)$$

In both of these equations, $U_{\alpha i}$ are the components of a unitary mixing matrix analogous to the CKM matrix for quarks (with $U_{\alpha i}^*$ denoting the component of the matrix's complex conjugate). For neutrinos, this is the *Pontecorvo-Maki-Nakagawa-Sakata* (PMNS) mixing matrix [44], U_{PMNS} :

$$U_{PMNS} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \quad (2.30)$$

This matrix is often parameterised in terms of three mixing angles: θ_{12} , θ_{13} and θ_{23} , which represent the Euler angles required to rotate the flavour eigenstates into the mass eigenstates, and a phase δ that represents any potential charge-parity (CP) violation in neutrino-mixing. Under this parameterisation, the PMNS matrix becomes:

$$U_{PMNS} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}s_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \quad (2.31)$$

where c_{ij} and s_{ij} ($i, j = 1, 2, 3$) are short-hand for $\cos(\theta_{ij})$ and $\sin(\theta_{ij})$ respectively.

Using the components of U_{PMNS} , the probability of a particular flavour eigenstate ν_α with energy E oscillating completely into another flavour eigenstate ν_β over a distance L is given by [30]:

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sum_i [U_{\alpha i} U_{\beta i}^*] - 4 \sum_{i>j=1}^3 \left[\text{Re}(U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right) \right] \\ + 2 \sum_{i>j=1}^3 \left[\text{Im}(U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}) \sin \left(\frac{\Delta m_{ij}^2 L}{2E} \right) \right] \quad (2.32)$$

where Δm_{ij}^2 is the *squared-mass difference* of mass eigenstates ν_i and ν_j :

$$\Delta m_{ij}^2 = m_i^2 - m_j^2 \quad (2.33)$$

(A derivation of this, starting from the eigenstate representation given in Equation 2.28, can be found in [30].)

It can be seen from Equation 2.32 that in order for $P(\nu_\alpha \rightarrow \nu_\beta) \neq 0$, there must be at least one $\Delta m_{ij}^2 \neq 0$ - i.e. at least one of the neutrino mass eigenstates must be non-zero, and therefore, the observation of neutrino flavour oscillations leads directly to the idea of massive neutrinos, in direct contradiction to the SM prediction (previously

discussed in Section 2.2.1) that neutrinos are massless.

In addition, neutrino flavour oscillations require that the leptonic family numbers no longer be conserved by weak interactions - once again, contradicting the SM. However, this is perhaps less radical an idea than massive neutrinos, since - as previously noted - the conservation of leptonic family numbers is only an assumption based on a lack of evidence in the contrary.

2.3.4 Current Measurements

Because:

$$\Delta m_{21}^2 + \Delta m_{32}^2 - \Delta m_{31}^2 = 0 \quad (2.34)$$

only two of the Δm_{ij}^2 terms are actually independent. In principle, it does not matter which two these are, but in practice Δm_{32}^2 and Δm_{21}^2 are chosen because they can be measured directly from atmospheric and solar neutrinos respectively. Atmospheric neutrinos allow the value of $|\Delta m_{32}^2|$ to be measured [45] (assuming that the neutrino masses are ordered as $m_1 < m_2 < m_3$ - which will be discussed in more detail in Section 2.3.5 below), and so $|\Delta m_{32}^2|$ is also denoted as Δm_{atm}^2 . The oscillations of solar neutrinos allows the value of Δm_{21}^2 to be measured, and therefore it is sometimes also written as Δm_{sol}^2 . Note that the sign of Δm_{21}^2 can be known from solar neutrino measurements, in contrast to measurements of Δm_{atm}^2 , which only yield the absolute value of Δm_{32}^2 .

Together with the three mixing angles θ_{12} , θ_{13} and θ_{23} , these two squared-mass differences form five of the six independent neutrino flavour oscillations parameters that can be measured by experiments. (The sixth parameter is the CP-violating phase δ featured alongside the mixing angles in the PMNS matrix.) Table 2.1 gives

the current best values of these five parameters, based on a global analysis of data gathered across a number of different experiments.

Quantity	Value	Hierarchy
Δm_{21}^2	$(7.53 \pm 0.18) \times 10^{-5} \text{ eV}^2$	Both
$ \Delta m_{32}^2 $	$(2.44 \pm 0.06) \times 10^{-3} \text{ eV}^2$	Normal
	$(2.52 \pm 0.07) \times 10^{-3} \text{ eV}^2$	Inverted
$\sin^2(2\theta_{12})$	0.846 ± 0.021	Both
$\sin^2(2\theta_{13})$	$(9.3 \pm 0.8) \times 10^{-2}$	Both
$\sin^2(2\theta_{23})$	$0.999^{+0.001}_{-0.018}$	Normal
	$1.000^{+0.000}_{-0.017}$	Inverted

Table 2.1: The current best measurements of the five independent parameters that describe neutrino flavour oscillations: the two square-mass differences Δm_{21}^2 and Δm_{32}^2 that feature in the oscillation probability equation (2.32), and the three mixing angles θ_{12} , θ_{13} and θ_{23} that parameterise the PMNS mixing matrix (Equation 2.31). All values are taken from [46]. The ‘‘Hierarchy’’ column refers to the neutrino mass hierarchies discussed in Section 2.3.5.

2.3.5 Mass Hierarchies

Since only the absolute value $|\Delta m_{32}^2|$ can be measured, it is not known if $m_3 > m_2$, or vica versa. This leads to two possible orderings of the neutrino masses, commonly known as the *normal* and *inverted hierarchies* (abbreviated to NH and IH respectively). These are depicted in Figure 2.1.

2.4 Beyond the SM: Massive Neutrinos

2.4.1 Introduction and Current Measurements

As previously discussed, it is clear from the experimental evidence of neutrino oscillations that neutrinos must in fact be massive, directly contradicting the SM. The question then becomes how neutrinos gain their masses, since - as noted in Section

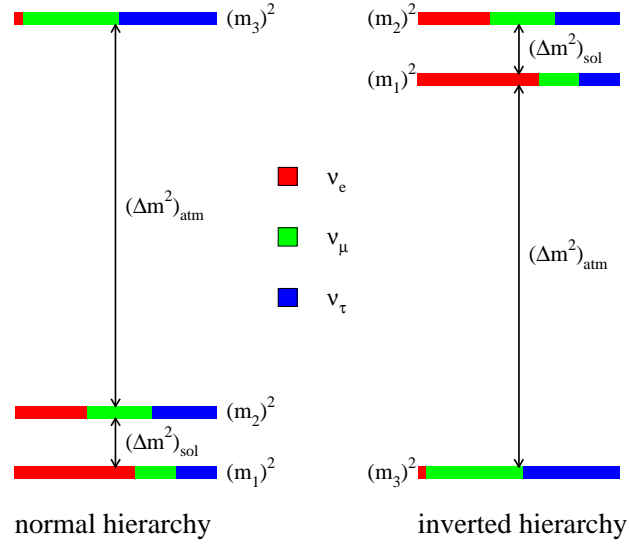


Figure 2.1: The two mass hierarchies that are possible given the current measurements of Δm_{21}^2 ($= \Delta m_{sol}^2$) and $|\Delta m_{32}^2|$ ($= \Delta m_{atm}^2$ in the normal hierarchy, and $\simeq |\Delta m_{31}^2| = \Delta m_{atm}^2$ in the inverted hierarchy, since $\Delta m_{sol}^2 \ll \Delta m_{atm}^2$). The colours represent the contributions to each mass eigenstate from the flavour eigenstates, and are equal to $|U_{\alpha i}|^2$ - for example, the red (ν_e) part of the m_1^2 bar is given by $|U_{e1}|^2$. This image was taken from [47].

2.2.1 - they cannot use the SM Higgs mechanism due to being purely left-handed particles.

Currently, the best limits on the ν_e mass come from measurements of the β decay of tritium (${}^3\text{H}$). This involves using a spectrometer to very accurately measure the decay's energy spectrum endpoint, since the neutrino mass is equal to the difference in energy between the spectrum endpoint and the decay's Q value of 18.5898keV [48]. Given the established limits on the squared-mass differences previously listed in Table 2.1, this is obviously not an easy task - such experiments are searching for a difference in energy on the order of just a few eV at the end of a multi-keV spectrum, but nevertheless, the current best limit on the effective mass of ν_e is [46]:

$$m_{\nu_e}^{eff} < 2\text{eV} \quad (2.35)$$

where:

$$m_{\nu_e}^{eff} = \sqrt{\sum_{i=1}^3 U_{ei}^2 m_i^2} \quad (2.36)$$

The limits on the effective ν_μ and ν_τ masses are much less strict, as they are calculated from much larger starting quantities and therefore lack the sensitivity of the low energy ${}^3\text{H}$ β decay measurements. The effective ν_μ mass $m_{\nu_\mu}^{eff}$ can be found using measurements of the π^\pm decay at rest, and results in [46]:

$$m_{\nu_\mu}^{eff} < 0.19\text{MeV} \quad (2.37)$$

while limits on $m_{\nu_\tau}^{eff}$ - the effective ν_τ mass - have only been set using data taken at LEP between 1991 and 1995 [46]:

$$m_{\nu_\tau}^{eff} < 18.2\text{MeV} \quad (2.38)$$

(Both $m_{\nu_\mu}^{eff}$ and $m_{\nu_\tau}^{eff}$ are defined analogously to $m_{\nu_e}^{eff}$ in Equation 2.36 above, but using $U_{\mu i}^2$ and $U_{\tau i}^2$ respectively in place of U_{ei}^2 .)

2.4.2 Dirac Neutrinos

It is possible to introduce a right-handed neutrino into the SM, with the caveat that it is *sterile* - that is, it does not interact via the weak, strong or electromagnetic forces, but only through gravity, thus explaining why such a particle has not been observed. In this case, neutrinos are still classed as *Dirac* particles - that is, a particle and its antiparticle are distinctly separate entities, as with all other fundamental fermions

- and they gain mass through the SM Higgs mechanism. The resulting Lagrangian mass term, analogous to Equation 2.12, is given by [45]:

$$\mathcal{L}_{mass,Dirac} = - \sum_{i=1}^3 m_{i,Dirac} (\overline{\nu_{iR}}\nu_{iL} + \overline{\nu_{iL}}\nu_{iR}) \quad (2.39)$$

with the individual neutrino masses taking the same form as Equation 2.13:

$$m_{i,Dirac} = \frac{c'_i v}{\sqrt{2}} \quad (2.40)$$

where v is once again the Higgs field's vacuum expectation value, ν_{iL} and ν_{iR} are the left- and right-handed components of the Dirac neutrino field ν_i ($= \nu_{iL} + \nu_{iR}$), $\overline{\nu_{ix}} = \nu_{ix}^+ \gamma^0$ is the adjoint chiral component and c'_i is the Yukawa coupling of each neutrino field i ($= 1, 2$ or 3) to the Higgs field.

However, this mechanism suffers from the same issue that it does under application to the other SM fermions - namely, that it results in a somewhat arbitrary set of c'_i that do not directly predict the particle masses. In addition, this method does not naturally explain why neutrinos have masses so much smaller than the other fundamental particles - i.e. why c'_i is many orders of magnitude smaller than the analogous couplings of the other particles.

2.4.3 Majorana Neutrinos

The left- (ψ_L) and right-handed (ψ_R) components of a general Dirac field ψ are usually treated as independent entities, but they can be related to each other through the *Majorana Relation*:

$$\psi_R = \xi \mathcal{C} \overline{\psi}_L^T = \xi \psi_L^C \quad (2.41)$$

where ξ is an arbitrary phase and \mathcal{C} is the charge-conjugation matrix. Fermionic particle fields that satisfy this relation must still of course satisfy the Dirac Equation that describes fermions, but with the added consequence that:

$$\psi = \psi^C \quad (2.42)$$

i.e. the charge-conjugated field is the same as the original field. In terms of physical particles, this corresponds to a particle being indistinguishable from its antiparticle - such particles are known as *Majorana* particles, after Ettore Majorana who originally proposed their existence [49]. A derivation of the above relations can be found in [45].

Being uniquely neutral out of all of the SM fermions, neutrinos are the only ones that can potentially be Majorana particles, with a Lagrangian mass term given by:

$$\mathcal{L}_{mass, Major} = - \sum_{i=1}^3 m_{i, Major} \left(\overline{\nu}_{iL}^C \nu_{iL} + \overline{\nu}_{iL} \nu_{iL}^C \right) \quad (2.43)$$

(analogous in form to that of a Dirac neutrino given in Equation 2.39), and individual masses of [45]:

$$m_{i, Major} = \frac{g_i v^2}{2\mathcal{M}} \quad (2.44)$$

where g_i are the dimensionless coupling constants of the neutrino mass eigenstates. The constant \mathcal{M} , which has dimensions of mass, is related to the scale at which Grand

Unified Theories become applicable ($\sim 10^{15}\text{GeV}$), and so Equation 2.44 provides a naturally small Majorana neutrino mass on the order of $\sim 10^{-2}\text{eV}$, without needing to introduce artificially small coupling constants or extra fundamental particles such as sterile right-handed neutrinos. (It should be noted that the idea of neutrinos being Majorana particles does not necessarily imply that they must be massive - if $g_i = 0$, then Majorana-type neutrinos may still be massless.)

One important additional detail to note is that if neutrinos are indeed Majorana particles, the total lepton number can no longer be a conserved quantity, since neutrinos and antineutrinos have different lepton numbers.

Majorana neutrinos also feature in certain *seesaw mechanisms* that can be used to generate very small neutrino masses. In these frameworks, the usual left-handed neutrino is accompanied by a massive right-handed sterile neutrino, but unlike the previously discussed simplest extension to the SM, the small left-handed neutrino mass in seesaw mechanisms is directly related to the large mass of the right-handed neutrino (hence the name “seesaw”: as the right-handed neutrino’s mass becomes larger, the left-handed neutrino’s mass becomes smaller). These mechanisms are of particular interest as they naturally provide a potential dark matter candidate particle in the massive sterile neutrino.

2.5 Double β Decay

2.5.1 Overview

As well as the aforementioned β decay, IBD and EC, another weak interaction that a nucleus can potentially undergo is *double β decay*. First proposed in 1935 by M. Goeppert-Mayer [50], this process involves the simultaneous β decays of two neutrons in the nucleus, resulting in an increase of 2 in the nuclear charge Z with no change

in the atomic number A , as well as four daughter particles that leave the nucleus: two e^- and two $\bar{\nu}_e$. The process is denoted by the shorthand of $2\nu\beta\beta$, and Figure 2.2 shows its quark-level Feynman diagram.




Figure 2.2: The Feynman diagram of double β decay when two e^- and two $\bar{\nu}_e$ are emitted ($2\nu\beta\beta$).

Double β decay can only occur in certain nuclei - specifically those where the mass of the original nucleus $m(Z, A)$ satisfies the following criteria [30]:

$$m(Z, A) > m(Z + 2, A) + 2m_e \quad (2.45)$$

$$m(Z, A) < m(Z + 1, A) + m_e \quad (2.46)$$

The first of these simply states that it is energetically possible for two β decays to occur simultaneously, and the second is the requirement that normal β decay is at the very least energetically unfavourable, if not outright impossible. The Q value of $2\nu\beta\beta$ - the total energy released in the reaction - is given (in natural units) by:

$$Q_{\beta\beta} = m(Z, A) - m(Z + 2, A) - 2m_e \quad (2.47)$$

but because there are four particles in the final state, the observable energy - carried by the two e^- - can take any value between 0 and $Q_{\beta\beta}$.

There are 35 nuclei that can theoretically undergo $2\nu\beta\beta$ [45], with $Q_{\beta\beta}$ ranging from 70.2keV to 3.35MeV. Experiments have observed $2\nu\beta\beta$ occurring in eleven of these, from which the various nucleus-specific half-lives of $2\nu\beta\beta$ have been calculated - varying between 6.8×10^{18} (^{100}Mo) and 2.0×10^{21} (^{238}U) years.

2.5.2 Neutrinoless Double β Decay

With the idea that neutrinos could potentially be Majorana particles ($\bar{\nu}_\alpha = \nu_\alpha$ for $\alpha = e, \mu$ or τ), it was proposed [51][52] that the $\bar{\nu}_e$ emitted by one of the β decays can be absorbed as a ν_e by the W boson originating from the other β decay, leaving only two e^- in the final state. This is known as *neutrinoless double β decay* ($0\nu\beta\beta$) and is depicted in Figure 2.3.




Figure 2.3: The Feynman diagram of neutrinoless double β decay when only two e^- are emitted ($0\nu\beta\beta$). The $\bar{\nu}_e$ emitted from the “top” β decay has been absorbed as a ν_e by the W boson emitted from the “bottom” d -quark decay.

As the difference between $2\nu\beta\beta$ and $0\nu\beta\beta$ only manifests once the decays have actually occurred, the two requirements for double β decay (Equations 2.45 and 2.46) still hold, and so all 35 nuclei that undergo $2\nu\beta\beta$ can in principle also perform $0\nu\beta\beta$. However, the kinematics of the final state of $0\nu\beta\beta$ are considerably different: since the e^- now carry all of the released energy (assuming negligible nuclear recoil), they now always have a total energy equal to $Q_{\beta\beta}$, which itself is the same as that given in Equation 2.47.

A more subtle requirement for $0\nu\beta\beta$ to occur, aside from the obvious assumption that neutrinos are Majorana particles, is that neutrinos *must* be massive - i.e. $g_i \neq 0$ in Equation 2.44. This is related to matching the helicities across the virtual neutrino: the particle that is absorbed by the second W boson must be a $\mathcal{H} = -1$ ν_e . Although the basic understanding would suggest that a $\bar{\nu}_e$ is emitted from the first W with $\mathcal{H} = +1$, and therefore a helicity flip is required, in fact there is a small probability of the first W emitting a $\mathcal{H} = -1$ ν_e - this probability is proportional to $\left(\frac{m_\nu}{E_\nu}\right)^2$ and therefore requires that $m_\nu \neq 0$ [45].

Effective Majorana Mass

As discussed in detail previously, the mediating $\nu_e/\bar{\nu}_e$ is a superposition of the mass eigenstates ν_i ($i = 1, 2$ or 3), and so has no definite mass itself. Each of the ν_i has a probability $\propto U_{ei}^2$ of interacting with the W boson, and together with the m_i terms required from helicity-matching as discussed above, this leads to the overall *effective Majorana mass*, $m_{\beta\beta}$, of the neutrino - that is, the mass of the virtual $\nu_e/\bar{\nu}_e$ that mediates $0\nu\beta\beta$ [45]:

$$m_{\beta\beta} = \sum_{i=1}^3 U_{ei}^2 m_i \quad (2.48)$$

(Note that this is *not* the same as the effective ν_e mass given previously in Equation

2.36.) This effective mass can be written in terms of the lightest neutrino mass in each of the hierarchies - m_1 in NH, or m_3 in IH:

$$m_{\beta\beta}^{NH} = U_{e1}^2 m_1 + U_{e2}^2 \sqrt{m_1^2 + \Delta m_{sol}^2} + U_{e3}^2 \sqrt{m_1^2 + \Delta m_{atm}^2} \quad (2.49)$$

$$m_{\beta\beta}^{IH} = (U_{e1}^2 + U_{e2}^2) \sqrt{m_3^2 + \Delta m_{atm}^2} + U_{e3}^2 m_3 \quad (2.50)$$

where $\Delta m_{sol}^2 \ll \Delta m_{atm}^2$ has been used for simplification. These relations are shown graphically in Figure 2.4.




Figure 2.4: The behaviour of the effective Majorana neutrino mass $m_{\beta\beta}$ (given in Equation 2.48) as a function of the lightest neutrino mass m_{min} ($= m_1$ in the normal hierarchy NH, and $= m_3$ in the inverted hierarchy IH). Also shown are the limits on $m_{\beta\beta}$ imposed by current $0\nu\beta\beta$ experiments, and the cosmological limit on m_{min} . This image was taken from [53].

$0\nu\beta\beta$ Decay Rate

The rate of the $0\nu\beta\beta$ process (equal to the inverse of the half-life $t_{1/2}^{0\nu}$) is in general given by [45]:

$$[t_{1/2}^{0\nu}]^{-1} = G^{0\nu} |\mathcal{M}^{0\nu}|^2 \frac{|m_{\beta\beta}|^2}{m_e^2} \quad (2.51)$$

where $G^{0\nu}$ is the phase space factor specific to a particular nucleus, $\mathcal{M}^{0\nu}$ is the nucleus' nuclear matrix element and $m_{\beta\beta}$ is the effective Majorana mass given in Equation 2.48.

The values of $G^{0\nu}$ have been calculated very precisely for several of the $0\nu\beta\beta$ -capable nuclei [54]. Unfortunately, the same cannot be said for $|\mathcal{M}^{0\nu}|$ - the value of this is much less well-known, even for a single nucleus - for example, the ^{76}Ge matrix element can take values between 1.50 and 4.59, depending on exactly which method is used [45], introducing a factor of almost 10 into the calculation of $[t_{1/2}^{0\nu}]^{-1}$ even when ignoring any uncertainties in $|m_{\beta\beta}|$. A summary of some of the nuclear models used for calculating $|\mathcal{M}^{0\nu}|$, together with a discussion of ways to improve the calculation, can be found in [55].

Current Measurements

A number of experiments have been performed over the past two decades that searched for $0\nu\beta\beta$ events in a variety of isotopes. As of 2015, no conclusive observation of $0\nu\beta\beta$ has been made, but lower limits on the half-life $t_{1/2}^{0\nu}$ can be set for the various isotopes - these are given in Table 2.2.

In 2006, an observation of $0\nu\beta\beta$ in ^{76}Ge was claimed by part of the Heidelberg-Moscow (HDM) collaboration [56], however this was disputed by a number of physicists. The GERDA experiment in 2013 also used ^{76}Ge , and found no evidence of $0\nu\beta\beta$ - disproving the Heidelberg-Moscow claim, and (together with previously taken

data from the HDM and IGEX collaborations) setting the current limit on the $0\nu\beta\beta$ half-life in ^{76}Ge . The second phase of GERDA - with improved detectors, electronics and background counters [57] - is currently undergoing commissioning.

The NEMO3 experiment, which ran from 2003 to 2011, was able to use six isotopes due to its modular design and set limits on the $0\nu\beta\beta$ half-life in all six, having observed no $0\nu\beta\beta$ events. The next generation of this experiment - SuperNEMO - will focus primarily on ^{82}Se , but will be ten times larger than NEMO-3, allowing it to potentially reach half-life limits on the order of 10^{26} years [58].

Two experiments - KamLAND-Zen and EXO-200 - have used ^{136}Xe as their potential $0\nu\beta\beta$ isotope. The former uses an isotope-loaded liquid-scintillator similar in concept to SNO+, while the latter uses enriched liquid xenon in a time projection chamber. (These are both in contrast to GERDA and NEMO-3, which used crystal-form isotopes.)

The CUORICINO experiment used ^{130}Te crystals as its $0\nu\beta\beta$ material - the same isotope that SNO+ will eventually use. Having taken data between 2003 and 2008, this experiment is the prototype stage of the eventual CUORE detector, which will have almost twenty times more ^{130}Te than CUORICINO operated with, and aims to reach a half-life limit on the order of 10^{26} years [59].

2.5.3 $0\nu\beta\beta$ at SNO+

As mentioned in Section 1.2.1, SNO+ will use ^{130}Te in its search for $0\nu\beta\beta$ during the loaded phase.

Natural Tellurium is a mixture of eight isotopes, of which ^{130}Te is the most abundant, with a relative fraction of 34.08% [12]. This, together with the relatively low price of bulk natural Tellurium, makes it reasonably easy and cheap to procure ^{130}Te in quantities large enough for a kilotonne-scale experiment such as SNO+. The ^{130}Te isotope itself has a $Q_{\beta\beta}$ of 2.526MeV [65], placing any potential $0\nu\beta\beta$ signal above

Isotope	$t_{1/2}^{0\nu}$ Lower Limit [years]	Experiment
^{48}Ca	1.3×10^{22}	NEMO3 [60]
^{76}Ge	3.0×10^{25}	HDM + IGEX + GERDA [61]
^{82}Se	2.1×10^{23}	NEMO3 [60]
^{96}Zr	8.6×10^{21}	NEMO3 [60]
^{100}Mo	1.1×10^{24}	NEMO3 [62]
^{116}Cd	1.6×10^{22}	NEMO3 [60]
^{130}Te	2.8×10^{24}	CUORICINO [63]
^{136}Xe	3.4×10^{25}	KamLAND-Zen + EXO-200 [64]
^{150}Nd	1.8×10^{22}	NEMO3 [60]

Table 2.2: The current lower limits on the half-life of $0\nu\beta\beta$ in a number of different isotopes, as measured by the respective experiments.

many of the lower energy background β decays that might be found in the experiment. (A more comprehensive discussion of the backgrounds to a potential $0\nu\beta\beta$ signal can be found in Section 6.1.1.) In addition, the $2\nu\beta\beta$ decay rate in ^{130}Te is considerably lower than some other potential $2\nu\beta\beta$ isotopes such as ^{150}Nd , making the identification of $0\nu\beta\beta$ potentially easier.

With an initial Phase I loading of 0.3% as discussed in Section 1.2.1, SNO+ can exclude the m_{min} vs. $m_{\beta\beta}$ parameter space above the IH within 3 to 5 years - a comparable result to CUORE, and SNO+ Phase II - with percent-level loading of ^{130}Te - can potentially allow the exploration of the entire IH region. However, current loading techniques and other limitations of the detector preclude easily reaching the NH region in the near future.

Chapter 3

Optical Scattering of Photons

‘Well, actually ...’

— Nazim Hussain, on (many) things that I thought I knew about, both in and outside Physics

3.1 General Principles

The scattering of a photon from a particle embedded in a medium is not simply the act of the photon “bouncing” off the particle. Instead, the scattering process is part of a more complex interaction between the photon’s intrinsic electromagnetic field and the particle’s electrons.

Consider a photon moving along the z -axis and linearly polarised in the x - y plane. The photon’s electric and magnetic fields are then given by:

$$\underline{E}(z, t) = E_0 \sin(k_0 z - \omega t) \underline{\hat{x}} \quad (3.1)$$

$$\underline{B}(z, t) = \frac{E_0}{c} \sin(k_0 z - \omega t) \underline{\hat{y}} \quad (3.2)$$

where k_0 is the photon's wavenumber, related to its wavelength in vacuum, λ_0 , by:

$$k_0 = \frac{2\pi}{\lambda_0} \quad (3.3)$$

This photon is incident on a single atom or molecule (hereafter called the *scattering particle*) of radius a and refractive index n_{par} , which is embedded in a bulk medium of refractive index n_{med} . (For simplicity, the scattering particle and bulk medium are both assumed to be non-absorbing, and therefore their respective refractive indices are real numbers. A more general situation which includes photon absorption would require complex refractive indices, but this is not relevant to this discussion of scattering.) The total light intensity incident on the scattering particle, I_0 , is equivalent to the power produced by the photon per unit area of the scattering particle, and is therefore given by the magnitude of the photon's time-averaged Poynting vector:

$$\begin{aligned} I_0 &= \frac{1}{\mu_0} |\langle \underline{E} \times \underline{B} \rangle| \\ &= \frac{(E_0)^2}{2\mu_0 c} \end{aligned} \quad (3.4)$$

The incident photon's electromagnetic field interacts with the scattering particle's orbital electrons, displacing them into an oscillating motion characterised by an angular frequency of ω - the same frequency as the photon's electromagnetic field. The electrons' oscillating motion leads to the scattering particle now possessing a periodic charge separation, known as an *induced dipole moment*. The induced dipole moment acts as a new source of electromagnetic radiation, emitting "scattered" light in an angular distribution that is dependent on the properties of both the incident photon and the scattering particle itself [66]. The majority of this scattered light is emitted

at the same frequency as the dipole's oscillation itself, i.e. at the same frequency and energy as the original incident photon - this is known as *elastic* scattering. A small fraction of the scattered light is *inelastic* - being emitted at a different frequency and energy to the incident photon. In this situation, the scattering particle is first excited to a higher energy state, thereby reducing the energy available to be carried by the subsequently emitted scattered photon. Inelastic scattering of photons is also known as *Raman scattering* [67], and will not be discussed further here.

3.2 Elastic Scattering of Photons

3.2.1 Overview

The elastic scattering of a single photon is most generally described by *Mie Theory*. Developed by Gustav Mie in 1908 [68], this is a general solution for the elastic scattering of photons from a perfectly spherical scattering particle of any size and absorbance. In this framework, the emission direction of the scattered light, \underline{k}' , is defined by two angles: θ and ϕ , as shown in Figure 3.1, and the *scattering plane* is defined as the plane that contains both \underline{k}' and \underline{k} - the direction of the incident photon. The polarisations of the incident and scattered photons are denoted by \underline{p} and \underline{p}' respectively.

Instead of defining the components of the incident photon's polarisation \underline{p} in the cartesian coordinate system, it is more convenient to consider the orthogonal components perpendicular and parallel to the scattering plane. There are then two extreme cases of incident polarisation: when \underline{p} is perpendicular to the scattering plane, and when \underline{p} is parallel to the scattering plane. Equations 3.5 and 3.6 respectively give the intensity of scattered light at some coordinates (r, θ, ϕ) with respect to the scattering particle in these two scenarios, formulated in terms of their respective *differential*




Figure 3.1: Diagram showing the notation used to describe the Mie scattering of photons. The incident photon travels in a direction \underline{k} (in this case, along the z -axis), and scatters within a scattering particle that is centred at the origin. The scattered photon is emitted in a direction \underline{k}' . The two scattering angles are: θ - the angle between \underline{k}' and \underline{k} , and ϕ - the angle between $\underline{k}'_{\text{perp}}$ (the projection of \underline{k}' onto the plane perpendicular to \underline{k} [in this case, the x - y plane]) and the incident photon's polarisation \underline{p} [in this case, the x -axis]. The scattered photon's polarisation \underline{p}' is not shown.

cross sections σ'_{\perp} and σ'_{\parallel} , which describe the power of scattered light in the direction defined by θ and ϕ .

$$I_{\perp}(r, \theta, \phi) = \frac{I_0}{r^2} \sigma'_{\perp}(\theta, \phi) \quad (3.5)$$

$$I_{\parallel}(r, \theta, \phi) = \frac{I_0}{r^2} \sigma'_{\parallel}(\theta, \phi) \quad (3.6)$$

Unpolarised incident light is an equal mixture of light polarised perpendicular and parallel to the scattering plane, and therefore the differential cross section of unpolarised light, σ'_{sca} , is simply the average of σ'_{\perp} and σ'_{\parallel} :

$$\sigma'_{sca}(\theta, \phi) = \frac{\sigma'_{\perp}(\theta, \phi) + \sigma'_{\parallel}(\theta, \phi)}{2} \quad (3.7)$$

and the intensity $I_{sca}(r, \theta, \phi)$ of unpolarised scattered light at a distance r from the scattering particle is:

$$I_{sca}(r, \theta, \phi) = \frac{I_0}{r^2} \sigma'_{sca}(\theta, \phi) \quad (3.8)$$

The exact forms of σ'_{\perp} and σ'_{\parallel} are dependent on the size of the scattering particle relative to the incident photon's wavelength - a comparison which can be quantitatively characterised using a dimensionless size parameter α , defined as [66]:

$$\alpha = \frac{2\pi a}{\lambda} \quad (3.9)$$

where λ is the incident photon's wavelength in the bulk medium, related to the wavelength in vacuum by:

$$\lambda = \frac{\lambda_0}{n_{med}} \quad (3.10)$$

When a is very small in comparison to the incident photon's wavelength both in vacuum (Equation 3.11) and within the scattering particle itself (Equation 3.12):

$$\alpha \ll 1 \quad (3.11)$$

$$n_{par}\alpha \ll 1 \quad (3.12)$$

the photon's electromagnetic field is effectively uniform across the entire diameter of

the particle, and the induced dipole moment is fully established within a time window shorter than the oscillation period of the photon's electromagnetic field [69]. When these two requirements, known as the *Rayleigh criteria*, are both met, the complicated Mie theory framework reduces to the much simpler *Rayleigh scattering* formulation of single photon elastic scattering [70] - a derivation of this limiting case starting from the general Mie framework can be found in [71].

At the other extreme, when a is much larger than the incident photon's wavelength in both vacuum and the scattering particle, the Mie solution converges to the standard calculations of geometric optics. Therefore, the phrase "Mie scattering" is generally used to describe the middle-ground situation, where a is comparable to the wavelength of the incident photon.

3.2.2 Mie Scattering

The differential cross sections σ'_{\perp} and σ'_{\parallel} for general Mie scattering are given by [66]:

$$\sigma'_{\perp}(\theta, \phi) = \frac{\lambda^2}{4\pi^2} i_1(\theta, \phi) \quad (3.13)$$

$$\sigma'_{\parallel}(\theta, \phi) = \frac{\lambda^2}{4\pi^2} i_2(\theta, \phi) \quad (3.14)$$

where $i_1(\theta, \phi)$ and $i_2(\theta, \phi)$ are angular intensity functions based on infinite series of the Legendre polynomials. The differential cross section of unpolarised Mie scattered light, based on combining these with Equation 3.7, is then:

$$\sigma'_{sca}(\theta, \phi) = \frac{\lambda^2}{8\pi^2} [i_1(\theta, \phi) + i_2(\theta, \phi)] \quad (3.15)$$

and is therefore also an infinite series.

The calculation of the angular intensity functions is a non-trivial task, and is nowadays achieved using purpose-designed computing algorithms, such as MIEV0 and MIEV1 [72], SuperMidi [73] and TUMiescat [74] among others.

3.2.3 Rayleigh Scattering I - Electrodynamical Approach

As noted previously, Rayleigh scattering is the limiting case of single photon scattering from a particle that is very small in comparison to the photon's wavelength. However, Rayleigh's original formulation was made in terms of fundamental electrodynamics, and only for an idealised collection of non-interacting particles, rather than as a special case of the general Mie framework.

A full derivation of Equations 3.16, 3.17 and 3.20, starting from basic electrodynamical principles, is provided in Appendix A.

The differential cross section $\sigma'_{sca}(\theta, \phi)$ and intensity $I_{sca}(r, \theta, \phi)$ of unpolarised light scattering from isolated, single particles are given in Equations 3.16 and 3.17 respectively.

$$\sigma'_{sca}(\theta, \phi) = \frac{8\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 (1 + \cos^2\theta) \quad (3.16)$$

$$I_{sca}(r, \theta, \phi) = \frac{I_0}{r^2} \frac{8\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 (1 + \cos^2\theta) \quad (3.17)$$

The *Rayleigh scattering length* l_{Ray} , equivalent to the mean value of d - the distance that a photon travels before undergoing Rayleigh scattering [75] - is defined by:

$$\begin{aligned}
l_{Ray} &= \langle d \rangle \\
&= \left(\iint \beta(r, \theta, \phi) d\Omega \right)^{-1}
\end{aligned} \tag{3.18}$$

where $\beta(r, \theta, \phi)$ is the *volume scattering function*:

$$\beta(r, \theta, \phi) = \frac{I_{sca}(r, \theta, \phi)}{I_0} \frac{r^2}{V} \tag{3.19}$$

i.e. the power of scattered light at a distance r and angle θ from the scattering particle per unit volume of the scattering medium per unit incident intensity [76]. The *Rayleigh ratio* R is defined as the volume scattering function at a scattering angle of $\theta = 90^\circ$ [77]. These together lead to the explicit Rayleigh scattering length for unpolarised light scattering from isolated, single particles:

$$l_{Ray} = \left(\frac{128\pi^5 a^6}{3\lambda^4} \frac{1}{V} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 \right)^{-1} \tag{3.20}$$

Real-world materials do not exist as a simple collection of single, isolated particles - the constituent atoms and/or molecules are related through their inter-particle forces. Therefore, attempting to formulate the elastic scattering of photons in a bulk material using only the most basic Rayleigh scattering framework described above will inevitably lead to discrepancies.

One example of this is the polarisation of light emitted at $\theta = 90^\circ$: Equation A.7 indicates that at this angle there should be no scattering of light with \underline{p} parallel to the scattering plane, and therefore all light at this angle should be polarised perpendicular to the scattering plane. However, this has been shown not to be true, even in gases [78], which possess only weak inter-molecular forces.

The main difference between single particles and bulk materials is in the electric polarisability, noted in Equation A.1. For a single particle, this is a scalar quantity, and the entire material is optically *isotropic*. However, in a bulk material the polarisability is a vector, and the material is now said to be *anisotropic*. This anisotropy changes $\sigma'_{\parallel}(\theta, \phi)$ (but not $\sigma'_{\perp}(\theta, \phi)$) by a factor called the *degree of polarisation* [77]:

$$\text{degree of polarisation} = \frac{1 - \delta}{1 + \delta} \quad (3.21)$$

where the quantity δ is known as the *depolarisation ratio*, and is defined as:

$$\delta = \frac{I_{\parallel}(\theta = 90^{\circ})}{I_{\perp}(\theta = 90^{\circ})} \quad (3.22)$$

The change in $\sigma'_{\parallel}(\theta, \phi)$ then modifies the previously derived single-particle differential cross section $\sigma'_{sca}(\theta, \phi)$, intensity $I_{sca}(r, \theta, \phi)$ and scattering length l_{Ray} :

$$\sigma'_{sca}(\theta, \phi) = \frac{8\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 \left(1 + \frac{1 - \delta}{1 + \delta} \cos^2 \theta \right) \quad (3.23)$$

$$I_{sca}(r, \theta, \phi) = \frac{I_0}{r^2} \frac{8\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 \left(1 + \frac{1 - \delta}{1 + \delta} \cos^2 \theta \right) \quad (3.24)$$

$$l_{Ray} = \left(\frac{128\pi^5 a^6}{3\lambda^4} \frac{1}{V} \left[\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right]^2 \frac{2 + \delta}{1 + \delta} \right)^{-1} \quad (3.25)$$

3.2.4 Rayleigh Scattering II - Density Fluctuation Approach

The molecules in any medium are in constant thermal motion, and this gives rise to short-lived random fluctuations in the local mass density. These in turn create localised charge density fluctuations, and optical scattering can be thought of as occurring from interactions between photons and these fluctuations, rather than between photons and individual molecules as in the electro-dynamical approach. The material's dielectric constant ε , which is a measure of the ability of the material to resist the formation of an internal electric field, will change due to the density fluctuations, and *Einstein-Smoluchowski Theory* [79][80] describes the relation between the change in ε and the Rayleigh ratio R_{iso} of a perfectly optically isotropic medium:

$$R_{iso} = \frac{\pi^2}{2\lambda^4} \left[\rho \left(\frac{\partial \varepsilon}{\partial \rho} \right)_T \right]^2 k_B T \kappa_T \quad (3.26)$$

where ρ is the bulk density of the medium, $\left(\frac{\partial \varepsilon}{\partial \rho} \right)_T$ is the change in ε caused by a given change in ρ at a constant temperature T , k_B is the Boltzmann Constant, and κ_T is the isothermal compressibility of the medium - a measure of the relative change in volume for a given change in pressure at constant temperature. Numerous experiments [75] have shown that the Eykman Equation [81]:

$$\rho \left(\frac{\partial \varepsilon}{\partial \rho} \right)_T = \frac{(n_{med}^2 - 1)(2n_{med}^2 + 0.8n_{med})}{n_{med}^2 + 0.8n_{med} + 1} \quad (3.27)$$

is very accurate as a relation between the changes in ε caused by density fluctuations in the medium and the medium's refractive index n_{med} .

As noted previously, bulk materials are generally anisotropic due to inter-molecular forces, and so Equation 3.26 is not strictly applicable. However, the Rayleigh ratio

R of an anisotropic medium can be related to that of an isotropic one, R_{iso} , by the *Cabannes factor* [82]:

$$R = \frac{6 + 6\delta}{6 - 7\delta} R_{iso} \quad (3.28)$$

where δ is the same depolarisation ratio defined in Equation 3.22.

Combining Equations 3.26, 3.27 and 3.28 leads to an explicit expression for R in an anisotropic medium, and combining this with Equations A.11 and 3.24 (with $\theta = 90^\circ$) leads to the corresponding volume scattering function. Substituting this into Equation A.10 leads to the *Einstein-Smoluchowski-Cabannes formula* for the Rayleigh scattering length l_{Ray} of photons of wavelength λ in an anisotropic medium of refractive index n_{med} , isothermal compressibility κ_T and depolarisation ratio δ at a temperature T :

$$l_{Ray} = \left(\frac{8\pi^3}{3\lambda^4} \left[\frac{(n_{med}^2 - 1)(2n_{med}^2 + 0.8n_{med})}{n_{med}^2 + 0.8n_{med} + 1} \right]^2 k_B T \kappa_T \frac{6 + 3\delta}{6 - 7\delta} \right)^{-1} \quad (3.29)$$

The differential cross section $\sigma'_{sca}(\theta, \phi)$ and intensity $I_{sca}(r, \theta, \phi)$ are still given by Equations 3.23 and 3.24 respectively.

3.3 Current Measurements

The optical scattering properties of pure liquid water have been measured numerous times over the past decades, with all experiments indicating that Rayleigh scattering is by far the dominant process. These experiments give a depolarisation ratio of $\delta \simeq 0.09$ at $\lambda = 436\text{nm}$ and $\delta \simeq 0.10$ at $\lambda = 546\text{nm}$ [77]. A different measurement

[83] reports a value of $\delta \simeq 0.05$ at a higher wavelength of $\lambda = 633\text{nm}$. (Unfortunately, no errors are available for any of these values.)

Previous studies [84] of the optical scattering of photons in LAB have indicated that the dominant process is also Rayleigh scattering, with no more than a negligible contribution from Mie scattering. More recently, the JUNO experiment has for the first time measured the depolarisation ratio of pure LAB at a single wavelength of 430nm, finding it to be $\delta = 0.31 \pm 0.01(\text{stat.}) \pm 0.01(\text{sys.})$ [85], with a resulting scattering length at this wavelength of $l_{\text{Ray}} = 28.2 \pm 1.0\text{m}$.

3.4 Simulation of Optical Scattering in RAT

3.4.1 Pure Water

The values of l_{Ray} used in RAT for the simulation of pure water were originally derived from SNO calibration data, and then imported for use in SNO+. These values, covering a wavelength range of $200 \leq \lambda \leq 800\text{nm}$, are shown in Figure 3.2, along with values calculated using Equation 3.29. In the latter case, due to the small number of available measurements, δ has been taken as a constant 0.09 for $\lambda < 436\text{nm}$, linearly interpolated between 0.09 and 0.10 for $436 \leq \lambda < 546\text{nm}$, linearly interpolated between 0.10 and 0.05 for $546 \leq \lambda < 633\text{nm}$, and as a constant 0.05 for $\lambda \geq 633\text{nm}$. The refractive indices in the given wavelength range are already present in RAT, and a temperature of 285K (12°C) has been used - this being the expected operating temperature of SNO+. At this temperature, the isothermal compressibility of water is $4.73474 \times 10^{-10} \text{Pa}^{-1}$ [86].

There are some clear differences between the values of l_{Ray} used in RAT and those calculated using Equation 3.29, particularly for wavelengths between 350nm and 650nm. As noted above, the RAT values were originally obtained from SNO calibration data,

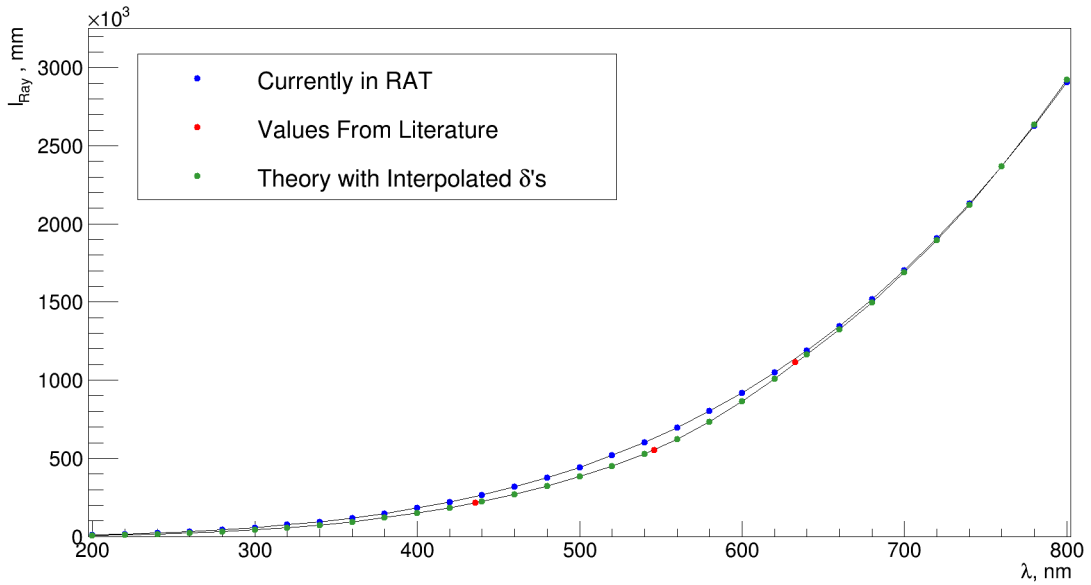


Figure 3.2: Values of the Rayleigh scattering length l_{Ray} of pure water at different wavelengths, originally derived from SNO's optical calibration data (blue). These values are now used in RAT for the simulation of scattering in pure water. Also shown are the values of l_{Ray} corresponding to the explicitly measured values of δ discussed in Section 3.3 (red), and theoretical values of l_{Ray} calculated using Equation 3.29 (green), with the various equation parameters discussed in the main text.

but were not themselves explicitly measured. This is because the individual contributions to the extinction of light from absorption and scattering could not be separated in SNO, and so only the total extinction length (the combination of absorption and scattering lengths) was measured at each wavelength. The scattering lengths were then inferred from these using the absorption spectrum of light in water. It seems likely then that some discrepancies are to be expected, as they have propagated through from any errors in the absorption model and extinction length measurements.

It should be noted however that, regardless of the source of these differences, they will not significantly impact any analysis of l_{Ray} in water, simply because at the wavelengths where the SNO+ PMTs are most efficient, both the theoretical and RAT values of l_{Ray} are much longer than the diameter of the SNO+ detector.

As noted previously in Equation A.10, l_{Ray} is the mean value of d - the distance that a photon travels before undergoing Rayleigh scattering, which itself should follow an exponential distribution. Figure 3.3 shows the distributions of d for three scenarios: photons of a single wavelength initially located at coordinates $(-7000, 0, 0)$ [in mm] and travelling towards positive x , photons initially located at coordinates $(0, 0, -7000)$ and travelling towards positive z , and photons initially located at coordinates $(6945, 2333, 4093)$ and travelling in a direction given by the unit vector $(-0.851, -0.276, -0.447)$. This third scenario is denoted as “off-axis”, and corresponds to a zero-width beam emitted from the FS025 optical fibre of the SMELLIE Calibration System discussed in Chapter 4.




Figure 3.3: The distance d travelled by a photon before it undergoes Rayleigh scattering in pure water, for three scenarios: photons travelling along the positive x axis (red), along the positive z axis (green), and emitted from the SMELLIE optical fibre FS025 as detailed in the main text (blue). Only photons that scattered within the detection medium have been included, and the distributions have been scaled so that they have equal areas in the range $0 \leq d \leq 12000$ mm.

It can be seen that the distribution of d is consistent across all three scenarios, indicat-

ing as expected that the initial direction of the photons does not affect the scattering. However, these distributions are all truncated by the physical limits of the detector: photons that scatter within the detection medium cannot have travelled for more than 12m (this being the diameter of the AV). The exception to this is for photons travelling in the positive z direction - due to the presence of the AV neck at high z , these photons can travel up to 19m before scattering. In principle, l_{Ray} could be extracted by simply fitting an exponential function to the distributions shown in Figure 3.3 and then calculating the mean value from each exponential's decay constant, but in practice the shallow slopes and truncation of the distributions makes the exponential fitting very unreliable.

Figures 3.4 and 3.5 show the normalised distribution of photons as a function of the scattering angle ϕ and the cosine of the scattering angle θ respectively, for the same three scenarios shown in Figure 3.3. It should be noted that ϕ cannot in fact be measured at SNO+, because (as previously shown on Figure 3.1) this angle is dependent on knowing the individual polarisation of each incident photon - something which cannot be achieved with the existing hardware. However, ϕ has still been discussed here, in order to provide a fuller picture of whether the RAT simulation of Rayleigh scattering as a whole is consistent with theory or not.

From comparing Figure 3.4 to Equation 3.24, it is clear that ϕ behaves as expected: since there is no dependence of intensity on ϕ , the distribution should be - and is - flat, regardless of the initial photon direction.

The distributions in Figure 3.5 indicate that, at least qualitatively, the correct behaviour is also present for θ : all three scenarios follow a $[1 + A \cos^2(\theta)]$ distribution, where A corresponds to the degree of polarisation previously defined in Equation 3.21. However, a quantitative examination of the fit function of each scenario results in $\delta = -0.00322 \pm 0.00420$ (x -axis scenario), 0.00782 ± 0.00353 (z -axis scenario) and




Figure 3.4: The Rayleigh scattering angle ϕ in pure water, for three scenarios: photons travelling along the positive x axis (red), along the positive z axis (green), and emitted from the SMELLIE optical fibre FS025 as detailed in the main text (blue). Also shown for reference is a flat fit function (black).




Figure 3.5: The cosine of the Rayleigh scattering angle θ in pure water, for three scenarios: photons travelling along the positive x axis (red), along the positive z axis (green), and emitted from the SMELLIE optical fibre FS025 as detailed in the main text (blue). Also shown is a fit function (black) of the form: $[1 + A \cos^2(\theta)]$, where A corresponds to the degree of polarisation previously defined in Equation 3.21.

0.00814 ± 0.00433 (off-axis scenario). (The errors on these δ values are comparatively large as a direct result of the δ 's themselves being so small.) All of these calculated δ 's are an order of magnitude smaller than the measured values discussed in Section 3.3, and indeed the value calculated for the x -axis scenario is unphysical - Equation 3.22 indicates that δ cannot be negative. In fact, these values of δ are more consistent with a simulated value of 0, possibly indicating that RAT does not account for the anisotropy of bulk materials.

3.4.2 LABPPO

The values of l_{Ray} that are used in RAT for LABPPO were measured specifically for SNO+ [87] at three institutions: Queen's University (Kingston, Canada), the University of Washington (Seattle, USA) and Brookhaven National Laboratory (New York, USA), all prior to 2012. These measurements, covering a wavelength range of $200 \leq \lambda \leq 600\text{nm}$ and shown in Figure 3.6, were found to be in agreement with each other as well as with the later JUNO measurement, and all indicate that photons undergo only Rayleigh scattering in LABPPO.

Figures 3.7, 3.8 and 3.9 show the distributions of d , ϕ and $\cos(\theta)$ respectively for the same three scenarios previously used for pure water.

As expected, the distribution of d is once again a truncated exponential in all three scenarios, but because l_{Ray} is shorter in LABPPO than in pure water, the exponential shapes are clearly steeper here than in Figure 3.3. However, the truncation still creates a high level of unreliability when attempting to extract the absolute value of l_{Ray} .

The distribution of ϕ is once again flat regardless of the initial photon direction, and $\cos(\theta)$ follows the accepted $[1 + A \cos^2(\theta)]$ shape as well. However, extracting absolute values of A from the three distributions yields corresponding δ values of -0.00139 ± 0.00271 (x -axis), -0.00112 ± 0.00244 (z -axis) and -0.00067 ± 0.00276




Figure 3.6: Measurements (blue) of the Rayleigh scattering length l_{Ray} of LABPPO at different wavelengths, originally measured at Queens University for SNO+ [87]. Also shown is the recent measurement of the Rayleigh scattering length at $\lambda = 430\text{nm}$ by the JUNO experiment [85] (red), which is in agreement with the measurements from Queen's.




Figure 3.7: The distance d travelled by a photon before it undergoes Rayleigh scattering in LABPPO, for three scenarios: photons travelling along the positive x axis (red), along the positive z axis (green), and emitted from the SMELLIE optical fibre FS025 as detailed in the main text (blue). Only photons that scattered within the detection medium have been included, and the distributions have been scaled so that they have equal areas in the range $0 \leq d \leq 12000\text{mm}$.




Figure 3.8: The Rayleigh scattering angle ϕ in LABPPO, for three scenarios: photons travelling along the positive x axis (red), along the positive z axis (green), and emitted from the SMELLIE optical fibre FS025 as detailed in the main text (blue). Also shown for reference is a flat fit function (black).




Figure 3.9: The cosine of the Rayleigh scattering angle θ in LABPPO, for three scenarios: photons travelling along the positive x axis (red), along the positive z axis (green), and emitted from the SMELLIE optical fibre FS025 as detailed in the main text (blue). Also shown is a fit function (black) of the form: $[1 + A \cos^2(\theta)]$, where A corresponds to the degree of polarisation previously defined in Equation 3.21.

(off-axis), all of which are unphysical and once again consistent with a simulated value of $\delta = 0$, lending further weight to the conclusion that RAT effectively treats all materials as optically isotropic ideal gases. (Once again, the error on each δ is large due to the small value of δ itself.)

Chapter 4

The SMELLIE Calibration System

‘So stupid.’

— Kate Pachal, on lasers (amongst other things)

4.1 The Scattering Module of ELLIE

4.1.1 Overview

Because the previously discussed scattering properties of light in water are known very precisely, they can be used as a benchmark calibration during the water phase of SNO+ - i.e. if the scattering properties of the water within the detector can be accurately measured in some way, any discrepancies from the known values could indicate potential issues with the water and/or detector, such as contamination from impurities. The same measurement procedure, once it has been shown to reliably work in water, can then be used to extract the scattering properties of LABPPO and Te_LABPPO during the scintillator and loaded phases of SNO+.

The system that has been developed to measure the detection medium’s scattering properties is the *Scattering Module* of the *Embedded LED/Laser Light Injection Entity* (SMELLIE). This is just one sub-system of the larger ELLIE calibration system,

which also includes sub-systems for measuring the SNO+ PMT timing (TELLIE) and detection medium's attenuation properties (AMELLIE).

The basic principle of SMELLIE is that it acts as a light-source with well-known properties - wavelength, position and direction - which is important, since (as discussed previously in Chapter 3) the angular distribution of scattered light in terms of θ is taken with respect to the incoming light's direction, and both the amount of scattering (related to the scattering length l_{Ray}) and the angular distribution are heavily dependent on the wavelength. In principle, once light from SMELLIE has been injected into the detector and observed by the PMTs, offline analysis will then be used to isolate the light that has been Rayleigh-scattered in the detection medium from light that underwent other optical processes, and the various properties of this isolated scattered light can then be calculated, allowing a determination of the bulk scattering properties of the detection medium to be made. (The offline analysis of SMELLIE light will be detailed in Chapter 5.)

SMELLIE itself consists of the hardware used to generate and inject light into the detector, together with the necessary software for controlling specific hardware components. The ex-situ hardware, shown schematically in Figure 4.1, is centred around four pulsed-diode laser heads operating at nominal wavelengths of 375, 405, 440 and 500nm, where the use of multiple wavelengths allows the wavelength dependence of the scattering to be investigated.

The full name of the ELLIE acronym would seem to indicate that either LEDs or laser heads can be used as light sources, however in practice each sub-system uses a light-source tailored for its own purposes. In the case of SMELLIE, laser heads offer several advantages over LEDs, such as inherently coherent output light that allows for a higher intensity and more effective collimation of the beam, as well as a much narrower wavelength distribution compared to an LED operating at the same nominal

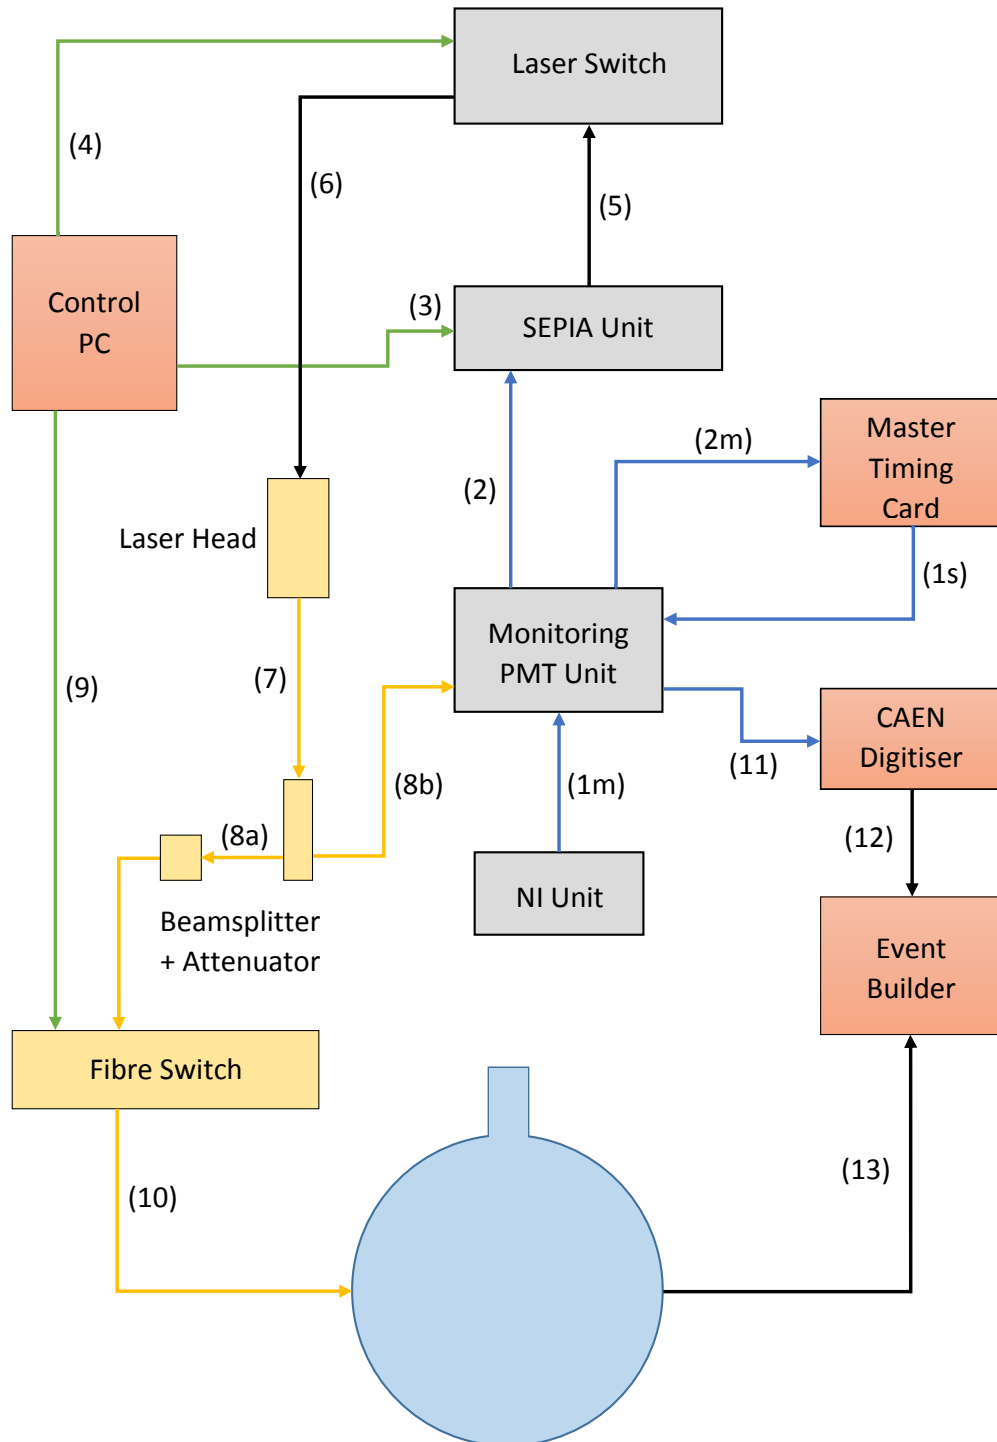


Figure 4.1: Schematic of the ex-situ SMELLIE hardware, as well as the various connections between them. Components that are not located on the ELLIE electronics rack are shown in orange, and those found on the rack are sub-divided into optical (yellow) and non-optical (grey) hardware. For clarity, only one laser head + beamsplitter + attenuator combination is shown. Inter-component connections are numbered from (1) to (13), with details given in Table 4.1.

Connection	Source	Signal Details	Destination	Comments
1m	NI Unit	TTL (0 to +5V)	Monitoring PMT Unit	Primary Signal (master mode)
1s	SNO+ Master Timing Card	ECL (-1.75 to -0.9V)	Monitoring PMT Unit	Primary Signal (slave mode)
2	Monitoring PMT Unit	Square pulse (-2 to +1V)	SEPIA Unit	Laser head Triggering Pulse
2m	Monitoring PMT Unit	ECL (-1.75 to -0.9V)	Master Timing Card	Only used in master mode
3	Control PC	USB	SEPIA Unit	Laser head intensity and repetition rate
4	Control PC	USB	Laser Switch	Laser Switch channel
5	SEPIA Unit	LEMO	Laser Switch	Laser head control signals
6	Laser Switch	LEMO	Laser head	Laser head control signals
7	Laser head	Optical fibre	Beamsplitter	
8a	Beamsplitter	Optical fibre	Fibre Switch (via Attenuator)	Low-fraction beamsplitter output
8b	Beamsplitter	Optical fibre	Monitoring PMT Unit	High-fraction beamsplitter output
9	Control PC	RS232	Fibre Switch	Fibre Switch channel
10	Fibre Switch	Optical fibre	In-situ Fibre Plate	
11	Monitoring PMT Unit	Sampled PMT Voltage pulse	CAEN Digitiser	
12	SNO+ CAEN Digitiser	Single voltage value	SNO+ Event Builder	
13	Detector	PMT readouts	SNO+ Event Builder	

Table 4.1: The connections between various ex-situ SMELLIE hardware components, as labelled on Figure 4.1. Further details on the master and slave operating modes are given in Section 4.1.2.

mean wavelength. More effective collimation is particularly important for SMELLIE, as the resulting narrow beam allows for significantly easier identification of light that has scattered away from the forward beam direction than a poorly collimated and wider beam would. In contrast, measurements of PMT timing as conducted by TELLIE require the near-simultaneous illumination of many PMTs across the PSUP, and therefore the wider beams created by the TELLIE LEDs are better suited to this task than the SMELLIE laser heads.

Each laser head's intensity, repetition rate and triggering are controlled through the SEPIA II Laser Driver Unit. However, the SEPIA Unit can only issue controls to a single laser head at a time, and so the Laser Switch is used to change control between laser heads in an automated fashion. Once the laser head has been fired, light travels into the Fibre Switch, where it is routed into one of twelve optical fibres that then enter the detector. In addition, a fraction of the laser head's light is routed into an ex-situ Monitoring PMT, which is used to measure the per-pulse performance of the laser heads in real time.

The SMELLIE system as a whole is controlled from a single source: the SMELLIE Control PC, which contains all of the necessary software for running each individual component remotely, as well as a general SMELLIE user interface that allows an operator to specify settings for all components and then run the entire system.

4.1.2 Operating Modes

SMELLIE is designed to be run in one of two possible operating modes: *master* and *slave*. The fundamental difference between these modes is the source of the signal that is used to activate SMELLIE - this will be referred to as the *primary signal* hereafter.

In master mode, the primary signal is provided by the National Instruments DAQ Unit (abbreviated to *NI Unit*) within SMELLIE itself. This device is capable of

performing many functions, more information about which can be found in [88], however within SMELLIE it is only used to produce a 0 to +5V TTL pulse as the primary signal in master mode. The choice of this signal is due to the requirements of various electronic components in the Monitoring PMT Unit further along in the system. In principle, the NI Unit is controlled by the SMELLIE Control PC via USB connection, however as noted it has only a single purpose within SMELLIE, and therefore very little actual control is required during operation. The primary signal from the NI Unit is simultaneously sent further into SMELLIE and to the SNO+ Master Trigger Card (abbreviated to *MTC*, corresponding to connections (1m) and (2m) respectively on Figure 4.1, and discussed further in Section 4.2.6), i.e. SMELLIE tells SNO+ when the laser head will be fired.

When operating SMELLIE in slave mode, the primary signal is sent from the MTC to SMELLIE (connection (1s) on Figure 4.1), i.e. SNO+ tells SMELLIE when to fire the laser head.

The choice of operating mode also has an impact on the various timing delays that must be introduced into the detector readout so that all of the information pertaining to a single SMELLIE event is collected and labelled as such. The timing delays between components are shown in Figures 4.2 and 4.3, for master and slave modes respectively.

In both operating modes, the first important time is known as the *coarse delay* - this is between the MTC issuing/receiving the primary signal, and issuing the global trigger (GT) to read out the detector's PMT information. The value of the coarse delay must be set such that the GT is issued just before light enters the detector - in master mode, this occurs 145ns after the primary signal is issued, but in slave mode this takes a longer time of 295ns. Taking into account the fixed cable delay of 110ns (in both modes) between the detector and the readout electronics, this results in a

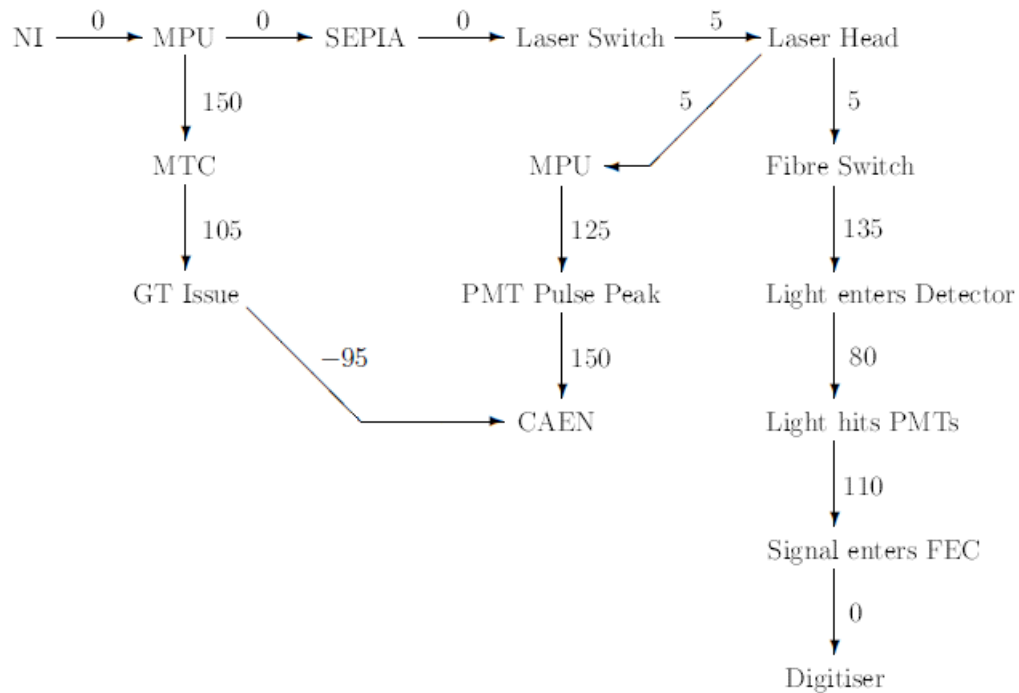


Figure 4.2: The timing delays (in ns) between the different components of SMELLIE, when operating in master mode. The primary signal is produced by the NI Unit, and sent to the MPU. From there, the signal is branched to the MTC - which then issues a global trigger (GT), and to SEPIA - which goes on to trigger the laser heads and send light into the detector.

coarse delay of $(145 + 110 - 150) = 105\text{ns}$ in master mode, and $(295 + 110) = 405\text{ns}$ in slave mode. The GT is therefore issued 255ns (master mode) or 405ns (slave mode) after the primary signal is sent.

The second variable time is the delay (referred to as the *CAEN delay*) between the issuing of the GT and the CAEN ADC beginning to record information for a single event. For the purposes of SMELLIE, CAEN must begin recording just before light enters the MPU from the laser head (once again, also taking into account the 150ns delay due to the cable between the MPU and CAEN). In both modes, the CAEN delay is -95ns , i.e. CAEN must begin recording 95ns before the GT is issued. (In master mode, the CAEN is calculated by: $(10 + 150 - 255) = -95$, and in slave mode: $(160 + 150 - 405) = -95$.)

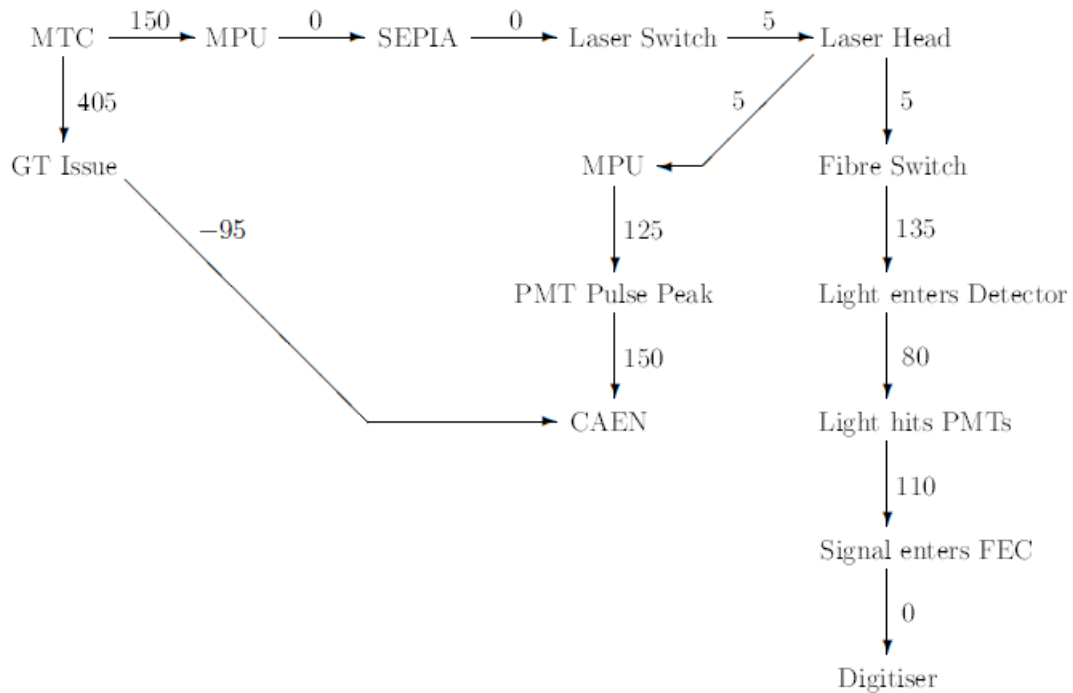


Figure 4.3: The timing delays (in ns) between the different components of SMELLIE, when operating in slave mode. The primary signal is produced by the MTC, and sent to the MPU. From there, the signal is sent to SEPIA, which goes on to trigger the laser heads and send light into the detector. The MTC also issues a global trigger (GT) at some time after sending the primary signal.

4.2 The SMELLIE Hardware

4.2.1 Laser Heads

The light-source of the SMELLIE system is a set of four pulsed-diode laser heads, manufactured by PicoQuant. This section will provide details about the laser heads as specifically related to their operation within SMELLIE, but more general information and specifications can be found at [89].

Figure 4.4 shows the four laser heads, and Table 4.2 gives a summary of their main specifications. All four produce a collimated output beam, with variable pulse intensity and a maximum possible repetition rate (the number of pulses produced per second) of 40MHz, both of which are controlled by the operator through the SEPIA




Figure 4.4: The four SMELLIE laser heads, arranged in increasing wavelength order from back (375nm) to front (500nm). The 405nm and 440nm laser heads have been refurbished, hence their differing external appearance from the remaining two. All four laser heads are shown coupled to orange-sheathed optical fibres.

II Laser Driver (discussed in Section 4.2.2).

Nominal Average Wavelength (nm)	375	405	440	500
Measured Average Wavelength (nm)	375.4	406.6	446.0	495.3
Average Power (mW) at 40MHz Repetition Rate	0.87	1.30	0.56	0.51
Single Pulse FWHM (ps) at 100% Intensity	54	49	63	132
Single Pulse Maximum Total Energy (pJ)	22	33	14	13

Table 4.2: Summary of the main specifications of the SMELLIE laser heads. All measurements were performed by PicoQuant and provided in documentation.

It can be seen that there are differences between the nominal and measured average wavelengths of all laser heads, with the nominally 440nm laser head having the largest

deviation of 1.36%. However, neither the functioning of the SMELLIE system nor any subsequent analysis of the data will be impacted in any meaningful way by such small differences in the light's wavelength, and therefore for convenience, the laser heads will be referred to by their nominal wavelengths hereafter. Figure 4.5 shows the measured wavelength distributions at varying output power for each laser head. The 375nm and 405nm distributions are very symmetric across a range of output powers, but the 440nm and 500nm laser heads have more asymmetric distributions with a slight decrease in peak wavelength as the output power decreases, thus explaining why these have a larger difference between nominal and measured average wavelengths than the 375nm and 405nm laser heads.

As noted previously in Section 1.2.3, the photocathode efficiency of the SNO+ PMTs is wavelength-dependent, peaking between approximately 400nm and 500nm, and the range of wavelengths used in SMELLIE has been chosen to take advantage of this. At the time when SMELLIE was originally proposed, the intention was to load the SNO+ detector with ^{150}Nd , which strongly absorbs light at very specific wavelengths: $< 370\text{nm}$, between 450 and 490nm, and $> 500\text{nm}$ [90], and therefore the individual laser head wavelengths were chosen on the basis of avoiding these absorption peaks. However, ^{130}Te has no such peaks, instead possessing a very smooth absorption spectrum [90]. This change does not affect the higher wavelength laser heads, but the observed output power of the 375nm laser head will be reduced, as ^{130}Te -loaded scintillator is more opaque to light at 375nm wavelength than ^{150}Nd -loaded scintillator would have been, and so fewer photons will travel across the detector unabsorbed. Further emphasis may also be put on the higher wavelengths during the loaded phase of the experiment due to the specific loading technique used, which may introduce extra absorption peaks and necessitate the use of additional wavelength shifters to avoid them.

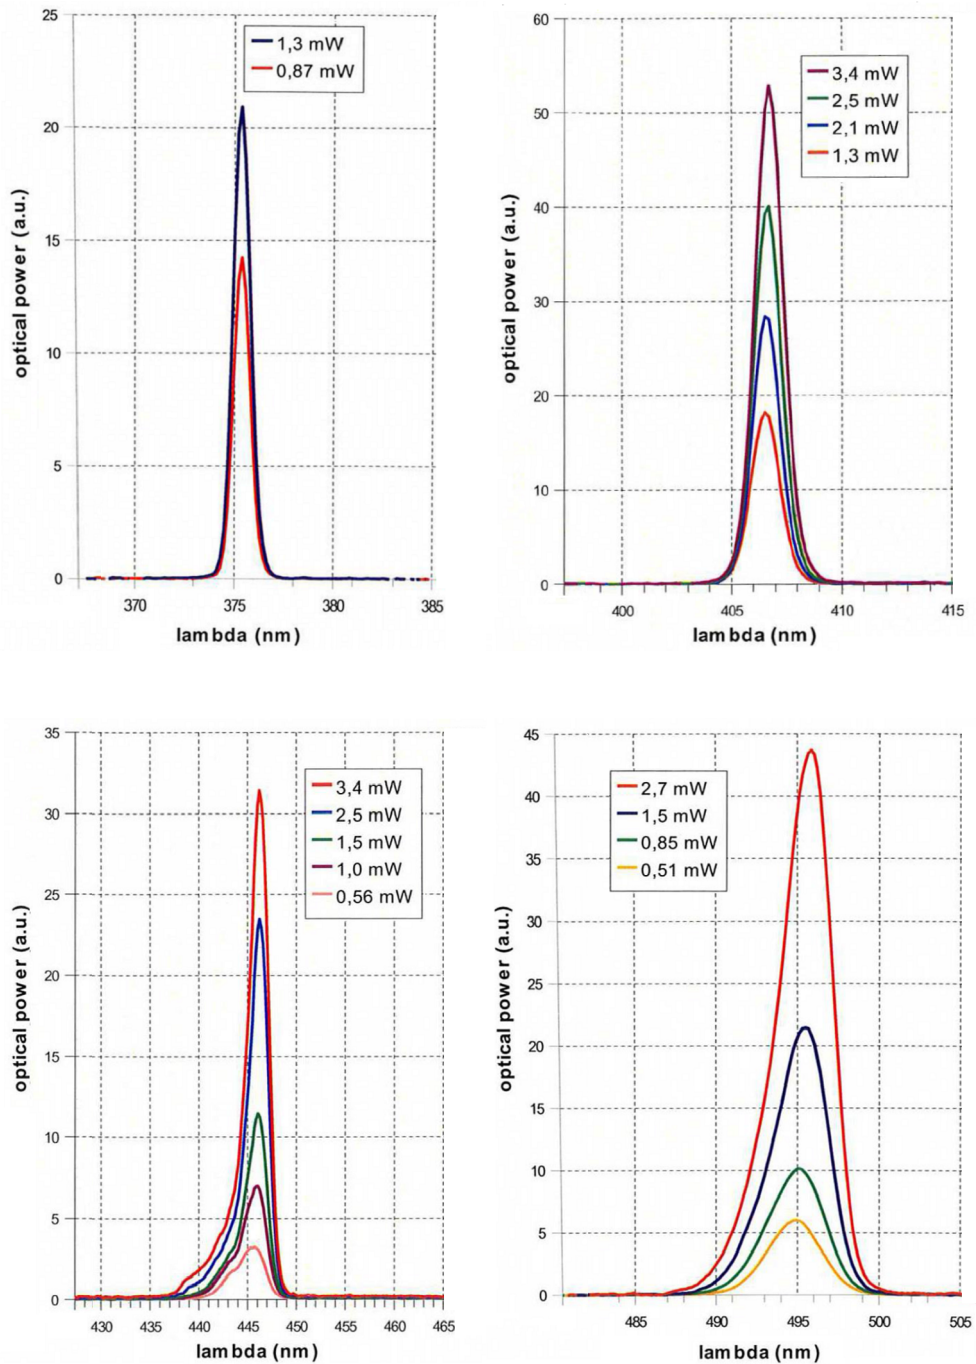


Figure 4.5: The wavelength distributions at varying output power, for the 375nm (top-left), 405nm (top-right), 440nm (bottom-left) and 500nm (bottom-right) laser heads. All measurements were performed by PicoQuant and provided in documentation.

Figure 4.6 shows the measured average single-pulse time profiles of each laser head. In all cases, the pulse begins no later than 2ns after the laser head is actually fired (indicated, but not shown, by time 0ns on each distribution's x-axis), and is very short, with a FWHM of < 200 ps. All four laser heads show some after-pulsing, most notably at 440nm wavelength.

4.2.2 Laser Head Control

As noted previously, each laser head requires two settings: the per-pulse intensity (as a percentage of the maximum possible output) and the repetition rate. These are both handled through the SEPIA II Laser Driver (abbreviated to *SEPIA* hereafter), shown in Figure 4.7. This device, manufactured by PicoQuant [91], contains a single driver module in the configuration used by SMELLIE, and so is only capable of sending commands to a single laser head at any one time. (A second driver module could in principle be added to SEPIA, together with a sequencer that would allow synchronised timing control across two laser heads. However, this is prohibitively expensive and, as will be discussed below, unnecessary.) The intensity and repetition rate are set from the SMELLIE Control PC (which is connected to SEPIA via USB), either through SEPIA's own bespoke interface or through the general SMELLIE user interface discussed previously.

As well as setting the intensity and repetition rate, SEPIA also controls the firing of the laser heads, through the use of either an internal (generated by SEPIA itself) or external (fed into SEPIA) trigger signal. There are several preset internal trigger rates available, from 40MHz to 2.5MHz, however all of these are too high for safe usage - at these rates the average light intensity entering the detector would cause significant damage to the SNO+ PMTs in the forward direction of the beam. Instead, an external trigger signal with a much lower frequency of between 1 and 10kHz is always used, which is low enough to avoid causing damage to the detector but still high enough to

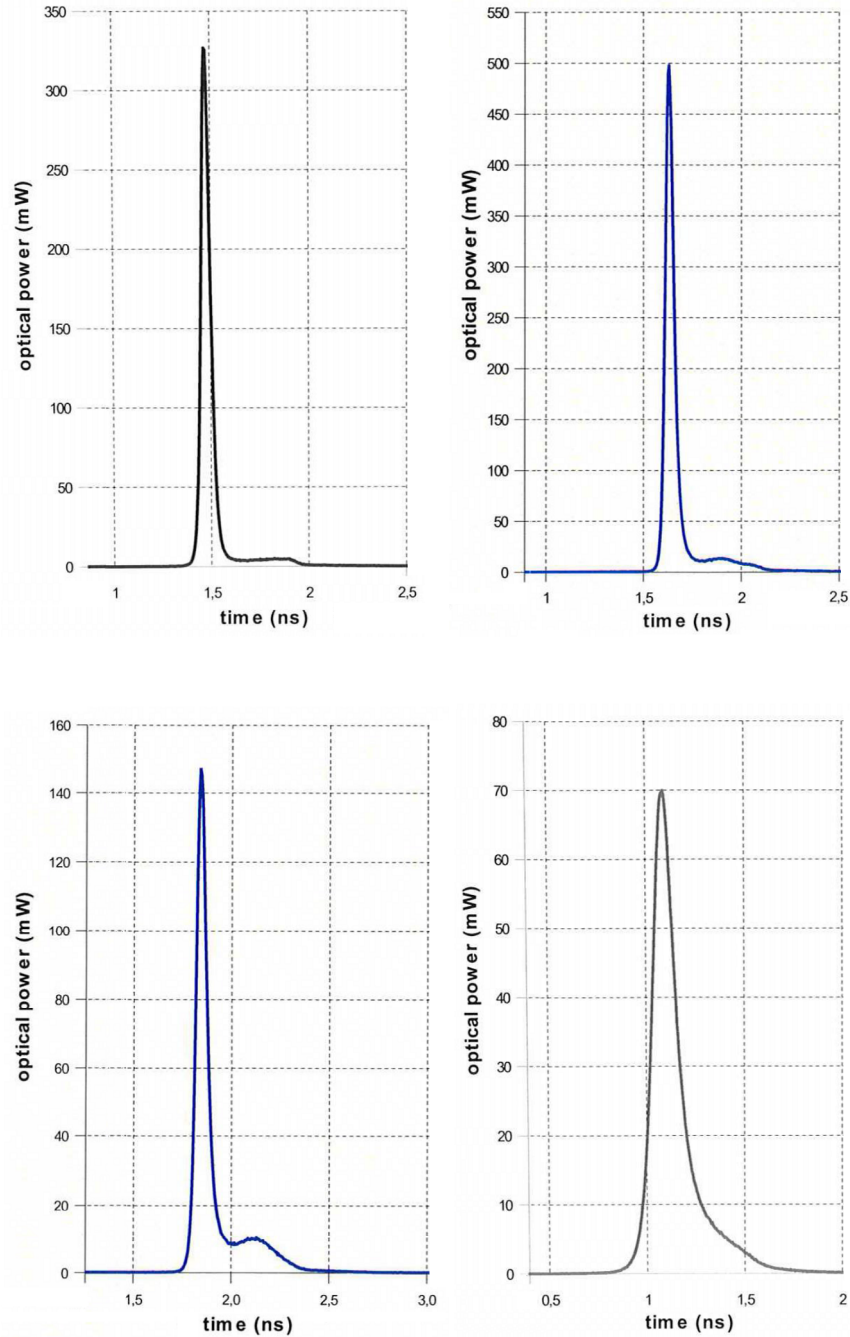


Figure 4.6: The measured time profile of a single pulse (averaged and normalised over many pulses) for the 375nm (top-left), 405nm (top-right), 440nm (bottom-left) and 500nm (bottom-right) laser heads. In all cases, the laser head was fired at time 0ns (not shown). All measurements were performed by PicoQuant and provided in documentation.



Figure 4.7: The Laser Switch (left) and the SEPIA II Laser Driver (right) connected by a LEMO cable (black). The Laser Switch's output sockets are located on the rear of the device. Also shown are the laser head interlock key and cables (orange) on SEPIA, the front panel controls on the Laser Switch, and the USB connections between the SMELLIE Control PC and each component.

produce enough photons for a reliable analysis to be performed. The external trigger signal is a square pulse with a low logic level of $-1V$ and a high of $+2V$, sent to SEPIA from the Monitoring PMT Unit (discussed in Section 4.2.6). These logic levels are not standard, but instead are set so as to satisfy SEPIA's own triggering requirement that the voltage (either on the rising or falling trigger-pulse edge) crosses the $-0.3V$ level. The frequency of this external signal is equivalent to the repetition rate of the laser heads.

Without any further hardware, changing between laser heads would require an operator to manually remove and reconnect each laser head's signal cable into the LEMO port on SEPIA. This is problematic, as not only is it a slow process, but the ELLIE electronics rack is located within the SNO+ deck clean room, and therefore access is very limited during normal detector operation. The use of a second driver module in

SEPIA would not solve this issue either, as this would allow control of only two laser heads at a time out of four. Instead, the Laser Switch, shown on Figure 4.7, has been developed as an automated alternative to manual connection.

The Laser Switch receives the electronic LEMO output signal of SEPIA, which would otherwise be connected directly to the desired laser head. Within the Laser Switch, mechanical relays are used to route this signal to one of six possible output sockets (referred to as *channels*), into which the four laser heads are connected based on a pre-determined mapping (with two channels spare). It should be emphasised that no optical signals are routed through the Laser Switch - only the voltages that comprise the LEMO signal. The Laser Switch is built around a daughter-card construction, with the route from input to each channel being controlled by a single removable card that contains all of the necessary electronics for only that channel. This modular design makes replacement of the relays very straight-forward, which is important as they will age with continuous use and replacement will eventually and inevitably be required. The choice of channel, and therefore the choice of which laser head is being controlled, can be made in one of two ways. The channel can be manually set using the control panel on the front of the Laser Switch - for safety and to avoid accidental channel-switching, this is a two-step process requiring the operator to first *set* the desired channel, and only then *execute* the change. The operator must also know the channel-to-laser head mapping prior to setting the channel, as this is not indicated on the Laser Switch. Alternatively, and during standard operation of SMELLIE, the laser head may be selected using its wavelength directly from the general SMELLIE interface - the correct channel will automatically be set according to the channel-to-laser head mapping.

4.2.3 Optical Paths within SMELLIE

Optical Fibres

The optical fibres used throughout SMELLIE can be separated into two groups. The fibres that route light between the laser heads and the Fibre Switch will be referred to as *rack fibres* - these are all 1 metre in length and number twelve: one between each of the four laser heads and its corresponding beamsplitter, one between each beamsplitter and its corresponding attenuator, and one between each attenuator and its corresponding Fibre Switch input socket. Those fibres which transport light from the Fibre Switch into the detector are designated as *detector fibres* - these also number twelve (for reasons that will be discussed in Section 4.2.5 below), but have a longer length of 45m. All fibres are *FC/APC* terminated (an acronym that denotes an angle-polished ferrule connector, which describes the angled face of the fibre's end).

Both groups of fibres are of the same type: multimode with a $50\mu\text{m}$ -diameter core, from Corning's InfiniCor SX*i* range [92][93]. This type of fibre has an optimal operating wavelength range of between 750 and 1450nm, since it is designed for use primarily in commercial Ethernet technology, which uses lasers operating at a wavelength of 850nm. This operating range is at significantly higher wavelengths than those used by SMELLIE, and so the fibres will not perform at their best, i.e. there will be some extra attenuation of light and therefore an unavoidable decrease in output light intensity. Extrapolating the known attenuation spectrum of these fibres [92] back into the SMELLIE wavelength range gives an approximate attenuation of 7 to 9dB per kilometre of fibre, where the attenuation is defined as:

$$\text{Attenuation (dB)} = -10 \times \log_{10} \left(\frac{\text{Output Intensity (W)}}{\text{Input Intensity (W)}} \right) \quad (4.1)$$

This corresponds to a negligible attenuation in each of the rack fibres, and no more

than 0.5dB of attenuation in each detector fibre.

Despite not being best-suited for the wavelengths used in SMELLIE, this type of fibre is nonetheless desirable as it contains very low intrinsic radioactivity. This is an important issue to consider, as the fibres that enter the detector are situated very close to the SNO+ PMTs, and therefore any radioactive decay products emanating from the fibres will enter the PMTs with no possibility of being removed or reduced. A study [94], the results of which are summarised in Table 4.3, has shown that the total radioactivity of all fibres used in SMELLIE is significantly lower than the pre-existing and irreducible background from all PMTs combined. The effect of the fibre radioactivity is in fact even smaller than these results would indicate, since the radioactive mass is evenly distributed throughout the entire length of each fibre, and so the majority of the radioactivity will be located away from the PMTs.

Component	^{238}U (g)	^{235}U (g)	^{232}Th (g)	^{40}K (g)	Total (g)
Single Detector Fibre	5.56×10^{-5}	6.83×10^{-8}	5.62×10^{-6}	2.20×10^{-2}	
12 Detector Fibres	6.672×10^{-4}	8.196×10^{-7}	6.744×10^{-5}	2.640×10^{-1}	0.2647
Single SNO+ PMT	1.00×10^{-4}	[NA]	1.00×10^{-4}	[NA]	
9438 SNO+ PMTs	9.438×10^{-1}	[NA]	9.438×10^{-1}	[NA]	> 1.8876

Table 4.3: The radioactive masses of four major isotopes that are present in a single SMELLIE detector fibre of total mass 313g, all twelve detector fibres, a single SNO+ PMT and all 9438 SNO+ PMTs. All values are taken from [94]. The PMT values include the radioactivity present in the glass and the coaxial cable, but information about ^{235}U and ^{40}K are not available for the PMTs.

Beamsplitters

The light leaving each laser head is immediately passed through a 2-input/2-output beamsplitter (with one input going unused), where the low-fraction output is routed towards the Fibre Switch and the high-fraction output leads to the Monitoring PMT Unit (discussed in Section 4.2.6). In principle, the Monitoring PMT should receive approximately the same per-pulse light intensity from all four laser heads, however each laser head has a different intrinsic power output even at a fixed intensity, as noted on Figure 4.5. Therefore, each laser head is connected to a separate beamsplitter, and the individual splitting ratios are chosen based on the laser head power, with the highest-power laser head being connected to the most asymmetric beamsplitter (10% of the light output going to the detector and 90% to the Monitoring PMT) and the lowest-power laser head sending light through a more symmetric beamsplitter (both the detector and Monitoring PMT receiving 50% of the light).

The beamsplitters used in SMELLIE are of the fused-fibre taper design, shown schematically in Figure 4.8. This type of splitter is created by heating, twisting and then drawing out two optical fibres such that there is a central region (the *coupling region*) where the fibre cores are intertwined with uniform but reduced radii. Either side of the coupling region are the *tapered regions*, which provide smooth transitions between the unchanged fibres and the coupling region.

For an optical fibre with a core of radius a and refractive index n_1 surrounded by cladding of refractive index n_2 , the *normalised frequency* V is given by [95]:

$$V = \frac{2\pi a}{\lambda} \times \sqrt{n_1^2 - n_2^2} \quad (4.2)$$

This dimensionless quantity is a measure of how many modes the fibre can support

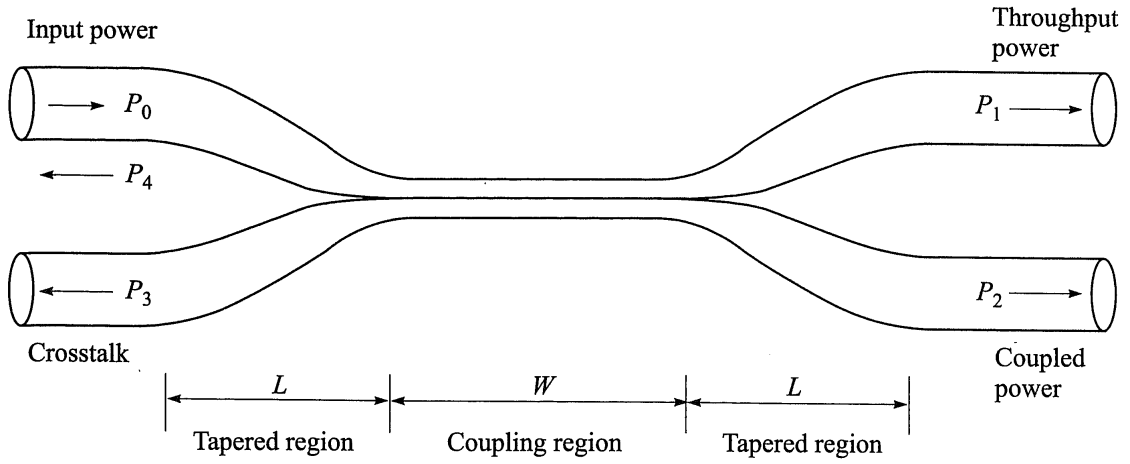


Figure 4.8: The internal design of a fused fibre taper splitter, adapted from [95], showing the primary (top) and secondary (bottom) fibres, the coupling region with length W and the two tapered regions each with length L . Light enters the splitter along the primary fibre with power P_0 and exits through the two outputs with powers P_1 and P_2 . Reflections and scattering in the fibres can cause light to travel the “wrong” way back towards the inputs, with powers P_3 and P_4 in the secondary and primary fibres respectively, however these are negligible effects.

at a wavelength of λ . Due to the smaller fibre radius in the coupling region, V is much smaller in the coupling region than it is in the unchanged primary fibre, and therefore fewer modes can be supported, with those that cannot be decoupled from the primary fibre’s core. However, some of this decoupled light can recouple into the secondary fibre’s core, with the exact fraction of recoupled light (the *coupled power* P_2) being dependent on the length W of the coupling region, as well as k - the coupling coefficient between the electromagnetic fields in the primary and secondary fibres [95]:

$$P_2 = P_0 \sin^2(kW) \quad (4.3)$$

Noting that the total output power must be the same as the total input power P_0 , and assuming negligible P_3 and P_4 , leads to an expression for the power remaining in

the primary fibre (the *throughput power* P_1):

$$P_1 = P_0 \cos^2(kW) \quad (4.4)$$

The splitting ratio of the beamsplitter is then the ratio of the coupled and throughput powers. The beamsplitters used in SMELLIE are designed for use at a wavelength of 800nm, and the fact that the laser heads operate below this means that the actual splitting ratios are different from the nominal values, and will also vary with wavelength, as shown in Table 4.4.

Wavelength (nm)	Nominal Splitting Ratio (% / %)	Measured Splitting Ratio (% / %)
375	50 / 50	55 / 45
	70 / 30	73 / 27
	90 / 10	91 / 9
405	50 / 50	68 / 32
	70 / 30	59 / 41
	90 / 10	78 / 22
440	50 / 50	52 / 48
	70 / 30	73 / 27
	90 / 10	93 / 7
500	50 / 50	60 / 40
	70 / 30	80 / 20
	90 / 10	95 / 5

Table 4.4: The nominal and measured splitting ratios of the beamsplitters used in SMELLIE, at the four laser head wavelengths and at 100% laser head intensity.

Attenuators

The light leaving the low-fraction output of each beamsplitter passes through an attenuator before entering the Fibre Switch, with each beamsplitter being connected to a separate attenuator. These attenuators allow for very fine adjustment of the light's intensity before it enters the Fibre Switch, allowing each laser head's input into the detector to be individually tuned in order to give the maximum amount of light that does not cause damage to the PMTs.

4.2.4 Fibre Switch

The final component through which light passes before entering the detector is the Fibre Switch, manufactured by Laser Components UK [96], and shown in Figure 4.9.



Figure 4.9: The front panel of the Fibre Switch, with the 5 input fibre sockets on the left and the 14 output sockets on the right. At the time of this picture, the detector fibres were not connected to their respective output sockets, but instead were temporarily housed under the pink foam in the foreground. The RS232 socket that connects the Fibre Switch to the SMELLIE Control PC is located on the rear.

Each attenuator is connected via a rack fibre to one of the Fibre Switch's input

FC/APC-terminated sockets (with one input left spare) using a pre-determined mapping between the input socket number and the laser head wavelength, and each of the twelve detector fibres is connected to one of the output FC/APC-terminated sockets (with two outputs left spare) based on a mapping between the detector fibre ID (discussed in Section 4.2.5) and the output socket number. The combination of a single input and single output, equivalent to a single laser head firing into a single detector fibre, is called a Fibre Switch *channel*. There are seventy possible channels numbered from 1 to 70 inclusive, where channel 1 is the combination of input 1 and output 1, channel 2 is equivalent to input 1 and output 2, etc. The choice of channel can only be made using the SMELLIE Control PC - there is no manual control of the Fibre Switch. The PC is connected to the Fibre Switch via RS232, and the operator must specify the wavelength and detector fibre ID. The correct channel number is then calculated using the aforementioned mappings and sent to the Fibre Switch in order to open the optical path between the desired input and output sockets.

Different channels require different optical paths through the Fibre Switch, and this is achieved through the use of three Micro-Opto-Electromechanical Systems (MOEMS) switching blocks, as shown schematically in Figure 4.10. Each block can be thought of as containing two sub-blocks which move with respect to each other in order to create an unbroken optical path between any given input and output socket combination.

For example, consider that the user requires input 4 and output 11. In this situation, sub-block (1) must first move “down” so that input 4 aligns with the connection going out of (1). The second movement is by sub-block (2), which must move “up” so that the connection coming from sub-block (1) aligns with the “lowest” of the four connections out of (2), leading to output sockets 11 to 14. The final movement - by sub-block (3) - is then “down”, so that the connection from (2) matches with output 11, creating a path from input 4 to output 11 through all three blocks. (It should be




Figure 4.10: Schematic of the internal design of the Fibre Switch, showing the 5 inputs (left), the three MOEMS switching blocks with each containing two sub-blocks (orange and green), the various inter-block optical connections, and the 14 outputs (right). This diagram has been adapted from documentation received via personal communications with Laser Components UK. Four of the 5 inputs are fed through the first switching block (1), with input 5 bypassing it entirely. The 14 outputs in sub-block (3) are grouped into 4 sections: 1 - 4, 5 - 8, 9 - 10 and 11 - 14, with each section being fed by a separate connection out of the second sub-block, labelled (2).

noted that this is not literally what occurs - it, and Figure 4.10, are simply schematic descriptions of the process.)

The movement of the sub-blocks is very fast, with the Fibre Switch taking $< 2\text{ms}$ to fully change to a new channel [96], along with a lock-out time of 20ms after each successful channel change during which no new commands will be accepted.

There will be some attenuation of the light as it travels through the Fibre Switch. Some of this loss can be due to human error, such as misaligned or over-tightened input and/or output connections, and some loss can be attributed to mechanical issues such as the switching block(s) failing to move the full required distance for a given channel, and thus being left slightly misaligned and allowing light to leave the optical path. In addition, the fibres used within the Fibre Switch cannot be perfectly transmitting at all wavelengths being used, and so there will be some wavelength-dependent absorption and scattering of light within these fibres. The measured total attenuation of each channel is shown in Figures 4.11 and 4.12, where the attenuation is given by Equation 4.1 with the input and output intensities being measured before and after the Fibre Switch respectively.

Light at wavelengths of 375nm and 405nm show the most overall attenuation, with losses of up to 2.5dB, but in contrast the 440nm and 500nm light transmits somewhat better, with losses no greater than 1.4dB.

There is a very clear increase in attenuation on outputs 5 to 8 across all inputs and wavelengths. Figure 4.10 indicates that these outputs share the same inter-block connection between the second and third blocks, and so it can be concluded that there may be some damage or opaqueness on this connection. These four outputs show consistent attenuation amongst themselves, indicating that the problem does indeed lie with the inter-block connection rather than with the four individual output paths within the third block.

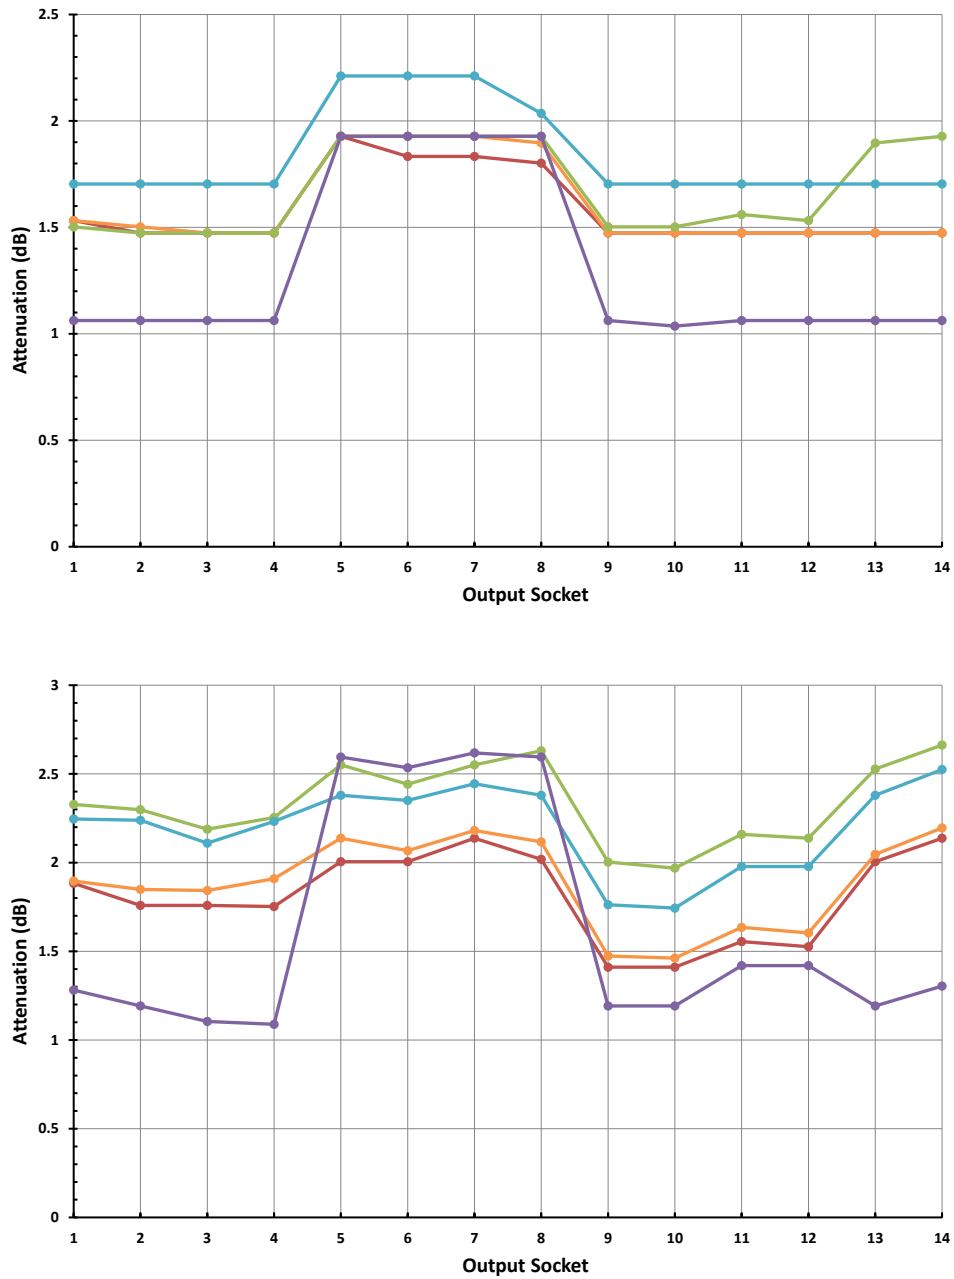


Figure 4.11: The attenuation of each output socket as defined in Equation 4.1 for light entering the Fibre Switch through input 1 (red), 2 (orange), 3 (green), 4 (blue) and 5 (purple) at a wavelength of 375nm (top) and 405nm (bottom).

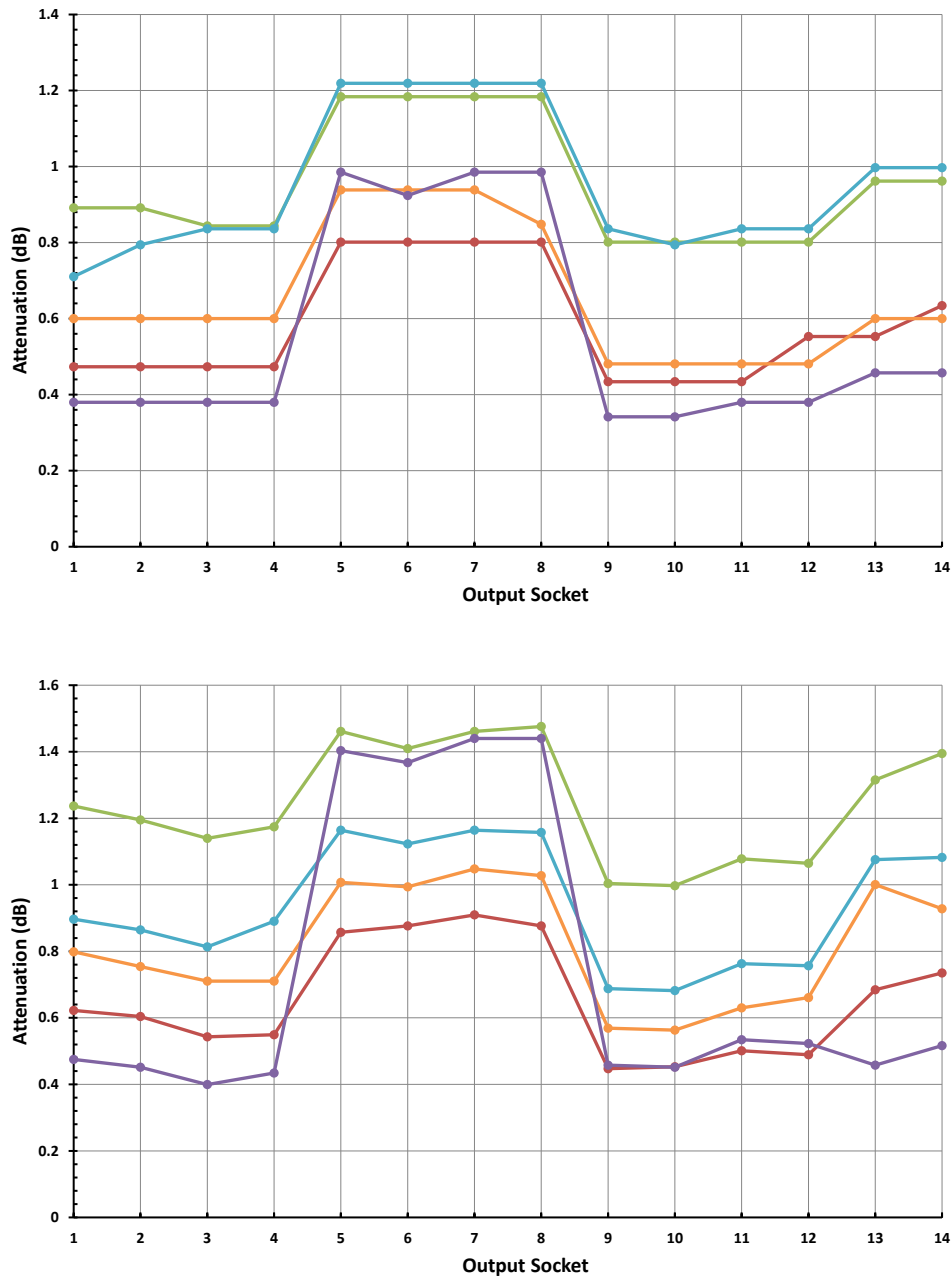


Figure 4.12: The attenuation of each output socket as defined in Equation 4.1 for light entering the Fibre Switch through input 1 (red), 2 (orange), 3 (green), 4 (blue) and 5 (purple) at a wavelength of 440nm (top) and 500nm (bottom).

Input 5 consistently has the smallest attenuation across all outputs and wavelengths. This can be explained by noting that input 5 connects via only two switching blocks rather than all three (shown on Figure 4.10), and so there is one fewer source of attenuation within the Fibre Switch for connections made using this input. The ordering of the inputs according to attenuation is also consistent across all outputs and wavelengths, with input 3 having the highest, followed by 4, 2, 1 and 5 in order of decreasing attenuation.

4.2.5 In-Situ Components

Collimators

The numerical aperture, NA , of any optical fibre is defined using the refractive indices of the fibre's core and cladding - n_{core} and n_{clad} respectively, and is also related to the half-angle of the fibre's emission cone (known as the *acceptance angle*, θ_{max}):

$$\begin{aligned} NA &= \sqrt{n_{core}^2 - n_{clad}^2} \\ &= n_{med} \sin(\theta_{max}) \end{aligned} \tag{4.5}$$

where n_{med} is the refractive index of the material into which the light from the fibre is being emitted. The fibres used in SMELLIE have $NA = 0.200 \pm 0.015$ [92], and so the acceptance angle of light exiting these fibres into air ($n_{med} = 1.0003$) is $\sim 11.5^\circ$. This is a considerably divergent beam of light, and so to reduce this spread, the detector-end of each detector fibre is bonded to a collimator - the internal design of which is shown in Figure 4.13, with Table 4.5 summarising the individual components.

The most important part of the collimators is the grin lens - this is the component that produces the actual collimation of the light. The grin lens is not a true lens, but rather




Figure 4.13: The internal design (to scale) of a single SMELLIE collimator, with the detector fibre shown in orange and other components as detailed in Table 4.5. Each collimator is 420mm in length.

Component Label	Component Name	Material	Information
(a)	Front Portion	Delrin	Front and Rear Portions bonded together with UV-cured glue and mechanically coupled using a pin/slot lock
(b)	Rear Portion	Delrin	
(c)	Grin Lens	Quartz	Graded refractive index, manufactured by crystallographers at University of Oxford
(d)	Optical Fibre	-	[see Section 4.2.3]
(e)	Glass Ferrule	Glass	Manufactured by ThorLabs
(f)	Glass Tube	Glass	
(g)	Strain Relief Tube	Stainless Steel	Holds (d) in position securely against (c) and centered in the collimator
(h)	Fibre Furnication Tube	Plastic	

Table 4.5: Summary of the internal components of a single SMELLIE collimator, as labelled on Figure 4.13.

a cylindrical quartz rod that has been manufactured such that the refractive index at the edges is smaller than that at the centre of the rod (known as a *graded index*), meaning light travels faster at the edges than at the centre. This then replicates the idea that light travels through less material, therefore leaving at a high velocity and after a shorter time, at the edge of a convex lens with a uniform refractive index than at its centre. This graded index grin lens design was chosen simply because it is considerably less difficult to manufacture a graded index quartz rod than such a small convex lens, and the focal length of the grin lens is purely dependent on its length alone.

Figure 4.14 shows the post-collimator beamspot profile of one detector fibre in air. The half-opening angle of the collimators has been calculated as $\sim 3^\circ$ - a value that is very consistent across all twelve detector fibres. Although there is some small variation of this angle between fibres due to slight variations in the manufacturing of the collimators, just the mere presence of the grin lenses means that there must be some collimation occurring - it is not physically possible for the emitted beam to have a half-opening angle of 11° or more.

Mounting Plates

The twelve detector fibres feed into the detector at four locations on the PSUP, with three fibres per location and each location consisting of a single mounting plate, shown in Figure 4.15, fixed to the outside of a PMT hex-cell. The mounting plates are designed such that the three fibres point in different directions: 0° , 10° and 20° with respect to the axis from the plate to the detector centre. This setup allows several different paths through the detector to be illuminated, each of which has a different length through the detection medium, AV and cavity - thus allowing trends within each region to be investigated.

Each of the four mounting plates is referred to using the number of the PSUP

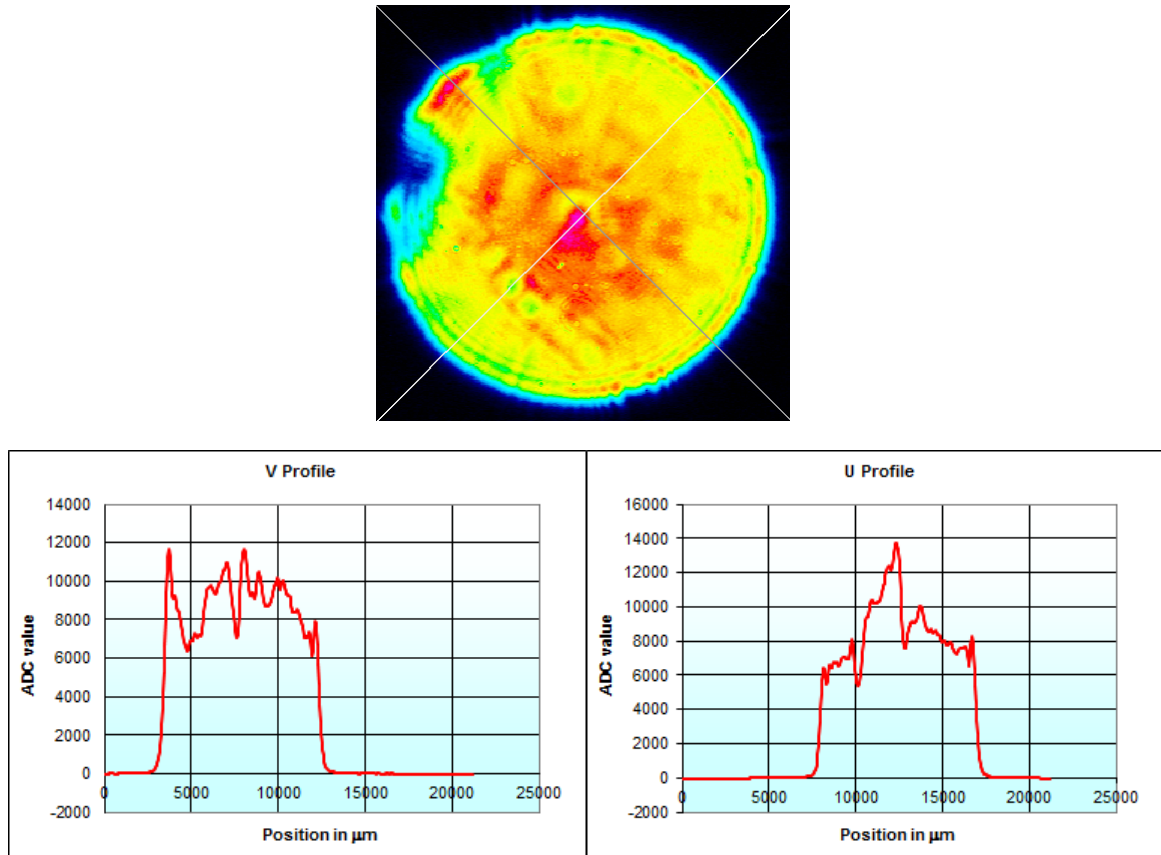


Figure 4.14: (top) The beamspot produced by one detector fibre in air, using the 500nm laser head and imaged using a CCD camera. (bottom) The intensity profiles of the beamspot along the V (left) and U (right) axes, corresponding to the top-left to bottom-right and bottom-left to top-right directions respectively on the CCD image. Three sets of such profiles were taken for each fibre at varying distances between fibre and camera, from which the half-opening angle of each fibre was calculated.

node to which the plate's parent hex-cell belongs (07, 25, 37 and 55), and each fibre on a plate is designated by its angle (0, 1 and 2 for 0° , 10° and 20° respectively). In this way, each of the twelve detector fibres has a unique identifying label, i.e. fibre FS137 is the 10° fibre mounted on a hex-cell at node 37 (with the "FS" denoting a SMELLIE optical fibre, in contrast to the fibres also present in the PSUP that are part of the other ELLIE sub-systems).

The plates and supporting components are manufactured from a polymer called Delrin (a commercial name for the chemical compound Polyoxymethylene [98]). This




Figure 4.15: Schematic of a mounting plate, adapted from [97], together with its associated components: (a) the fence, (b) the retainer bar and (c) the strain relief bar. The positions of the three fibres are indicated in orange, and on the left-hand image the 10° , 0° and 20° positions are at left, centre and right respectively.

material was chosen due to its very low intrinsic radioactivity, the measurement of which has been performed and is summarised in Table 4.6 with comparison to the radioactive backgrounds originating in the SNO+ PMTs. The effect of the four mounting plates is seen to be negligible. In addition, Delrin has been found to resist deformation under prolonged contact with and submersion in water, and therefore the fibre angles will remain unchanged throughout the duration of SNO+, and the plates themselves will remain firmly attached to their respective hex-cells.

4.2.6 Ex-Situ Monitoring

The energy of a single laser head pulse varies from pulse to pulse, with the amount of variation being strongly dependent on the laser head's operating intensity. During operation at medium and high intensities the variation will be on the order of a few percent, but this will increase at low intensities, i.e. in the regime close to and below each laser head's lasing threshold, where the output becomes unstable. A large variation at medium and high intensities can also be a sign of abnormal laser head

Component	^{238}U (g)	^{235}U (g)	^{232}Th (g)	^{40}K (g)	Total (g)
Single Mounting Plate	5.34×10^{-7}	9.68×10^{-10}	3.21×10^{-9}	4.78×10^{-4}	
4 Mounting Plates	2.136×10^{-6}	3.872×10^{-9}	1.284×10^{-8}	1.912×10^{-3}	0.001914
Single SNO+ PMT	1.00×10^{-4}	[NA]	1.00×10^{-4}	[NA]	
9438 SNO+ PMTs	9.438×10^{-1}	[NA]	9.438×10^{-1}	[NA]	> 1.8876

Table 4.6: The radioactive masses of four major isotopes that are present in a single SMELLIE mounting plate (including the supporting components), all 4 plates and components, a single SNO+ PMT and all 9438 SNO+ PMTs. All values are taken from [94]. The PMT values include the radioactivity present in the glass and the coaxial cable, but information about ^{235}U and ^{40}K are not available for the PMTs.

behaviour. Even a small pulse energy variation will have an effect on the number of photons entering the detector, and therefore on the number and properties (i.e. spatial and time distributions, charge, etc.) of any triggered PMTs. However, the detector's response alone does not indicate if such changes are due to variations in the laser head or changes occurring within the detector itself, and therefore it is important to know the laser head output before light enters the detector - a task which is accomplished using the ex-situ Monitoring PMT Unit (abbreviated to *MPU*), shown in Figure 4.16. The various external connections to and from the MPU are described in Table 4.7.

The MPU's PMT receives light from the high-fraction output of each laser head's beamsplitter, and the raw output of the PMT is then slowed and shaped. This is required due to the very short duration of the laser head pulses, as previously shown in Figure 4.6 - such pulses are simply too fast for measurements to be performed on them using commercially available electronics. The post-shaping PMT pulse is shown on Figure 4.17.

The maximum voltage generated by the PMT is directly related to the laser head




Figure 4.16: The Monitoring PMT Unit, containing the PMT (top-left), the analog electronics board (middle), the digital electronics board (bottom) and the PMT's manual Gain Control (bottom-left).

intensity, and so a per-pulse measurement of this maximum voltage is equivalent to measuring the per-pulse intensity, with an appropriate conversion factor that can be determined. This measurement is performed in two stages by a sampling circuit - this first *samples* the voltage on the PMT pulse for a length of time equal to the rise time of the PMT pulse, at the end of which the sampling circuit switches to a *hold* state,

Connection Name	Source Component	Signal Details	Destination Component
Light Input	Laser heads	Light pulses via optical fibre	-
PMT Pulse	-	PMT Output (0 to 2V)	SNO+ CAEN Digitiser
Sampling Pulse	-	Sampling circuit Output (0 to 2V)	SNO+ CAEN Digitiser
NI TTL Input	NI Unit	TTL (0 to +5V)	-
MTC ECL Output	-	ECL (-1.75 to -0.9V)	SNO+ MTC
MTC ECL Input	SNO+ MTC	ECL (-1.75 to -0.9V)	-
SEPIA Output	-	Square pulse (-2 to +1V)	SEPIA

Table 4.7: The connections to and from the MPU, separated into those associated with ex-situ laser head monitoring (top) and those related to trigger signal transmission and reception (bottom).

whereby the output of the sampling circuit is now fixed at the last sampled voltage, i.e. the maximum PMT voltage, which is then read out. The entire output of the sampling circuit (the *sampled pulse*) is also shown on Figure 4.17.

Both the PMT and sampled pulses are passed to the SNO+ CAEN digitiser, which prepares the data for inclusion in the overall event datastream. The CAEN digitiser can only accept signals up to +2V, and so it is important that the Monitoring PMT's gain be adjusted so that the held voltage from the sampling circuit does not exceed this limit. Each laser head requires a slightly different PMT gain setting, due to their aforementioned differing power outputs even at a single intensity.

As well as monitoring the per-pulse laser head energy, the MPU serves as a transceiver for the various trigger signals required by the SMELLIE system in both operating




Figure 4.17: The PMT pulse after being slowed and shaped (blue), and the sampled pulse (yellow). The sampled pulse lags behind the PMT pulse by 50ns during the sampling state due to various delays in the electronics and the sampling chip itself, and then holds at the maximum PMT pulse voltage. The height of the PMT pulse, and therefore the final held voltage, is dependent on the PMT’s gain setting, which can be manually adjusted.

modes. This functionality is unconnected to the laser head monitoring - it is simply more efficient to locate the electronics for both functions within the same unit, as both sets of electronics require power and an isolated and stable environment, and space on the ELLIE electronics rack is limited.

In master mode, the primary signal (the 0 to +5V TTL pulse from the NI Unit) is received by the MPU and branched into two signals. One branch is converted into an ECL pulse (-1.75 to -0.9V) and sent to the SNO+ MTC. This conversion from TTL to ECL is required due to the 50m distance between the ELLIE electronics rack and the MTC - a TTL signal travelling over such a long cable will degrade significantly, whereas ECL signals are specifically designed to travel faster and with significantly less degradation over long cable lengths [99]. The second branch remains as a square pulse, but the low and high voltage levels are converted to -2V and +1V respectively,

as required to trigger SEPIA.

In slave mode, the primary signal received by the MPU is an ECL signal from the MTC, which is simply converted to the square -2V to +1V signal required by SEPIA - there is no requirement to split the signal when in slave mode.

Chapter 5

Scattering Measurements in Water

‘Have you figured out that red splodge yet?’

— Luca Cavalli, on being scientifically rigorous

5.1 Simulation of SMELLIE in RAT

5.1.1 The ELLIE Monte-Carlo Generator

The *ELLIE generator* [100] has been specifically developed to simulate the emission of light into the detector from the various ELLIE systems previously outlined in Section 4.1.1.

When dealing with the SMELLIE sub-system in particular, the generator does not actually simulate any of the ex-situ or in-situ hardware discussed in Chapter 4, but instead begins with a point-like source of optical photons that substitutes for the emission end of a generic SMELLIE detector fibre. This photon source is then assigned a fibre ID and a wavelength, based on which specific detector fibre and laser head is required for the simulation. In principle, the fibre IDs used by the generator are supposed to follow the same naming scheme as that which was discussed for the real fibres in Section 4.2.5, but in reality some IDs have been changed, as will be

discussed in Section 5.1.2 below. The choice of fibre ID then determines the position, direction and angular distribution of the photon source. These quantities correspond respectively to the position of the real mounting plate that contains the desired fibre, the fibre's forward direction with respect to the detector centre, and the fibre's post-collimator intensity distribution as discussed in Section 4.2.5. All of this information, for the twelve real detector fibres, is stored in the ELLIE generator's RAT database file. The chosen wavelength of the photon source determines the properties that are related to the real laser heads: the wavelength (previously shown in Figure 4.5) and timing (Figure 4.6) distributions of the emitted photons, which are also all stored in the generator's database. The final property of the source is the number of photons emitted per event (equivalent to the number of photons per pulse of the real laser head), which will be denoted by n_γ , and is related to the laser head intensity.

5.1.2 Comparisons between Real and Simulated Data

A number of studies have been conducted to test the performance of the ELLIE generator in comparison to real SMELLIE data taken in the detector [101][102]. These studies were conducted using data taken in early 2014, when the SNO+ detector was filled almost entirely with *air*, with only a small volume of water at the bottom of the cavity. Therefore, the simulations tested here also used an air-filled simulated detector. Only six SMELLIE detector fibres - those located at PSUP nodes 37 and 55 - were installed in the detector at the time of these performance studies, with the remaining six at nodes 07 and 25 to be installed in the second half of 2015.

Figure 5.1 shows the distribution of triggered PMTs across many laser head pulses emitted from fibre FS155 at a wavelength of 407nm, in both real data (top) and the original iteration of the ELLIE generator (bottom).

It can be seen that there were significant discrepancies between the simulated and




Figure 5.1: A flat map of the PSUP showing the locations of all PMTs that were triggered across 6000 laser head pulses emitted from fibre FS155 in real data (top) and 3000 simulated events using the original ELLIE generator (bottom). The colours represent the frequency at which a particular PMT was triggered, summed across all pulses, with blue denoting PMTs that were triggered 1 - 5 times, and red representing PMTs that were triggered 2500 - 3000 times. Different numbers of events have been used in order to take into account the different number of photons per pulse in real and simulated data, such that the background triggering frequency of non-beamspot PMTs is consistent across both plots. Only PMTs that were triggered by the SMELLIE system have been included in the real data, and the blank area of PMT coverage is due to two electronics crates being turned off at the time that this data was taken.

real data: the forward beamspots (indicated by the red regions of PMTs) were in completely different places on the PSUP, and the PMT triggering frequency within and immediately around the beamspot (directly related to the angular intensity distribution of the beam) was also somewhat different, with the simulated beamspot being slightly smaller and more clearly defined at its edges than that in data.

The fact that the beamspots were located at very different positions indicated that the intended fibre directions had not been adhered to during installation - i.e. the real fibre FS155 was not pointing in the same direction as its simulated counterpart. Such discrepancies were noted between five of the six installed detector fibres and their simulated versions, and are due to the real fibres having been installed in the incorrect slots on their respective mounting plates. By matching the PSUP flat maps across various pairs of real and simulated fibres, it was possible to determine the actual directions of the real fibres - a summary of which is given in Table 5.1.

Real Fibre ID	FS037	FS137	FS237	FS055	FS155	FS255
True Direction	10°	0°	20°	10°	20°	0°

Table 5.1: The IDs of the six real SMELLIE detector fibres that have been installed in the detector as of writing, and their true directions calculated using comparisons with the simulated fibres of known directions. Only fibre FS237 has the same direction as intended - the other five fibres are found to have been installed at incorrect angles in the detector.

To correct for these errors in the fibre installation, the intended fibre directions in the generator's database have been changed to match the true directions, i.e. the simulated FS155 fibre now points at 20° with respect to the detector centre, in the same way as the real FS155 fibre does. This does now mean that the fibre direction can no longer be assumed from the first number in the fibre ID, but the mapping

between the fibre IDs and directions given in Table 5.1 is clearly documented in the ELLIE generator's database. The effect of these changes to the ELLIE generator is shown in Figure 5.2 - the simulated beamspot is now located in the same place on the PSUP as the one in real data.

It is known from installation reports that the SMELLIE mounting plates have been installed on the wrong PMT hex-cells at each of the two nodes - a situation that is depicted for node 37 in Figure 5.3.

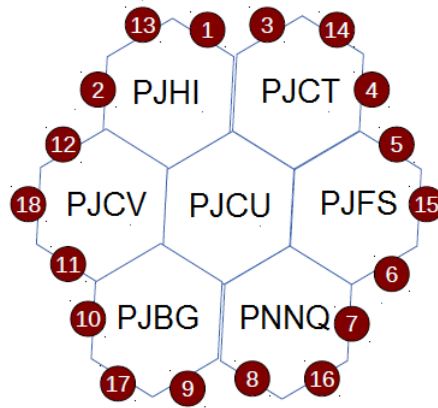


Figure 5.3: A schematic of the 7 PMT hex-cells (labelled by their unique 4-letter IDs) at PSUP node 37 along with the 18 possible positions at which the SMELLIE mounting plate at this node could have been installed, adapted from [102]. The intended position was that numbered 8 (along the bottom), but position 10 (along the lower-left side) is the actual one.

Correcting for this installation error is easy, as the mounting plate positions stored in the generator's database can simply be changed from the intended to the actual coordinates. The effect on the simulation is negligible, as seen by comparing Figures 5.2 and 5.4 (before and after the mounting plate position correction respectively).

The final remaining discrepancy between simulation and data occurs in the angular




Figure 5.2: A flat map of the PSUP showing the locations of all PMTs that were triggered across 6000 laser head pulses emitted from fibre FS155 in real data (top) and 3000 simulated events using the ELLIE generator with the corrected fibre directions (bottom). In comparison to Figure 5.1, correcting the fibre direction has made a significant difference, as the real and simulated beamspots are now found in the same place on the PSUP.




Figure 5.4: A flat map of the PSUP showing the locations of all PMTs that were triggered across 6000 laser head pulses emitted from fibre FS155 in real data (top) and 3000 simulated events using the ELLIE generator with the corrected fibre directions and positions (bottom). In comparison to Figure 5.2, correcting the position of the mounting plate in simulation has made no noticeable difference.

profile of the SMELLIE beamspot.

It can be seen in Figure 5.4 that the beamspot in data shows a central, high triggering frequency (i.e. high intensity) region approximately 4 PMTs in diameter, surrounded by a concentric pattern of progressively lower intensity rings, with the entire beamspot (i.e. those PMTs that were triggered more than a total of 5 times) being approximately 10 PMTs wide, or $\sim 2\text{m}$. Using a simple trigonometrical calculation, this corresponds to a fibre half-opening angle of 3.25° - a value which is similar across all six installed real fibres, and completely consistent with the measured half-opening angles of the fibres noted previously in Section 4.2.5. In contrast, the simulated beamspot is somewhat smaller - with a diameter of 8 PMTs (corresponding to a half-opening angle of 2.60°), and exhibits a more top-hat shaped angular profile - with a wider central high intensity region and much more well-defined edges than the real beamspot (as expected, since the simulation utilises the top-hat shaped beam profiles previously shown in Figure 4.14).

The reasons for these differences - in particular the presence of the low intensity tails at the edges of the real beamspot - are yet to be fully understood, but one possible source is the CCD camera that was used to image the beamspots and subsequently extract the angular intensity profiles that are used in the ELLIE generator. It is possible that the CCD camera was not sensitive to the low intensity tails of the profile, thereby resulting in the top-hat functions seen previously. However, the SNO+ PMTs are sensitive to such tails when considering large numbers of SMELLIE events. Another possibility is an unaccounted-for optical effect in the detector - such as diffuse scattering within the AV acrylic - which is currently not being simulated in RAT. It can then be surmised that the low intensity tails of the real beamspot are caused by a small number of photons whose trajectories were very slightly altered away from the central beam region as they passed through the AV.

Presently, the solution to the beamspot width discrepancy is to extract the beam's

angular intensity distribution from the real data, and then use this in the simulation rather than the previously measured intensity distributions from Figure 4.14. However, this procedure has not really worked, as it results in a significantly wider and lower intensity profile, with tails extending out to a half-angle of 11° - as shown in Figure 5.5. While it does now introduce the lower intensity tails into the simulation as desired, it can be seen that the central region of the simulated beamspot is no longer as high in PMT triggering frequency (and intensity) as the data, and it is unclear as to why the distribution seemingly extracted from the real data deviates so much from it.

The ongoing discrepancy between the real and simulated angular intensity distributions will have an impact on the analysis of scattering in angular beams, which will be discussed in Section 5.5. The analysis framework has been developed and tested under the scenario of a pencil beam - i.e. with an opening angle of 0° , and it has been subsequently shown that applying it directly to a simulated SMELLIE beam (where the simulation uses the 22° angular intensity distribution shown in Figure 5.5) results in a considerable loss in scattering reconstruction performance, mostly due to the photons that are emitted from the fibre at larger angles. However, as noted above, the simulated beamspot intensity profile is significantly wider than the real one, meaning that there are considerably more photons emitted at larger angles in simulation than in the real SMELLIE system, and so the contribution of these large-angle photons to the simulated reconstruction performance will be larger than expected. Therefore, the pencil beam analysis may actually perform better when applied to the real SMELLIE data than might be expected from the simulated results shown in Section 5.5.




Figure 5.5: A flat map of the PSUP showing the locations of all PMTs that were triggered across 6000 laser head pulses emitted from fibre FS155 in real data (top) and 3000 simulated events using the ELLIE generator with the corrected fibre directions and positions, and an angular intensity profile that was extracted from the real data (bottom). This version of the ELLIE generator is used in all further work presented here.

5.2 Isolation of Scattered Light

5.2.1 Comparison of Different Optical Processes

In order to perform an accurate and reliable study of the optical scattering in water, it is first necessary to obtain as pure a sample as possible of this scattered light. Unfortunately, there are a number of competing optical processes that a photon could be subject to as it travels across the detector. These include absorption in any of the media encountered, reflection from the various AV surfaces, reflection from the PMTs' light concentrators, and optical scattering in the cavity water or the AV acrylic. In addition, light can undergo combinations of these processes, or potentially none of them.

All of the various permutations, apart from those that end with the absorption of the photon, eventually lead to the photon entering and triggering a PMT, but the process(es) it underwent along its path will influence which PMT it enters and at what time relative to its emission from the fibre. For example, if a photon does not reflect or scatter anywhere along its path, it will trigger a PMT almost directly across the detector from the fibre at an easily calculable time (since the distances travelled in the various media will be well-known for a purely refracted path between fibre and PMT). On the other hand, if the photon underwent a single reflection from the first AV surface that it encountered, it must then enter a PMT fairly close to the fibre's location, and at a very short time after emission, since the distance from fibre to AV and back to the PSUP is much shorter than the path across the entire detector.

Although the PMT trigger times are independent of the SMELLIE hardware (instead being measured relative to the start of the event trigger window as previously discussed in Section 1.1.2), the positions of any triggered PMTs will be dependent on the location of the fibre from which the photons are emitted. However, instead of using the PMT coordinates relative to the centre of the detector, an alternative

fibre-independent measure of a PMT's position is the relative angle between that PMT and the fibre's forward direction, shown as angle α on Figure 5.6.




Figure 5.6: An illustration of the SNO+ detector (not to scale) showing a SMELLIE fibre located at position \underline{f} (with respect to the centre of the detector). A photon is emitted from the fibre, and (in this example) scatters at a location \underline{s} in the detection medium, before going on to trigger a PMT located at position \underline{p} , as shown by the green line. The orange line indicates the path the photon would have taken had it not scattered. The scattering angle θ is shown (green), and α (purple) indicates the angle between the fibre's forward direction and the triggered PMT at \underline{p} .

The following figures 5.7 to 5.20 are derived from a single simulation of a water-filled detector, into which 5000 laser head pulses with $n_\gamma = 2000$ were delivered through fibre FS025 from the 405nm laser head. The photons' angular distribution is that of

a simulated SMELLIE beam, rather than an idealised pencil beam of zero angular width. The effect of varying n_γ will be explored in Section 5.2.3. Photon tracking information has been used to identify the various optical processes that each photon was subject to.

Figure 5.7 shows the calibrated PMT trigger time vs. angle α for PMTs that were triggered by photons that did not reflect or scatter anywhere along the path between the SMELLIE fibre and PMT (these will hereafter be referred to as *beamspot* PMTs for convenience).




Figure 5.7: The calibrated trigger time vs. angle α for beamspot PMTs - those triggered by photons which did not reflect or scatter anywhere in the detector as they travelled between the SMELLIE fibre and the PMT.

The vertical striations present in this and subsequent trigger time vs. α distributions are an artifact of the binning, and a result of the PMT positions (and therefore α) taking discrete rather than continuous values.

Taking this into account, the beamspot PMTs have an even distribution in α up to

$\sim 26^\circ$, which is consistent with the simulated SMELLIE beam of full angular width of 22° being subsequently widened slightly due to the refraction that occurs as it passes across the AV surfaces. The beamspot PMTs are almost all triggered within a well-defined 100ns-wide time window between 300ns and 400ns. However, within this window the distribution is not even, as shown in Figure 5.8.




Figure 5.8: The calibrated trigger times of beamspot PMTs. This distribution is simply the projection of Figure 5.7 onto the trigger time axis.

The shape of this distribution is a direct consequence of the *photoelectron transit time* - the time that a photoelectron takes to travel from the PMT's photocathode on the inside surface of the glass face to the first dynode at the bottom of the PMT's bulb. In RAT, this time is not calculated, but is determined for every individual photoelectron by sampling a transit time probability distribution which is shown on Figure 5.9, along with a normalised and shifted version of Figure 5.8 for comparison.

Figure 5.10 shows the calibrated trigger time vs. α for PMTs triggered by photons that were subject to a single reflection from one of the AV surfaces (*AV-reflected*

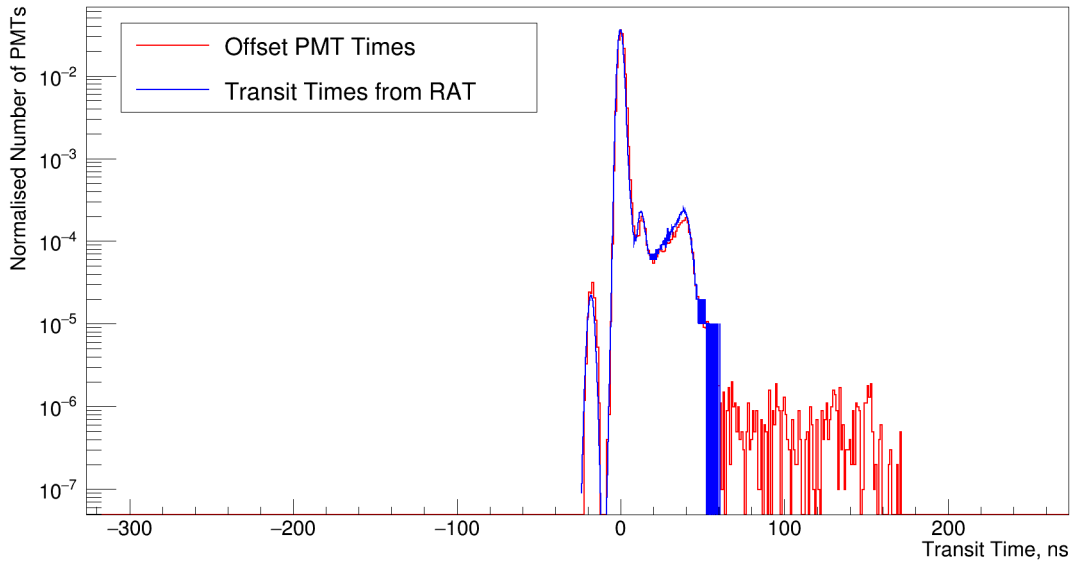


Figure 5.9: The photoelectron transit time probability distribution as stored in RAT (blue) and the calibrated trigger times of beamspot PMTs (red), where the latter distribution is the same as that shown in Figure 5.8, but has been normalised and shifted in time so that the two distributions' highest points overlap. The slight differences elsewhere in the distributions are caused by the additional contributions to the PMT distribution at transit times > 60 ns, beyond the bounds of the transit time probability distribution.

PMTs).

The AV-reflected PMTs are split into two time-separated populations. The first is found at earlier trigger times than the beamspot PMTs and very close to $\alpha = \pi$ - i.e. exactly opposite the beamspot PMTs, very close to the position of the fibre itself. This population is caused by reflections from the two AV surfaces closest to the fibre: photons are reflected back towards the fibre from either AV surface without ever entering the detection medium, and therefore travel along a very short path between fibre and PMT, resulting in an early trigger time. The second AV-reflected PMT population occurs at considerably later times than the first - similar to and slightly later than the beamspot PMTs, and over a wider range of α . This population is a result of reflections from the two AV surfaces that photons can encounter *after* passing through the detection medium: reflection redirects the photons back towards




Figure 5.10: The calibrated trigger time vs. angle α for AV-reflected PMTs - those triggered by photons that were reflected from exactly one of the four possible AV surfaces that could be encountered between fibre and PMT.

the fibre, but due to the greater distances involved they can go on to potentially enter any PMT located in the same hemisphere of the detector as the fibre itself.

Figure 5.11 shows the calibrated trigger time vs. α for PMTs triggered by photons that first underwent a single reflection from another PMT's light concentrator (*concentrator-reflected* PMTs).

The distribution of concentrator-reflected PMTs can be explained through a similar argument to that which was previously made for the second population of AV-reflected PMTs. A photon must have first travelled across the entire detector before encountering a light concentrator, explaining why the vast majority of these PMTs are not triggered until after the beamspot PMTs. Due to the geometry of the concentrators discussed previously in Section 1.1.1, the photon can then in principle be reflected in any direction, and therefore concentrator-reflected PMTs can occur at any angle α . There is also a small population of these PMTs that overlaps with




Figure 5.11: The calibrated trigger time vs. angle α for concentrator-reflected PMTs - those triggered by a photon that was first reflected from the light concentrator of a different PMT.

the beamspot PMTs at low α and trigger times of ~ 330 ns - these are caused by photons that reflect from a beamspot PMT's concentrator directly into that same concentrator's PMT, and effectively therefore look identical to beamspot PMTs.

Figures 5.12 and 5.13 show the calibrated trigger time vs. α for PMTs triggered by photons that scattered once in the cavity water or AV (*cavity-scattered* PMTs) and once in the detection medium (*innerAV-scattered* PMTs) respectively.

The two time-separated populations of cavity-scattered PMTs reflect the fact that a photon can scatter in the cavity either before entering (early-time population) or after exiting (late-time population) the detection medium, and can therefore trigger a PMT either before or after the beamspot PMTs are triggered. There is a higher density of cavity-scattered PMTs at both very low (scattering angle $\theta = 0^\circ$) and very high ($\theta = 180^\circ$) angle α , indicative of the $[1 + A \cos^2(\theta)]$ distribution of θ in Rayleigh scattering. This is more clearly seen in photons that scattered in the cavity




Figure 5.12: The calibrated trigger time vs. angle α for cavity-scattered PMTs - those triggered by photons that were scattered exactly once in either the AV acrylic or the cavity water.




Figure 5.13: The calibrated trigger time vs. angle α for innerAV-scattered PMTs - those triggered by photons that were scattered exactly once in the detection medium.

before entering the detection medium than in photons that scattered after leaving the medium. This is due to absorption of the photons: a photon that was scattered backwards after exiting the detection medium has a much longer path and more chance of being absorbed than a photon that was scattered in a forward direction. Photons that were scattered in a forward direction either before or after the detection medium will travel along paths very similar to those that did not scatter or reflect at all, hence why the two populations converge at the same trigger times and α 's as beamspot PMTs.

The innerAV-scattered PMTs follow a broadly similar distribution: occurring at any angle α and at trigger times before, similar to and after those of beamspot PMTs, with the forward-scattered (low α) photons leading to PMTs triggering at the same time and α as beamspot PMTs due to their similar paths. Since there is only a single detection medium, there is no division of the innerAV-scattered PMTs into multiple time-separated populations.

Figure 5.14 shows the calibrated trigger time vs. α for PMTs triggered by photons that underwent multiple optical processes (*multiple-effect* PMTs) between fibre and PMT.

The multiple-effect PMTs occur at any value of angle α , as the combinations of multiple scatterings and reflections serves to homogenate any angular distributions produced by these individual processes. Multiple-effect PMTs are triggered across a much wider range of trigger times as well, as multiple processes serve to both shorten and/or lengthen the paths taken by photons before they reach the PMTs. Some densities can still be seen in the multiple-effect distribution, which are indicative of particular combinations - for example, there is one dense region that strongly resembles the distribution of concentrator-reflected PMTs, but is in fact caused by a combination of forward-scattering either in the cavity or the detection medium




Figure 5.14: The calibrated trigger time vs. angle α for multiple-effect PMTs - those triggered by photons that underwent multiple optical processes between fibre and PMT, for example scattering in both the AV acrylic and cavity water, or reflection from a PMT light concentrator after being scattered in the detection medium.

followed by reflection from a PMT concentrator.

Figure 5.15 shows the calibrated trigger time vs. α for *noise* PMTs - these are not triggered by photons emitted from the SMELLIE system, but are instead randomly triggered throughout the detector due to thermionic emission of photons from the PMT cathodes, electronic noise, or internal PMT electric field breakdowns.

This homogenous and isotropic background of triggered PMTs is always present, occurring at any angle α (since the noise is fibre-independent), and at any calibrated trigger time between 100 and 500ns. This lower time limit is related to the *cable delay* between the PMTs and the readout electronics: if a photon triggers the PMT at time t_p (the *actual* trigger time, which is itself the incidence time of the photon on the PMT's photocathode plus the photoelectron's transit time to the PMT's dynode), the *calibrated* trigger time used in all of the previous distributions is then the time at which the PMT's signal actually reaches its corresponding PMTIC, i.e. t_p plus the




Figure 5.15: The calibrated trigger time vs. angle α for noise PMTs - those that are randomly triggered throughout the detector, independent of any physical processes occurring in the detection medium. The lack of noise PMTs at angles $\alpha < 0.4$ is discussed in the main text.

time that the signal takes to travel along the PMT's cable to the PMTIC. This cable delay is approximately the same for all real PMTs, and is treated as a constant time within RAT.

Figure 5.16 shows the combined calibrated trigger time vs. angle α distribution for all PMTs in a single simulation of 5000 laser head pulses emitted in an angular beam at a wavelength of 405nm from fibre FS025, with $n_\gamma = 2000$.

It can be seen that some PMTs that should be associated with a particular process are misidentified as having been triggered by a different process - for example, the apparent beamspot PMTs at $\alpha < 0.4$ and trigger time < 300 ns should be noise PMTs (thus explaining the lack of noise PMTs at $\alpha < 0.4$ on Figure 5.15), and the expected multiple-effect PMTs at $\alpha < 0.4$ and trigger time > 400 ns are actually depicted as beamspot PMTs. This is due to multiple photons being incident on the same PMT within a single trigger window: in the time that it takes photon 1 to travel from the




Figure 5.16: The calibrated trigger time vs. angle α for all triggered PMTs in a single simulation of 5000 laser head pulses emitted in an angular beam at a wavelength of 405nm from fibre FS025 into a water-filled detector, with 2000 photons per pulse. The PMTs are colour-coded according to the optical process(es) each photon was subject to before it triggered the PMT.

fibre to a given PMT along a particular path, that same PMT can first be triggered by a different photon 2 that was emitted later, but arrives earlier having undergone some other optical process. The PMT's trigger time will therefore be that of photon 2, but tracking the photon paths through the simulation will lead to this PMT being hit by photon 1, and therefore associated with photon 1's optical process(es).

Qualitatively, there is an approximately triangular region between the two populations of cavity-scattered PMTs which is predominantly populated by innerAV-scattered PMTs. Isolating this region using cuts on the calibrated PMT trigger time and α angle would therefore provide the purest sample of photons that were scattered exactly once in the detection medium.

However, it should be noted that the calibrated PMT trigger time is extremely dependent on the group velocity of light of the specific detection medium being used - and therefore by extension, the medium's refractive index. The refractive indices of water and LABPPO are quite different at the wavelengths used in SMELLIE - 1.344 and 1.505 respectively, at 400nm - and therefore there will be a non-negligible difference between the times that a photon takes to travel across the entire detector before and after SNO+ moves from the water phase to the scintillator phase, meaning that any cuts based on the absolute calibrated PMT trigger times would need to be re-calculated and re-optimised for each phase of the experiment. While this is not a fundamentally difficult issue, a more robust method would be to instead reformulate the parameter space in terms of a different time-based parameter - one that is not directly affected by a change in the detection medium's optical properties.

5.2.2 Time Residuals

Instead of using the unaltered calibrated PMT trigger time, a more material-independent quantity known as the *time residual*, t_{res}^i , can be defined as:

$$t_{res}^i = t_{pmt}^i - t_{flight}^i \quad (5.1)$$

where t_{pmt}^i is the calibrated trigger time of the i^{th} PMT as previously used, and t_{flight}^i is the time of flight between the i^{th} PMT and the known SMELLIE fibre position. This is calculated using RAT's LightPath class [103], which uses a direct path between PMT and fibre (previously shown in purple on Figure 5.6) and takes into account both refraction at the various surfaces as well as the different group velocities of light in the various materials encountered along this path, thereby making the overall time residual independent of the detection medium being used.

Figure 5.17 shows the t_{res}^i vs. α distribution for the same simulation previously depicted in Figure 5.16, once again separated and coloured according to which optical process (if any) was encountered by the photon that triggered each PMT.

The PMTs located at high α have experienced the smallest shift in time (only a few ns) when moving from t_{pmt}^i to t_{res}^i , and in contrast, the low- α PMTs (beamspot PMTs and the forward-scattered cavity- and innerAV-scattered PMTs) have shifted by more than 60ns. The difference between these shifts is due to the use of the direct path in the calculation of t_{flight}^i : PMTs that are very close to the fibre position (large α) have a very short direct path - significantly shorter than the actual path that the photon took - resulting in a small value of t_{flight}^i and therefore only a small difference between t_{pmt}^i and t_{res}^i . In contrast, PMTs that are located further away from the fibre (at small α) have much longer direct paths which are more similar in both length and time to the actual paths that the photons took. Therefore, these PMTs have a larger t_{flight}^i term and thus a large change from t_{pmt}^i to t_{res}^i . Figure 5.17 also shows that these α -dependent time shifts have “straightened out” the previously curved contributions from the individual optical processes, such that the predominantly innerAV-scattered




Figure 5.17: The time residual (defined in Equation 5.1) vs. angle α for all triggered PMTs in a single simulation of 5000 laser head pulses emitted in an angular beam at a wavelength of 405nm from fibre FS025 into a water-filled detector, with 2000 photons per pulse. The PMTs are colour-coded according to the optical process(es) each photon was subject to before it triggered the PMT.

region of PMTs in the centre can now be bounded by simple linear functions rather than, for example, higher-ordered polynomial ones.

Two linear cuts have been simultaneously applied to the t_{res}^i vs. α parameter space in order to find the region that contains the purest selection of innerAV-scattered PMTs - hereafter referred to as the *acceptance region*. The acceptance region is bounded from the bottom (Equation 5.2) and top (Equation 5.3):

$$t_{res}^i > c_{bottom} + m_{bottom}\alpha \quad (5.2)$$

$$t_{res}^i < c_{top} + m_{top}\alpha \quad (5.3)$$

with the fact that α cannot be greater than π acting a natural bound on the right side.

A balance should be made between maximising the acceptance region's *purity* - the fraction of all PMTs within the region that are innerAV-scattered, and its *efficiency* - the fraction of all innerAV-scattered PMTs that are found in the region, as it is easy to fall into the undesirable situation of maximising the purity while leaving the efficiency so low that any results become statistics-dominated. Therefore, the free parameters of the cuts - the two intercepts c_{bottom} and c_{top} and the two gradients m_{bottom} and m_{top} - have been simultaneously optimised such that the value of the parameter N , given in Equation 5.4, is maximised:

$$N = \frac{\text{number of innerAV-scattered PMTs within acceptance region}}{\sqrt{\text{number of other PMTs within acceptance region}}} \quad (5.4)$$

In the context of the acceptance region, this quantity is equivalent to the commonly used "signal over square-root of background" or *signal significance*, in that the region's signal is the innerAV-scattered PMTs, and the background is any other contribution

to the region.

5.2.3 Acceptance Cuts I: Varying n_γ

Table 5.2 shows the optimised values of the four free parameters that maximise N , along with the corresponding values of purity and efficiency, for a single fibre and wavelength combination, but varying values of n_γ - the number of photons in each laser head pulse. For each n_γ , the specific contributions to the acceptance region from the various PMT classifications are also given.

The optimised parameters are extremely consistent across different values of n_γ , which is expected since increasing n_γ is equivalent to increasing the density of the entire t_{res}^i vs. α distribution equally, apart from the contribution from noise PMTs - increasing n_γ does not alter the shape of the distributions from the individual optical processes, and therefore the acceptance region should remain in the same position. The efficiency also remains very stable at $\sim 7\%$. This very low value is due to the acceptance region not including any of the forward-scattered innerAV-scattered PMTs found at low α . As seen in Figures 5.7 to 5.14, every optical process apart from AV-reflected PMTs has a significant population in the same area of the t_{res}^i vs. α distribution as forward-scattered innerAV-scattered PMTs - therefore if the acceptance cuts were to include these PMTs, the purity of the region would be decreased significantly. Even at this very low efficiency, Table 5.2 shows that there are certainly enough innerAV-scattered PMTs remaining within the acceptance region (many thousands in a run of 5000 laser head pulses with 2000 photons per pulse) for statistical fluctuations not to be a concern.

The cuts completely eliminate the contribution from beamspot PMTs at all values of n_γ , but - although the acceptance region is clearly dominated by innerAV-scattered PMTs - there is still a contribution of at least 10% from the noise PMTs at the lower values of n_γ . This can be reduced relative to the other contributions by increasing

	n_γ	500	1000	1500	2000	2500	3000
	C_{bottom}	240	242	242	241	242	242
Optimised	m_{bottom}	11	11	11	11	11	11
Parameters	C_{top}	220	220	220	220	220	220
	m_{top}	51	50	51	51	50	51
	N	91.3026	152.535	191.601	230.995	263.636	289.535
Results	Purity	0.686	0.767	0.775	0.783	0.798	0.795
	Efficiency	0.074	0.068	0.069	0.072	0.068	0.070
	Beamspot	0 (0.00)	0 (0.00)	0 (0.00)	0 (0.00)	0 (0.00)	0 (0.00)
Number	AV-reflected	100 (1.80)	88 (0.96)	198 (1.44)	338 (1.79)	345 (1.57)	493 (1.81)
(Percentage)	Concentrator	18 (0.32)	37 (0.40)	100 (0.73)	186 (0.99)	255 (1.16)	368 (1.35)
of PMTs	-reflected						
	Cavity	188 (3.38)	243 (2.64)	400 (2.91)	648 (3.44)	661 (3.00)	870 (3.20)
in	-scattered						
Acceptance	InnerAV	3814 (68.61)	7053 (76.74)	10661 (77.50)	14753 (78.34)	17563 (79.83)	21632 (79.49)
Region	-scattered						
	Multiple	472 (8.51)	875 (9.52)	1455 (10.58)	1962 (10.42)	2263 (10.23)	2954 (10.85)
	-effect						
	Noise	966 (17.38)	895 (9.74)	943 (6.85)	945 (5.02)	914 (4.15)	897 (3.30)

Table 5.2: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs, all as functions of varying n_γ - the number of photons per laser head pulse. Each simulation contained 5000 laser head pulses emitted in an angular beam at a wavelength of 405nm from fibre FS025.

n_γ - this increases the total number of PMTs that were triggered by photons, but since the noise is independent of the SMELLIE hardware, the absolute number of noise PMTs in the acceptance region remains almost unchanged, and therefore its relative contribution decreases. However, this means that at larger n_γ the highest contribution to the acceptance region apart from the innerAV-scattered PMTs becomes the multiple-effect PMTs, and since this contribution appears to increase with n_γ , increasing it further will not increase the purity of the acceptance region.

The best value of n_γ to use in the real detector is currently unknown, and can only be accurately determined using real data, but the consistency of the acceptance cut parameters across different values of n_γ means that - at least for the purposes of isolating a region of the detector that is predominantly populated by scattered light - it does actually not matter what value of n_γ is used.

In subsequent work presented here, a choice of $n_\gamma = 1500$ has been made unless explicitly stated otherwise - this is an arbitrary decision, motivated only by the facts that: 1) a simulation of 5000 such pulses at a wavelength of 405nm gives ~ 10000 innerAV-scattered PMTs in the acceptance region, and 2) the total simulation time decreases when a smaller value of n_γ is used.

5.2.4 Acceptance Cuts II: Varying Fibre & Wavelength

Tables B1 to B8 show the optimised values of the four free parameters that maximise N (given by Equation 5.4), along with the corresponding values of purity and efficiency, for a single value of n_γ , but varying fibre / wavelength combinations. For each combination, the specific contributions to the acceptance region from the various PMT classifications are also given.

For a given SMELLIE fibre, the optimised acceptance cut parameters are extremely

consistent across different values of wavelength. However, the wavelength does have an effect on the purity of the acceptance region: $\sim 78\%$ at 375nm and 405nm, but dropping slightly to $\sim 75\%$ at 440nm and then further down to $\sim 66\%$ at 500nm. The efficiency of the region is also affected by the wavelength: holding at between 5 and 6% for the highest wavelength, but increasing slightly to between 6 and 7% as the wavelength decreases.

Varying the wavelength also has a clear effect on the individual contributions to the acceptance region from the various optical processes, for any given fibre. As the wavelength increases, the scattering length of light in the detection medium (pure water) also increases, meaning that fewer photons will undergo Rayleigh scattering, and therefore fewer will be available to contribute to the acceptance region - as can be seen, with ~ 13000 innerAV-scattered PMTs in the region at 375nm, but ~ 3000 at 500nm. For the same reason, the absolute number of cavity-scattered and multiple-effects PMTs also decreases with increasing wavelength, although the relative contribution of the former remains fairly constant, whereas that of the latter is almost halved between 375nm and 500nm. A knock-on effect of these is that, although the absolute number of AV-reflected PMTs in the region decreases with increasing wavelength, the relative contribution increases from $\sim 1.5\%$ at 375nm to $> 3\%$ at 500nm. The final effect of the varying wavelength is that, although the absolute number of noise PMTs remains fairly constant (since the acceptance region's total area in the t_{res}^i vs. α parameter space also stays stable due to the consistent cut parameters), the relative contribution from noise increases by almost a factor of 4 due to the constant or decreased contributions from other optical processes.

For a fixed wavelength, it can be seen that the optimised parameters for the “bottom” acceptance cut are very consistent across all fibres, but the “top” cut's parameters are slightly different for the 20° fibres compared to the 0° and 10° ones, resulting

in a smaller area of the t_{res}^i vs. α parameter space being covered by the acceptance region. The acceptance region's purity is fairly stable regardless of the fibre's direction, but the efficiency decreases slightly for the 20° fibres - an effect of there being fewer innerAV-scattered photons contained in the smaller acceptance region for this direction.

The acceptance cuts fail to completely eliminate the contribution from beamspot PMTs at 20°, due to the beamspot having a larger spread in α at this fibre direction. (The need to minimise this contribution is the reason why the optimised acceptance region is slightly smaller in area at 20°.) The relative contribution from innerAV-scattered PMTs decreases slightly due to the smaller acceptance region and the increased contribution from beamspot PMTs, as does that of noise PMTs, at the cost of an increased contribution from multiple-effects PMTs.

It can be seen that at wavelengths of 440nm and 500nm, both the absolute and relative contributions from AV-reflected PMTs are higher at 20° than at 0° and 10°, unlike at 375nm and 405nm - where the AV-reflected PMTs contribution is fairly consistent regardless of fibre direction. This is likely due to a combination of two effects: the higher emission angle means that the AV-reflected PMTs will no longer be found at $\alpha \simeq \pi$ (i.e. very close to the fibre position), but rather at values of t_{res}^i and α which overlap with those of innerAV-scattered PMTs, and the higher wavelengths result in more total internal reflection from the AV surface due to the different refractive indices.

Although many of the fibre / wavelength combinations have been seen to share identical or very similar optimised acceptance cut parameters, the individual values for each combination will be used in subsequent work, rather than a single sets of parameter covering all possible fibre and wavelength choices.

5.2.5 Post-Acceptance Cuts PMTs

Figure 5.18 shows the t_{res}^i vs. α distribution after application of the acceptance cuts given in Equations 5.2 and 5.3 with the parameters for $n_\gamma = 2000$ from Table 5.2.

For convenience, the following nomenclature will be used throughout subsequent discussion of this analysis:

- *accepted PMT* - any triggered PMT that pass the acceptance cuts and falls in the acceptance region
- *accepted-scattered PMT* - an accepted PMT that is also innerAV-scattered, i.e. triggered by a photon that was scattered exactly once in the detection medium
- *accepted photon* - a photon that went on to trigger an accepted PMT, regardless of which optical processes the photon was subject to. Each accepted PMT is associated with a single accepted photon.
- *accepted-scattered photon* - an accepted photon that scattered exactly once in the detection medium, i.e. that triggered an accepted-scattered PMT. Each accepted-scattered PMT is associated with a single accepted-scattered photon.
- *scattered photon* - a photon that was scattered exactly once in the detection medium. Due to the imperfect efficiency of the acceptance cuts, not all scattered photons are accepted-scattered photons.

Figure 5.19 shows the distribution of d - the distance that a photon travels before undergoing Rayleigh scattering - for all accepted-scattered photons. Also shown is the pre-acceptance cuts distribution of d for scattered photons, that was previously shown as “off-axis” in Figure 3.3. As expected given the efficiency of the acceptance region, there is a general reduction in the number of photons across the entire distribution,




Figure 5.18: The time residual (defined in Equation 5.1) vs. angle α for all triggered PMTs that pass the acceptance cuts given by Equations 5.2 and 5.3 with the optimised parameters given in Table 5.2, in a single simulation of 5000 laser head pulses emitted in an angular beam at a wavelength of 405nm from fibre FS025 into a water-filled detector, with 2000 photons per pulse. The PMTs are colour-coded according to the optical process(es) each photon was subject to before it triggered the PMT.

but this is not constant across all values of d , with the greatest reduction occurring at very small and very large d - i.e. photons that either scattered very soon after entering the detection medium or just before leaving it.




Figure 5.19: The distance d travelled by a photon before it undergoes Rayleigh scattering in pure water, for all scattered photons (blue) and the subset of these photons that triggered PMTs found in the acceptance region defined by Equations 5.2 and 5.3, i.e. accepted-scattered photons (red). Note that a logarithmic scale has been used.

Figure 5.20 shows the distribution of the cosine of the scattering angle θ for accepted-scattered photons, along with the pre-acceptance cut distribution previously depicted in Figure 3.5. As discussed previously in Section 3.4.1, the other scattering angle ϕ has not been dealt with any further in this analysis, as its dependence on the initial polarisation of the photons means that it cannot be measured in SNO+.

As with the distribution of d , there is an overall reduction in the number of photons across the entire range of $\cos(\theta)$, but the reduction is not equal: all photons with $\cos(\theta) \simeq 1$ (corresponding to $\theta \simeq 0^\circ$, or scattering in the forward direction) are rejected by the acceptance cuts, as expected from the previous discussion, and the reduction at other values is not consistent - resulting in the post-acceptance cuts




Figure 5.20: The cosine of the Rayleigh scattering angle θ in pure water, for all scattered photons (blue) and the subset of these photons that triggered PMTs found in the acceptance region defined by Equations 5.2 and 5.3, i.e. accepted-scattered photons (red). Note that a logarithmic scale has been used.

distribution of $\cos(\theta)$ no longer following a $[1 + A \cos^2(\theta)]$ distribution.

5.2.6 Extraction of l_{Ray}

Varying l_{Ray} Scaling Factor

Table 5.3 shows the optimised values of the four free parameters that maximise N (given by Equation 5.4), along with the corresponding values of purity and efficiency, for a single fibre and wavelength combination, fixed values of n_γ and the acceptance cut parameters, but a varying value of s_l - the l_{Ray} scaling factor used in RAT. For each s_l , the specific contributions to the acceptance region from the various PMT classifications are also given.

It can be seen that increasing the value of s_l does not affect the absolute contributions from the beamspot PMTs (which are completely eliminated by the acceptance cuts)

	s_l	0.1	0.2	0.5	0.67	1.0	1.5	2.0
	N_{acc}	2583	3967	7707	9774	13757	19771	25453
Results	N	30.121	58.828	122.968	153.465	191.601	237.174	269.502
	Purity	0.443	0.595	0.729	0.760	0.775	0.784	0.784
	Efficiency	0.070	0.071	0.070	0.069	0.069	0.069	0.070
	Beamspace	0 (0.00)	0 (0.00)	0 (0.00)	0 (0.00)	0 (0.00)	0 (0.00)	0 (0.00)
Number	AV-reflected	156 (6.04)	156 (3.93)	219 (2.84)	178 (1.82)	198 (1.44)	223 (1.13)	227 (0.89)
(Percentage)	Concentrator -reflected	21 (0.81)	27 (0.68)	59 (0.77)	72 (0.74)	100 (0.73)	112 (0.57)	145 (0.60)
in	Cavity -scattered	72 (2.79)	109 (2.75)	190 (2.47)	257 (2.63)	400 (2.91)	575 (2.91)	795 (3.12)
Acceptance	InnerAV -scattered	1143 (44.25)	2359 (59.47)	5619 (72.91)	7430 (76.02)	10661 (77.50)	15500 (78.40)	19965 (78.45)
Region	Multiple -effect	221 (8.56)	336 (8.47)	675 (8.76)	917 (9.38)	1455 (10.58)	2431 (12.30)	3375 (13.26)
	Noise	970 (37.55)	980 (24.70)	945 (12.26)	920 (9.41)	943 (6.85)	930 (4.70)	946 (3.72)

Table 5.3: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs, all as functions of varying s_l - the l_{Ray} scaling factor. Each simulation contained 5000 laser head pulses with $n_\gamma = 1500$, emitted from fibre FS025 at a wavelength of 405nm, and the acceptance region was bounded using the optimised parameters given for this fibre / wavelength / n_γ combination in Table 5.2.

or the noise PMTs (which remain stable, as they are not affected by any changes in the scattering properties of the detection medium).

Of the remaining contributions, which all increase, the number of cavity-scattered, innerAV-scattered and multiple-effects PMTs increase by much larger factors than the AV- and concentrator-reflected PMTs - this is because increasing s_l decreases the value of l_{Ray} being used in the simulation, leading to a higher probability of a photon undergoing Rayleigh scattering within the detection medium or cavity, and therefore more such photons in the fixed acceptance region.

In real data, the absolute and relative contributions from the various PMT classifications will of course not be directly known, but it is still possible to relate the (unknown in real data) value of s_l to a quantity that will be known: the total number of all PMTs that are contained in the acceptance region, denoted by N_{acc} and named the *acceptance number*. The values of N_{acc} for each value of s_l are also given in Table 5.3.

In fact, because N_{acc} will itself also be dependent on both n_γ and the number of laser pulses per SMELLIE run (both of which are likely to change during the running of the experiment), it is more robust to express s_l in terms of the fraction of all emitted photons that go on to hit PMTs in the acceptance region - this quantity will be referred to as the *acceptance fraction*, and is denoted by f_{acc} :

$$f_{acc} = \frac{N_{acc}}{n_\gamma \times \text{number of pulses}} \quad (5.5)$$

Relationship between s_l and f_{acc}

Figure 5.21 shows the relationship between s_l and f_{acc} , for 5000 laser head pulses with $n_\gamma = 1500$, emitted through SMELLIE fibre FS025 at a wavelength of 405nm.




Figure 5.21: The value of the l_{Ray} scaling factor, s_l as a function of the acceptance fraction f_{acc} , given by Equation 5.5, for 5000 laser head pulses with $n_\gamma = 1500$, emitted at a wavelength of 405nm from fibre FS025 (green points). Also shown is a second-order polynomial fit function (red line).

It can be seen that s_l and f_{acc} can be related extremely accurately by a second-order polynomial of the form:

$$s_l = A + Bf_{acc} + Cf_{acc}^2 \quad (5.6)$$

where the coefficients A , B and C are constants, and s_l is treated as the unknown quantity in order to mirror the situation that will be presented in real data. A linear function could also possibly be used, but a second-order form is more accurate and does not add any significant computational overhead to the calculation of s_l .

Tables 5.4, 5.5 and 5.6 show the values of A , B and C respectively, for all fibre / wavelength combinations, with $n_\gamma = 1500$. It can be seen that A (the value of s_l when $f_{acc} = 0$) does not vary significantly from fibre to fibre for a given wavelength, but

	375nm	407nm	446nm	495nm
FS007	-0.079511	-0.104098	-0.147190	-0.251652
FS107	-0.071176	-0.095917	-0.131842	-0.300010
FS207	-0.071240	-0.108134	-0.142307	-0.336578
FS025	-0.075012	-0.103039	-0.154546	-0.228130
FS125	-0.071257	-0.104322	-0.135608	-0.303161
FS225	-0.072536	-0.101192	-0.151321	-0.259313
FS137	-0.076210	-0.092954	-0.159300	-0.256159
FS037	-0.068190	-0.083568	-0.160266	-0.271470
FS237	-0.064922	-0.105545	-0.153284	-0.332569
FS255	-0.073156	-0.086231	-0.149821	-0.319321
FS055	-0.070612	-0.101095	-0.133995	-0.298299
FS155	-0.066657	-0.088784	-0.140747	-0.282726

Table 5.4: The coefficient A of Equation 5.6, relating s_l and f_{acc} . Each fibre / wavelength combination contained 5000 laser head pulses with $n_\gamma = 1500$, and utilised the individual optimised acceptance cut parameters given in Tables B1 to B8.

there is a trend of increasing A with increasing wavelength, caused by the previously discussed decrease in N_{acc} as the wavelength - and therefore the scattering length - increases. This increasing coefficient value against wavelength can also be seen for B and C .

It should be noted that the absolute values of A , B and C are dependent on the value used for n_γ , since the latter is directly related to f_{acc} , and therefore so is the relating function between s_l and f_{acc} .

Using the pre-calculated values of A , B and C for a given n_γ and the easily calculated f_{acc} from Equation 5.5, it is therefore straight-forward to calculate s_l using Equation 5.6.

As an example, consider the situation where the true scattering length of pure water at a wavelength of 407nm has been set to be $l_{Ray}^{true} = 162.055\text{m}$. A simulation has

	375nm	407nm	446nm	495nm
FS007	449.400	569.514	934.236	1762.11
FS107	446.196	616.339	917.302	2085.64
FS207	573.761	770.061	1100.280	2617.31
FS025	432.895	582.627	962.461	1665.45
FS125	473.602	596.824	927.840	2083.78
FS225	512.177	736.680	1103.560	2472.27
FS137	435.849	559.492	940.437	1853.67
FS037	480.233	587.968	1035.000	2039.69
FS237	506.303	665.741	1036.150	2426.15
FS255	466.007	580.403	917.192	2149.28
FS055	494.504	670.137	1021.170	2353.52
FS155	503.900	652.852	1045.700	2366.99

Table 5.5: The coefficient B of Equation 5.6, relating s_l and f_{acc} . Each fibre / wavelength combination contained 5000 laser head pulses with $n_\gamma = 1500$, and utilised the individual optimised acceptance cut parameters given in Tables B1 to B8.

	375nm	407nm	446nm	495nm
FS007	9865.0	10995.8	15197.6	227171.0
FS107	18718.9	11860.6	38681.4	55278.1
FS207	23812.8	7617.2	46982.0	17634.6
FS025	14650.9	15960.4	19801.0	299230.0
FS125	15403.2	14182.0	42020.9	61936.2
FS225	26217.3	17895.1	37274.3	201127.0
FS137	13218.8	22925.5	10191.3	227556.0
FS037	16864.0	20783.8	-2250.7	204354.0
FS237	22416.9	18666.6	34930.5	-46757.7
FS255	14275.6	22308.3	25035.2	-58407.5
FS055	19617.4	19025.7	45208.1	99279.9
FS155	21749.7	18094.4	24343.3	83768.6

Table 5.6: The coefficient C of Equation 5.6, relating s_l and f_{acc} . Each fibre / wavelength combination contained 5000 laser head pulses with $n_\gamma = 1500$, and utilised the individual optimised acceptance cut parameters given in Tables B1 to B8.

been performed using this true scattering length, with 5000 laser head pulses of 1500 photons each being emitted from SMELLIE fibre FS025. Applying the acceptance cuts (Equations 5.2 and 5.3) using the optimised cut parameters for this fibre / wavelength / n_γ combination as given in Table B3 leads to $N_{acc} = 11521$, and therefore $f_{acc} = 0.00153613$. Using the coefficients given in Tables 5.4, 5.5 and 5.6, this leads to a calculated scaling factor of $s_l = 0.8296 \pm 0.0175$ (the error having been propagated through from the errors in A , B , C and N_{acc}). In RAT, the assumed scattering length of water at 407nm with $s_l = 1.0$ is $l_{Ray}^{s=1} = 195.483\text{m}$, and so applying the calculated s_l to this leads to a calculated l_{Ray}^{calc} of $(162.176 \pm 3.412)\text{m}$ - completely consistent with l_{Ray}^{true} .

The results of the study show that such a technique can indeed be used successfully to find s_l and l_{Ray} to a very high precision, but it does not give any information regarding the scattering angle θ - this requires a more detailed analysis of individual photons.

5.3 General Analysis Technique

5.3.1 Introduction

Knowing the position \underline{s} at which a photon scattered is fundamental to any study of the detection medium's scattering properties, as both d and θ can be easily calculated using the scattering position together with the known initial and final photon positions - the SMELLIE fibre and triggered PMT positions respectively. Therefore, the first step in any analysis must be to reconstruct \underline{s} using only the information available from the detector readout, as shown in Figure 5.22. The only initial bounds on \underline{s} are that it must be somewhere within the detection medium, and also somewhere within the cone of the SMELLIE beam.




Figure 5.22: An illustration of the SNO+ detector (not to scale) showing a SMELLIE fibre located at position \underline{f} (with respect to the centre of the detector). A zero angular-width pencil beam is depicted by the orange line, entering the detection medium at position \underline{p}_1 and exiting it at position \underline{p}_2 , and a SMELLIE beam of non-zero angular width is indicated by the yellow shaded region. A photon emitted from the fibre can potentially scatter at some position \underline{s} . The various other times and positions shown here are discussed in the main text.

5.3.2 Path Time Calculation

Consider a photon that is emitted at time t_0 from the SMELLIE fibre located at position \underline{f} with respect to the detector centre. (Note that t_0 is *not* equal to 0ns, since that corresponds to the issue of the global trigger from the SNO+ MTC as discussed in Sections 1.1.2 and 4.1.2, which must occur some time before the emission of light

from the fibre.) If the photon travels directly to a beamspot PMT, the resulting photoelectron is incident on the PMT's first dynode at a time t_{di} given by:

$$t_{di} = t_0 + t_{dp,true} + t_t \quad (5.7)$$

where $t_{dp,true}$ is the true time of the direct path across the detector and t_t is the photoelectron's transit time from photocathode to dynode as previously discussed in Section 5.2.1.

As noted previously, the calibrated PMT trigger time (which is the actual observable available from the detector readout, not t_{di}) is the sum of t_{di} and the cable delay between the PMT and the electronics (denoted by t_c), and so the calibrated PMT trigger time t_{dc} of this particular beamspot PMT can be expressed as:

$$\begin{aligned} t_{dc} &= t_{di} + t_c \\ &= t_0 + t_{dp,true} + t_t + t_c \end{aligned} \quad (5.8)$$

Rearranging this equation therefore yields a formula for calculating the total time difference between the calibrated PMT trigger time and the time of the direct path across the detector:

$$t_{dc} - t_{dp,true} = t_0 + t_t + t_c \quad (5.9)$$

For a beamspot PMT (and only for a beamspot PMT), $t_{dp,true} = t_{flight}^i$ as calculated using RAT's Lightpath class, and so the left-hand side of Equation 5.9 is equivalent

to the right-hand side of Equation 5.1 - Equation 5.9 describes the time residual of beamspot PMTs, $t_{res}^i|_{sp}$. Figure 5.23 shows the distribution of $t_{res}^i|_{sp}$, which (as expected) has a distribution almost identical to that of the calibrated beamspot PMT trigger times previously shown in Figure 5.8. Fitting a Gaussian function to the distribution shown in Figure 5.23 gives a mean $\mu_{res,sp} = 251.97\text{ns}$ and a standard deviation $\sigma_{res,sp} = 1.60\text{ns}$. This value of $\mu_{res,sp}$ is indeed consistent with the distribution of calibrated beamspot PMT trigger times, which can be seen to peak at $\sim 330\text{ns}$, along with a simple calculation of the flight time ($\sim 80\text{ns}$) of a photon travelling in a straight line across the entire detector.




Figure 5.23: The time residuals, calculated using Equation 5.1, of beamspot PMTs - those triggered by photons which did not reflect or scatter anywhere in the detector as they travelled between the SMELLIE fibre and the PMT.

Now consider a photon that is emitted at the same time t_0 from the same SMELLIE fibre, but goes on to scatter at some position \underline{s} in the detection medium before creating a photoelectron that hits the first dynode of an innerAV-scattered PMT at time t_{si} . The time that the photon travelled for between emission and scattering is

denoted by $t_{bs,true}$, and $t_{as,true}$ is the time between scattering and hitting the PMT's photocathode. Therefore, t_{si} can be expressed as:

$$t_{si} = t_0 + t_{sp,true} + t_t \quad (5.10)$$

where $t_{sp,true}$ is the true total time of the scattered path across the detector, equal to:

$$t_{sp,true} = t_{bs,true} + t_{as,true} \quad (5.11)$$

and t_t is the photoelectron's transit time as before. The calibrated PMT trigger time t_{sc} is given by:

$$\begin{aligned} t_{sc} &= t_0 + t_{sp,true} + t_t + t_c \\ &= t_{sp,true} + t_{res}^i|_{sp} \end{aligned} \quad (5.12)$$

using Equation 5.9. Rearranging this equation, and using $\mu_{res,sp}$ as a representative value of the $t_{res}^i|_{sp}$ distribution, gives an expression for the true total path time of a photon that scattered once in the detection medium:

$$t_{sp,true} = t_{sc} - \mu_{res,sp} \quad (5.13)$$

where both terms on the right-hand side can be extracted directly from the detector readout.

5.3.3 Scattering Position Reconstruction

The two equations 5.11 and 5.13 together provide a way to reconstruct the entire path of a scattered photon - from the fibre at position \underline{f} , to the scattering position \underline{s} , and on to the innerAV-scattered PMT located at \underline{p}_s .

Consider the hypothesis that this photon underwent Rayleigh scattering at a position \underline{s}_{trial} . The `Lightpath` class in RAT can be used to calculate the path that the photon travelled along before scattering - from \underline{f} to \underline{s}_{trial} , and after scattering - from \underline{s}_{trial} to \underline{p}_s . From these, the respective path times $t_{bs,trial}$ and $t_{as,trial}$ can be calculated, and using Equation 5.11, the total path time $t_{sp,trial}$ for this scattering hypothesis can be found. If this value of $t_{sp,trial}$ matches the value of $t_{sp,true}$ obtained for this PMT using Equation 5.13 (independent of the true and hypothesised scattering positions), then this suggests that the hypothesised scattering position \underline{s}_{trial} is in fact the same as the true scattering position \underline{s} . On the other hand, if $t_{sp,trial}$ and $t_{sp,true}$ do not have the same value, then a difference choice of \underline{s}_{trial} must be made.

The reconstruction of \underline{s} can therefore be considered as the minimisation of Δt - the absolute difference between $t_{sp,trial}$ and $t_{sp,true}$:

$$\Delta t = |t_{sp,trial} - t_{sp,true}| \quad (5.14)$$

through the variation of \underline{s}_{trial} .

There are two approximations that have been implicitly made within the reconstruction method outline above, which could potentially affect its overall performance.

A photon can in principle be incident anywhere on the 20cm-diameter PMT face, but the exact location of incidence cannot be recorded by the SNO+ detector - the position of the triggered PMT is given by the centre of the PMT's face. Therefore, a photon path's endpoint as extracted from the detector readout can only be accurate

to within 10cm of the true endpoint. A simple straight-line approximation (ignoring refraction) shows that even for the extreme (and non-realistic) case in which a photon scattered at one side of the detection medium and was then incident on a PMT located on the other side of the medium, a change in the path's endpoint of 10cm gives an increase in the path length of only 0.338mm, and therefore an increase in the post-scattering path time t_{as} of ~ 1.51 ps in pure water. In comparison, the path from fibre to a beamspot PMT takes approximately 80ns, and therefore the change in the total path time due to the approximated photon incidence position is negligible.

As seen in Figure 5.23 above, there is a spread in $t_{res|dir}^i$ around the mean value. This spread is caused by the distribution of photoelectron transit times previously shown in Figure 5.9, and though it is two orders of magnitude smaller than the peak height of the $t_{res|dir}^i$ distribution, it is certainly not negligible: given the large number of scattered photons remaining even after application of the acceptance cuts, there will be a corresponding spread of tens of nanoseconds in the calculated values of $t_{ps,true}$ from Equation 5.13. This is comparable in magnitude to the ~ 80 ns path time from fibre to beamspot PMT, and so could result in large discrepancies between the reconstructed and true scattering positions - for example, if the true value of $t_{res|dir}^i$ is just one $\sigma_{res,sp}$ away from $\mu_{res,sp}$, this leads to a difference between the reconstructed and true total path lengths (and scattering positions) of approximately 35cm, and similarly a deviation of three $\sigma_{res,sp}$ (covering the $t_{res|dir}^i$ distribution's highest peak) will give a difference between reconstructed and true scattering positions of 1m.

5.4 Pencil Beam Analysis

5.4.1 Reconstruction Method

To test the fundamental performance of the general scattering position reconstruction method outlined above, the simple case of a pencil beam can be used - i.e. photons

are emitted from the SMELLIE fibre with no angular distribution, and therefore they all have the same initial direction when entering the detection medium, as depicted on Figure 5.22. The true scattering position \underline{s} of any scattered photons must therefore lie along the line between \underline{p}_1 (the entry point to the detection medium) and \underline{p}_2 (the exit point from the medium of any photons that did not undergo any scattering or reflections), and so the hypothesised scattering position can be formulated as:

$$\underline{s}_{trial} = \underline{p}_1 + d_{trial} \left(\frac{\underline{p}_2 - \underline{p}_1}{|\underline{p}_2 - \underline{p}_1|} \right) \quad (5.15)$$

where d_{trial} is a trial value of the previously defined d - the distance that the photon travelled before scattering, i.e. a scalar quantity taking values (in millimetres) between 1.0 and $|\underline{p}_2 - \underline{p}_1|$ - the maximum distance between \underline{p}_1 and \underline{p}_2 .

The minimisation of Δt as given in Equation 5.14 is therefore done with respect to d_{trial} , using a simple variable step-size line search.

The initial value of d_{trial} is taken as $0.5 |\underline{p}_2 - \underline{p}_1|$, i.e. exactly halfway between \underline{p}_1 and \underline{p}_2 , and this gives \underline{s}_{trial} and Δt for the current value of d_{trial} - denoted by $\underline{s}_{trial}^{curr}$ and Δt_{trial}^{curr} respectively. The value of d_{trial} is immediately increased by some step size Δd , which is initially set at +500mm for this first iteration of the search, but in subsequent iterations can be either positive (increasing d_{trial} , moving \underline{s}_{trial} away from \underline{p}_1) or negative (decreasing d_{trial} , moving \underline{s}_{trial} closer to \underline{p}_1). A new Δt_{trial}^{curr} is calculated, and the value from the previous iteration is stored as Δt_{trial}^{prev} .

The exact magnitude and sign of Δd in the current iteration of the search, and therefore the values of d_{trial} and \underline{s}_{trial} in the next iteration, is determined based on Δt_{trial}^{curr} and Δt_{trial}^{prev} , as detailed in Table 5.7. Using these rules, the trial values of d_{trial} oscillate around the true value of d with decreasing amplitude, and so Δt_{trial} also decreases.

Δt_{trial}^{curr}	Δt_{trial}^{prev}	Δd		Effect on \underline{s}_{trial}
		magnitude	sign	
< 0	< 0	unchanged	+	moves away from \underline{p}_1 by an unchanged Δd
< 0	> 0	$\div 10$	+	moves away from \underline{p}_1 by a smaller Δd
> 0	> 0	unchanged	-	moves towards \underline{p}_1 by an unchanged Δd
> 0	< 0	$\div 10$	-	moves towards \underline{p}_1 by a smaller Δd

Table 5.7: The rules for determining the step size Δd that is used in the current iteration, and the resulting change on the trial scattering position \underline{s}_{trial} for the next iteration, based on the relative values of the time difference Δt_{trial} as calculated using Equation 5.14 for the current and previous values of \underline{s}_{trial} .

The search ends either when Δt_{trial} decreases below a tolerance value of 0.0001ns (this being the maximum precision of the Lightpath class), or when 500 iterations of the search have been performed, at which point the current values of d_{trial} and \underline{s}_{trial} are taken as d_{reco} - the reconstructed distance that the photon travelled before scattering, and the reconstructed scattering position \underline{s}_{reco} respectively.

5.4.2 Reconstruction Performance I: s

Four scenarios have been considered when testing the performance of the reconstruction method, based on the previously discussed assumptions about the position of each triggered PMT and the path time of each scattered photon:

1. using the true position of incidence on the PMT (denoted as the *true PMT position*), and the *true path time* of the scattered photon
2. using the position at the centre of the PMT face (the *central PMT position*), and the true path time
3. using the true PMT position and the *calculated path time* from Equation 5.13

4. using the central PMT position and the calculated path time

The method's performance in each scenario can be characterised by the reconstruction *bias*, μ , and *resolution*, σ , given by:

$$\mu = \sqrt{\mu_x^2 + \mu_y^2 + \mu_z^2} \quad (5.16)$$

$$\sigma = \sqrt{\sigma_x^2 + \sigma_y^2 + \sigma_z^2} \quad (5.17)$$

with μ_q and σ_q being the mean and standard deviation respectively of a Gaussian fit to the distribution of $(q_{reco} - q_{true})$ for each coordinate q ($= x, y$ or z) of the scattering position. Table 5.8 gives the calculated bias and resolution of each reconstruction scenario for the previously used combination of fibre FS025 and light at a wavelength of 405nm, and the results for only scenarios (1) and (4) across all fibre / wavelength combinations are given in Tables C1 (bias) and C2 (resolution).

Reconstruction Scenario		\underline{s} Bias	\underline{s} Resolution
PMT Position	Path Time	(mm)	(mm)
True	True	+ 0.468	1.027
Central	True	+ 0.829	12.050
True	Calculated	+ 36.706	277.591
Central	Calculated	+ 35.345	279.318

Table 5.8: The bias and resolution, calculated using Equations 5.16 and 5.17 respectively, of \underline{s}_{reco} (the reconstructed scattering position) of accepted-scattered photons emitted in a pencil beam from fibre FS025 at a wavelength of 405nm, in the four reconstruction scenarios discussed in the main text.

In all fibre / wavelength combinations, scenario (1) produces the smallest bias and resolution - with the former varying between just 0.3 and 0.4mm, and the latter consistently around ~ 1 mm, proving that the underlying principle of the reconstruction method is effective: given the correct end position and path time of the accepted-scattered photon, the entire path itself can be reconstructed to within just a few millimetres.

Scenario (2) shows that, as discussed previously, using the position at the centre of the PMT instead of the true incidence position of the photon on the PMT does not have a large effect on the reconstruction performance. The poorer resolution is due to the maximum $(q_{reco} - q_{true})$ of any of the coordinates now being 10cm, instead of the 1cm maximum that was seen in scenario (1) - this is consistent with the idea that a photon can be incident on a PMT at most 10cm from the centre of its face. Even so, the bias and resolution both remain very small compared to the detector itself.

Comparing scenarios (1) and (3) shows that the use of the calculated path time in place of the true path time - and thereby incorporating the distribution of the photoelectron transit times t_t into the reconstruction - leads to a considerable loss in performance. As discussed previously, even a small difference between $\mu_{res,sp}$ used in Equation 5.13 and the true value of $t_{res|dir}^i$ leads to a large change in the reconstructed scattering position.

Scenario (4) is the situation that will be available in real data from the SNO+ detector, and it can be seen that the reconstruction method in this scenario performs very similarly to scenario (3) - as expected since it has already been shown that the loss in performance due to using an approximate photon incidence position is insignificant when compared to the loss incurred from using a calculated path time. Nevertheless, the performance in scenario (4) is still fairly good across all fibre / wavelength combinations - with a bias of between 30 and 40mm (compared to the detection medium's diameter of 12m) and a resolution of no more than 300mm - a

similar length scale to the diameter of a single PMT.

5.4.3 Reconstruction Performance II: d

In a similar way to \underline{g} , the reconstruction bias and resolution can be calculated for d by fitting a Gaussian function to the distribution of $(d_{reco} - d_{true})$, and taking the mean and standard deviation respectively. These values are given in Table 5.9 for the single fibre / wavelength combination and four reconstruction scenarios discussed previously. Tables C3 and C4 give the bias and resolution respectively for scenarios (1) and (4) across all fibre / wavelength combinations.

Figures 5.24 and 5.25 show the comparison between the true and reconstructed d for each accepted-scattered photon in scenarios (1) and (4) respectively, in the FS025 / 405nm combination.

Reconstruction Scenario		d Bias	d Resolution
PMT Position	Path Time	(mm)	(mm)
True	True	+ 0.354	0.742
Central	True	- 0.832	11.980
True	Calculated	+ 29.755	254.481
Central	Calculated	+ 28.824	257.174

Table 5.9: The bias and resolution of d_{reco} (the reconstructed distance travelled before a photon undergoes Rayleigh scattering) of accepted-scattered photons emitted in a pencil beam from fibre FS025 at a wavelength of 405nm, in the four reconstruction scenarios discussed in the main text.

The reconstruction of d in scenario (1) is, as expected, extremely good overall, with biases and resolutions both of no more than 1.0mm in scenario (1) across all fibres and wavelengths. There are almost no accepted-scattered photons with $d_{reco} < d_{true}$, and only a small number that reconstruct with $d_{reco} > d_{true}$. In the latter case, the




Figure 5.24: The true and reconstructed d (the distance travelled before a photon undergoes Rayleigh scattering) of accepted-scattered photons emitted in a pencil beam from fibre FS025 at a wavelength of 405nm, when the true position of incidence on the PMT and the true total path time of the photon have been used in the reconstruction method. The block-like pattern is an artifact of the binning.




Figure 5.25: The true and reconstructed d (the distance travelled before a photon undergoes Rayleigh scattering) of accepted-scattered photons emitted in a pencil beam from fibre FS025 at a wavelength of 405nm, when the position at the centre of the PMT's face and the calculated total path time of the photon have been used in the reconstruction method.

corresponding accepted-scattered PMTs are all in fact located either near the AV neck or at the very bottom of the PSUP. However, it should be noted that having an accepted-scattered PMT at these positions does not guarantee that the corresponding photon's d will always mis-reconstruct - for most of the accepted-scattered PMTs located at very low or very high z , the reconstruction of d is performed extremely well. The same mis-reconstruction has been observed with other simulated SMELLIE fibres at different nodes and emission angles, and so a likely common cause is an intermittent error in the calculation of the refracted path of the photon by the RAT Lightpath class when dealing with PMTs located at very low or very high z .

This mis-reconstruction is no longer an issue in scenario (4) - although the general trend clearly still follows $d_{reco} \simeq d_{true}$, there are now many more accepted-scattered photons with d_{reco} significantly different from d_{true} , as expected due to the use of $\mu_{res,sp}$ in place of the true $t_{res}^i|_{dir}$. There is a larger population of photons with $d_{reco} \gg d_{true}$ than with $d_{reco} \ll d_{true}$, thereby accounting for the clear positive bias of this scenario in Table 5.8, and there is also a small number of photons with $d_{reco} \sim 0\text{mm}$, regardless of the value of d_{true} .

Nevertheless, the reconstruction of d even in scenario (4) is still performed very well: compared to the scales of both the detector and the expected Rayleigh scattering length of $\sim 195\text{m}$ at a wavelength of 400nm , the d_{reco} bias of between 20 and 40mm and resolution of 250 to 300mm are both very small.

5.4.4 Reconstruction Performance III: θ

The reconstruction of θ is trivial once \underline{s}_{reco} has been calculated, since (as seen in Figure 3.1) the angle is dependent only on the pre-scattering and post-scattering directions of the photon.

As with d , the reconstruction bias and resolution can be calculated for θ by finding the mean and standard deviation of a Gaussian fit to the distribution of $(\theta_{reco} -$

θ_{true}). These values are given for FS025 / 405nm combination in Table 5.10 in the four reconstruction scenarios discussed previously, and in Tables C5 (bias) and C6 (resolution) for all fibre / wavelength combinations in scenarios (1) and (4).

Reconstruction Scenario		θ Bias	θ Resolution
PMT Position	Path Time	($^{\circ}$)	($^{\circ}$)
True	True	+ 0.146	0.636
Central	True	- 0.003	0.757
True	Calculated	+ 0.265	2.054
Central	Calculated	+ 0.221	2.167

Table 5.10: The bias and resolution of θ_{reco} (the reconstructed Rayleigh scattering angle θ) of accepted-scattered photons of accepted-scattered photons emitted in a pencil beam from fibre FS025 at a wavelength of 405nm, in the four reconstruction scenarios discussed in the main text.

A similar situation to that of \underline{g} is observed in the reconstruction of θ . Scenarios (1) and (2) give the best reconstruction performance - biases of $\sim 0.1^{\circ}$ and resolutions of $\sim 0.7^{\circ}$ - due to their use of the true total path time, and their similarity shows that the use of the central PMT position does not have a large effect on θ reconstruction. Introducing $\mu_{res,sp}$ in place of the true $t_{res}^i|_{dir}$, as in scenarios (3) and (4), results in a slight decrease in reconstruction performance - with biases and resolutions of $< 1.0^{\circ}$ and $< 3.0^{\circ}$ respectively, with θ_{reco} being calculated at up to 57° more than θ_{true} , but only at the level of 1 or fewer photons in every 10^3 .

Figures 5.26 and 5.27 show the comparison between the true and reconstructed θ in scenarios (1) and (4) respectively for accepted-scattered photons in the FS025 / 405nm combination. The same population of mis-reconstructed photons that was noted in scenario (1) of the d reconstruction is also present here, and Figure 5.26 shows that these photons seem to preferentially reconstruct at $\theta \simeq 57^{\circ}$ and $\theta \simeq 114^{\circ}$. This mis-reconstruction is qualitatively expected due to the interdependence of \underline{g} and

θ , but the fact that θ takes very specific values when mis-reconstructing may imply the existence of a “default” direction within the Lightpath class which is used when the refracted path cannot be calculated - such as if errors are encountered when using PMTs near the neck or the bottom of the PSUP.

In a similar way to the reconstruction of d , there is a larger population of photons with $\theta_{reco} \gg \theta_{true}$ than with $\theta_{reco} \ll \theta_{true}$ in scenario (4), and this is related to the analogous observation in d : an overestimated d_{reco} will naturally lead to an overestimated θ_{reco} .

5.4.5 Reconstruction of Other Accepted Photons

As noted in Section 5.2.5, due to the imperfect purity of the acceptance cuts, there will be a certain fraction of other photons - denoted hereafter as *accepted-other* photons - present in the acceptance region alongside the accepted-scattered ones. The reconstruction method discussed above should not in principle be applied to such photons, since by definition they did not scatter exactly once in the detection medium, but of course in practice (when applying the analysis to real data), the method will be blind to the distinction between accepted-scattered and accepted-other photons - and therefore it is important to consider how the latter reconstruct.

Figures 5.28 and 5.29 show the reconstructed values of d and θ respectively, for all accepted-scattered photons in scenario (4) using 5000 pulses of light with $n_\gamma = 1500$ and a single fibre / wavelength combination, as well as the the separate contributions from the various accepted-other photons that underwent different optical processes between emission and triggering the corresponding accepted PMT.

It can be seen that there are certain values of both variables at which the various types of accepted-other photons are much more likely to reconstruct than others.

The majority of cavity-scattered photons reconstruct with $80 < d_{reco} < 100\text{mm}$,




Figure 5.26: The true and reconstructed θ (the Rayleigh scattering angle) of accepted-scattered photons emitted in a pencil beam from fibre FS025 at a wavelength of 405nm, when the true position of incidence on the PMT and the true total path time of the photon have been used in the reconstruction method. The block-like pattern is an artifact of the binning.




Figure 5.27: The true and reconstructed θ (the Rayleigh scattering angle) of accepted-scattered photons emitted in a pencil beam from fibre FS025 at a wavelength of 405nm, when the position at the centre of the PMT's face and the calculated total path time of the photon have been used in the reconstruction method.




Figure 5.28: The reconstructed d of all accepted photons in a simulation of 5000 laser head pulses with $n_\gamma = 1500$, emitted at a wavelength of 405nm from fibre FS025, separated by the optical process(es) each photon was subject to before triggering the corresponding accepted PMT. The summation of all contributions is shown by the dashed line. The reconstruction method used the PMT's central position and the photon's calculated total path time.




Figure 5.29: The reconstructed scattering angle θ of all accepted photons in a simulation of 5000 laser head pulses with $n_\gamma = 1500$, emitted at a wavelength of 405nm from fibre FS025, separated by the optical process(es) each photon was subject to before triggering the corresponding accepted PMT. The summation of all contributions is shown by the dashed line. The reconstruction method used the PMT's central position and the photon's calculated total path time.

with only a few having $d_{reco} \gg 100\text{mm}$. This is due to the underlying assumptions made by the reconstruction method: as noted in Equation 5.15, the method has been limited such that it only searches for potential scattering positions within the detection medium, and therefore, a photon that truly scattered outside the medium, in the cavity or AV, will be reconstructed just a few millimetres inside it - either at very low or very high d , in order to match the calculated and expected total photon path times as closely as possible. Additionally, the overall shape of the cavity-scattered photons' θ_{reco} distribution, which has contributions predominantly at $\sim 0.5\text{rads}$ and $\sim 3.0\text{rads}$, is consistent with the distribution of angle α for the accepted cavity-scattered photons, shown previously in Figure 5.18.

Almost all of the AV-reflected photons remaining in the acceptance region reconstruct with $\theta_{reco} > 3.0\text{rads}$, corresponding to one of two large d_{reco} peaks: $d_{reco} \sim 500\text{mm}$ or $d_{reco} > 11500\text{mm}$. These values of d_{reco} and θ_{reco} are as expected, since photons that reflected back from any of the AV surfaces will be treated as though they scattered almost directly backwards ($\theta_{reco} \sim \pi$), and - as noted for the cavity-scattered photons - the closest matching path time will be found at values of d_{reco} very close to the AV surfaces, i.e. either edge of the detection medium.

The reconstruction of multiple-effects photons exhibits some very interesting behaviour. As well as the expected underlying flat distribution of this quantity, a large number of these photons can be found in two populations: ($d_{reco} > 11500\text{mm}$, $\theta_{reco} > 3.0\text{rads}$) - caused by photons undergoing only reflections from exactly two AV surfaces, and ($d_{reco} \simeq 9250\text{mm}$, $\theta_{reco} \simeq 1.5\text{rads}$) - due to photons that first undergo a reflection at the first AV surface encountered and are subsequently reflected from a PMT concentrator. These peaks are in fact caused by a combination of using a pencil beam and the deterministic nature of reflections as they are simulated in RAT: both the amount of reflected light and the direction of the photon after reflection are completely calculable using the Fresnel equations that describe the behaviour of

light as it moves between different media. Therefore, once it has been determined that a particular photon will reflect at, for example, an AV surface, it will take the same path across the detector as any other photon that also reflected at that same surface (since in a pencil beam, all photons begin with the same initial direction), and this then leads to specific combinations of reflections - from AV surfaces and/or PMT concentrators - being reconstructed with specific values of d_{reco} and θ_{reco} .

Many of the other fibre / wavelength combinations exhibit broadly similar behaviour to that shown in Figures 5.28 and 5.29 - for example, the distribution of cavity-scattered photons at very low d_{reco} is always present, as are the multiple-effects and AV-reflected photons at high d_{reco} .

However, there are some details that are more fibre-specific - for example, fibres FS007 and FS025 both show the same broad multiple-effects peak centred at approximately $d_{reco} \sim 9000\text{mm}$, regardless of wavelength, but this is not seen elsewhere, and although all fibres have at least one very prominent peak of multiple-effects and/or AV-reflected photons in their distribution of θ_{reco} , the exact values of these peaks varies greatly between fibres (but is fairly stable with wavelength). The fibre's emission direction also affects the highest possible value of d_{reco} : for the 0° fibres, this is very close to 12000mm , but decreases to $\sim 10000\text{mm}$ at 10° and $\sim 8000\text{mm}$ for the fibres at 20° - simply due to the shortening of the longest possible path across the detection medium at higher emission angles.

In addition, as the wavelength increases, there is a reduction in the number of innerAV-scattered photons in comparison to the accepted-other photons, which is caused by the larger scattering length at higher wavelengths as previously discussed in Section 5.2.4.

5.5 Angular Beam Analysis

5.5.1 Introduction

In principle, the same analysis as that discussed for the pencil beam can still be applied to a non-zero angular-width beam such as that produced by the real-world SMELLIE hardware - that is, the scattering position \underline{s} of each photon can be reconstructed using the known trigger time of the PMT that the photon was incident on, as well as the known emission location, and from this the values of d and θ can then be reconstructed. However, the step from a one-dimensional search along the pencil beam to a three-dimensional search within the cone of the angular beam requires that the method used previously for the reconstruction of \underline{s} be modified somewhat.

5.5.2 Application of Pencil Beam Analysis

To understand how much of an effect going from a pencil beam to an angular beam has on the ability to reconstruct the scattering properties of a photon, the entire pencil beam analysis developed in Section 5.4 has, without modification, been directly applied to a SMELLIE beam simulated with its full angular width. The resulting reconstruction biases and resolutions are given in Table 5.11 for \underline{s} and d , and Table 5.12 for θ , in two of the reconstruction scenarios previously used:

1. using the true position of incidence on the PMT (the *true PMT position*), and the *true path time* of the scattered photon
4. using the position at the centre of the PMT face (the *central PMT position*), and the *calculated path time* from Equation 5.13

It can be seen that the reconstruction of both \underline{s} and d suffer significant losses in performance, with the resolution of both now clearly exceeding 500mm in scenario 4 (the “real data” scenario). This is simply due to how the analysis works: although

Reconstruction Scenario		\underline{s}		d	
PMT Position	Path Time	Bias (mm)	Resolution (mm)	Bias (mm)	Resolution (mm)
True	True	+ 17.078 (+ 0.468)	1012.380 (1.027)	- 18.044 (+ 0.354)	472.802 (0.742)
Central	Calculated	+ 81.800 (+ 35.345)	1159.860 (279.318)	+ 33.503 (+ 28.824)	580.691 (257.174)

Table 5.11: The bias and resolution of \underline{s}_{reco} (the reconstructed scattering position) and d_{reco} (the reconstructed distance travelled before a photon undergoes Rayleigh scattering) of accepted-scattered photons emitted in an angular beam but analysed using the unmodified pencil beam method, in the two reconstruction scenarios discussed in the main text. In brackets for comparison are the previously calculated corresponding values when using a pencil beam, taken from Tables 5.8 and 5.9.

Reconstruction Scenario		θ	
PMT Position	Path Time	Bias (°)	Resolution (°)
True	True	+ 0.106 (+ 0.146)	1.237 (0.636)
Central	Calculated	+ 0.120 (+ 0.221)	2.405 (2.167)

Table 5.12: The bias and resolution of θ_{reco} (the reconstructed Rayleigh scattering angle θ) of accepted-scattered photons emitted in an angular beam but analysed using the unmodified pencil beam method, in the two reconstruction scenarios discussed in the main text. In brackets for comparison are the previously calculated corresponding values when using a pencil beam, taken from Table 5.10.

it is still indeed doing the task that it is designed to - that is, finding a location at which the calculated total path time is equal to the true total path time - an angular beam differs from a pencil beam in that there are a number of such locations instead of a single, unique one. This is illustrated qualitatively in Figure 5.30.

The pencil beam analysis inherently assumed that \underline{s} lies along the centre of the angular beam, and so will always find it as such (the scattering position of the orange




Figure 5.30: An illustration of the SNO+ detector (not to scale) showing three paths (purple, orange and green) that a scattered photon could have taken after being emitted from a SMELLIE fibre with an angular beam (shown by the yellow shaded region). Each path has a different scattering position \underline{s} , but all three paths would take the same total time to travel along.

path in Figure 5.30), when in reality the true \underline{s} is almost certainly not located in the centre of the beam (the scattering positions of the purple and green paths). Depending on the distance from fibre to \underline{s} , this will then lead to the observed large discrepancies between the true and calculated \underline{s} and d - with the differences becoming larger as d_{true} increases due to the increasing beam width, and thus the increasing number of possible off-centre scattering locations.

Although the θ reconstruction performance also decreases relatively speaking, the absolute performance is still very good for the angular beam - with a bias of $< 1^\circ$ and resolution of $< 4^\circ$ even in the real-data scenario.

5.5.3 Possible Methods for Improvement

In principle, two paths with the same starting and ending positions and total time along path could still possibly be distinguished by applying constraints on their starting and ending directions - for example, the emission direction from the SMELLIE fibre, or the angle at which the photon is incident on the final triggered PMT.

There are only two remaining observables associated with each PMT: the charges QHS and QHL (previously described in Section 1.1.2), but these are solely dependent on the behaviour and performance of the dynode stack, which in turn only requires that an electron be emitted from the PMT's photocathode - a process that will occur regardless of the photon's original incidence angle. Therefore, there is no way to calculate the incidence angle of a photon on a SNO+ PMT purely from the remaining available PMT-related observables, and so unfortunately the photon's path cannot be constrained any further at its endpoint.

However, the fibre's individual angular intensity profile (previously discussed in Section 4.2.5) may provide a way to constrain the photon initial emission direction, effectively reducing the angular beam to a pencil beam for a single photon, but where the beam's direction is not necessarily the same as the fibre's forward direction. By treating the intensity profile as the probability distribution of the photon's emission angle from the fibre, and then sampling from this distribution for each photon to be reconstructed, the direction of the path before scattering is constrained to a single direction. Of course, at the level of single photons, randomly sampling from the emission angle distribution is certainly not a reliable method of finding the true emission angle, but over many hundreds of thousands of photons (such that will be

present in a standard SMELLIE data run), the emission angle distribution will be recovered. In this way, the reconstruction method can take into account the angular distribution of photons, at least on a basic level and over many photons.

Of course, fundamentally the pencil beam method has been designed and optimised to work for a very specific 0° opening angle approximation of an angular beam. A much larger improvement in reconstruction performance would likely come from implementing a new method - perhaps based around similar principles, but one that has been built specifically to reconstruct the scattering properties of photons emitted in any angular beam, regardless of whether its angular distribution has a non-zero width or not.

5.6 Systematic Uncertainties

As noted throughout the previous sections, this analysis is based purely on simulation (due to a simple lack of enough real SMELLIE data), and so there has been no quantitative discussion of the impact of systematic variations on the results - such systematic errors are simply not present in the simulation. However, it is still important to identify potential sources of uncertainty even at this stage, if only to qualitatively explore their possible impact on the performance of the scattering analysis.

5.6.1 Detector-level Systematics

Some sources of uncertainty are found outside the SMELLIE system, in other components of the wider detector configuration.

Certain uncertainties are associated with the properties of the detector materials - for example, there will be some error on the refractive indices of the acrylic in the AV as well as the water (and eventually, the LABPPO), and since these are important

quantities in the calculation of the time residual t_{res}^i and thus the scattering position \underline{s} , these errors will produce some uncertainty on the final calculation of \underline{s} and θ . There will be similar errors intrinsically present in many other physical parameters associated with the materials - such as the temperature, or the purity of the water within the cavity - all of which will produce systematic uncertainties in this analysis (and indeed, all analyses performed on real data).

In principle, consideration must be given to the physical uncertainties of the overall detector itself - such as the position of the AV with respect to the PSUP. In simulation, this is assumed to be the ideal situation - i.e. the AV is perfectly centred and stationary within the PSUP, but the possibility of this not being the case must be considered, as this will certainly affect the path that a photon takes between emission from a SMELLIE fibre and entering a PMT. For example, if the AV is positioned slightly higher in the PSUP than centre, the angles at which photons are incident on the AV from any one of the fibres will be different, leading to different reflected and refracted paths through the detector, and thus changing which PMTs are hit and triggered by light from a particular fibre. This in turn changes where in the t_{res}^i vs. α parameter space the optimal acceptance region will be found.

Another uncertainty is the exact number of noise PMTs present at any particular moment. As noted in Section 5.2.1 previously, the exact noise level in the detector has a non-negligible impact on the eventual choice of n_γ - the number of photons per SMELLIE laser head pulse, and also on the purity of the acceptance region. This in turn has an impact on both the calculation of l_{Ray} , as in Section 5.2.6, as well as the performance of the scattering position reconstruction: if the number of noise PMTs in the detector (and acceptance region) is higher than anticipated, then the overall performance of the scattering reconstruction could potentially be reduced due to the large contribution from these noise PMTs. Fortunately, the noise level of the detector during normal operation is fairly straight-forward to actually measure - by simply

running the detector with no global triggers enabled, and counting the number of PMTs that are individually triggered within a given time window.

One final possible source of uncertainty is the timing calibrations of the SNO+ PMTs, of which there are two types, both of which will eventually be used in tandem. The first type of calibration is known as Electronics Calibration (*ECA*) and involves calibrating the parameters that are used to convert the raw voltage from each PMT into a corresponding trigger time (this is the timing amplitude conversion, or TAC, mentioned previously in Section 1.1.2). The second calibration is the PMT Calibration (*PCA*), which is the calibration of differences between individual PMT electronics channels - for example, correcting for the fact that each channel has a slightly different route through the DAQ electronics, and therefore a slightly different readout time. Although these two calibrations are somewhat independent of each other, any errors or uncertainty on either one will propagate through to the final calibrated PMT trigger times used throughout this analysis. For example, if there is a bias on the calibration that results in over-estimated trigger times, this will result in the values of t_{sc} and $t_{sp,true}$ (Equation 5.13) being too large, and therefore the minimisation of Δt (Equation 5.14) occurring at an incorrect value of s_{reco} .

5.6.2 SMELLIE Systematics

As well as those systematics associated with the detector, there are a number of uncertainties originating within the SMELLIE system itself.

As mentioned previously in Section 4.2.1, the laser head intensity does naturally fluctuate slightly, and so this will produce an uncertainty on the number of photons entering the detector per laser head pulse. (The MPU measures this intensity fluctuation, but this measurement still must be related to the number of photons in the pulse - a relationship which also contains its own error.) This therefore leads to an uncertainty on the number of accepted-scattered PMTs in the acceptance region which

is separate from the natural statistical fluctuation, thereby introducing a source of error into the calculations of l_{Ray} and the scattering position \underline{s} .

The number of photons entering the detector is also related to the intrinsic intensity losses within the various optical paths between laser head and detector - i.e. the rack and detector fibres, as well as the Fibre Switch. Although the intensity losses for each of these components have been measured, there is of course some uncertainty on these measurements, which will eventually lead to corresponding errors on l_{Ray} , \underline{s} and θ .

The light emitted from the SMELLIE fibres is assumed to be completely unpolarised - i.e. each emitted photon has a completely random polarisation direction that can take any value, but it is possible that the light does in fact have a small overall polarisation - i.e. the polarisation has a small bias towards a particular direction. This could potentially be a property (or even a defect) of the laser heads, or due to effects occurring within the optical fibres. If this is true, it will affect the distribution of the scattering angle θ , since (as shown in Section 3.2.3) only perfectly unpolarised light - where the contributions from all polarisation directions cancel out completely - results in a $[1 + A \cos^2(\theta)]$ distribution (with A as the degree of polarisation given by Equation 3.21). Any underlying increase in the contribution from a specific polarisation direction will therefore result in a deviation from this expected distribution.

5.7 Summary

A Monte-Carlo generator has been developed that simulates the emission of light from the SMELLIE Calibration System within the wider RAT software framework. This generator has been tested against real data gathered during the air fill phase of SNO+, and in its original form a number of discrepancies between simulation and

data were found. Two of these - the incorrect positions and directions of some of the simulated SMELLIE fibres - have now been fully corrected for. A third discrepancy - between the angular widths of the simulated and real beams - has not yet been fully understood, but has been approximately accounted for in the simulation.

Using the SMELLIE generator, a study has been performed that investigates the potential to isolate only those photons that underwent Rayleigh scattering once within the detection medium - these being the most useful photons for any subsequent analysis of the scattering properties of the medium itself.

There are a number of other possible optical processes that a photon can be subject to between it being emitted from the SMELLIE fibre and potentially entering and triggering a PMT - such as scattering in the cavity water or in the AV material, reflections from any one or more of the boundary surfaces, and combinations of all of these. In water in particular, the probabilities of scattering and absorption are quite low due to water's long scattering and absorption lengths compared to the size of the SNO+ detector, and so photons can also undergo no optical deviations at all aside from refractions at the various surfaces encountered.

It has been found that applying simultaneous acceptance cuts (given by Equations 5.2 and 5.3) on the time residual t_{res}^i of each triggered PMT (calculated using Equation 5.1) and the angle α between the PMT and SMELLIE fibre (shown on Figure 5.6) results in a region of the t_{res}^i vs. α parameter space that is dominated by PMTs that were triggered by photons that underwent a single Rayleigh scattering in the detection medium.

The purity of this acceptance region is not perfect - the largest impurity stems from either noise PMTs or PMTs that were triggered by photons undergoing multiple optical processes, depending on how many photons are contained in each laser head pulse. However, the total population of the acceptance region - including PMTs

that were triggered by non-scattered photons - can still be related very accurately to the scaling factor of the detection medium's scattering length l_{Ray} by a second-order polynomial (shown in Figure 5.21), and therefore l_{Ray} itself can be calculated to within 1% of the true value, once the polynomial's coefficients are known.

The efficiency of the cuts that define the acceptance region is also very low - on the order of just a few percent, and this has an effect on the post-acceptance cuts distributions of the two most useful scattering properties of a single photon: d - the distance travelled by the photon prior to it undergoing Rayleigh scattering, and θ - the Rayleigh scattering angle. According to theory, the distributions of these quantities should follow specific and well-known shapes - exponential and $[1 + A \cos^2(\theta)]$ respectively, but the application of the acceptance cuts significantly distorts the distributions (shown in Figures 5.19 and 5.20).

Although some of the information that describes the bulk scattering properties of the detection medium has therefore been lost, it is still possible to reconstruct the scattering position of individual photons, using only PMT information and the known SMELLIE hardware parameters. A method has been developed to do this, using the fact that the true total time $t_{sp,true}$ that the photon took to travel from fibre to scattering position and then on to the PMT can be calculated using only quantities available in the SNO+ detector readout, and independently of the actual scattering position (Equation 5.13). By calculating the total path time at various trial scattering positions in the detector and then minimising the difference between the trial time and $t_{sp,true}$, it is therefore possible to reconstruct the true scattering position.

The basic principles of the reconstruction method have been tested against a simple pencil beam of photons, and found to perform very well - with a bias on the reconstructed scattering position \underline{s} of $< 4\text{cm}$ and a resolution of $< 30\text{cm}$ (Tables 5.8, C1 and C2) under the same conditions as would be found in real data. The

largest contribution to the difference between the true and reconstructed scattering positions is the photoelectron transit time - the time that the photoelectron takes to travel from the photocathode (located on the inside of glass face of the PMT) to the dynode stack within (Figure 5.9). In simulation, this time is randomly sampled from a preset distribution for every individual photoelectron, and it is not possible to calculate this value from the selection of observables available in the detector readout. In addition, the underlying distribution of photoelectron transit times has a width of a few nanoseconds, which is not insignificant compared to the total path time, and so introduces a non-negligible error into the calculation of $t_{sp,true}$ for each individual PMT. Nevertheless, the scattering position reconstruction is still extremely good in the context of the detector's physical scale, and the subsequent reconstruction of d and θ have correspondingly good performance (Tables 5.9, 5.10, and C3 to C6).

It has been found that applying the scattering position reconstruction method to all photons in the acceptance region results in very specific values of d and θ being calculated for some photons that did not undergo just a single Rayleigh scattering in the detection medium. This is a consistent feature across all fibre / wavelength combinations, and is due to the use of a pencil beam with a fixed initial photon direction, leading to certain paths being preferred over others.

The real SMELLIE photon beams have a non-zero angular distribution, and so the reconstruction method developed for pencil beams has been tested against non-zero width angular beams. There is a significant loss in performance for the reconstruction of the scattering position \underline{s} itself, as well as d (Table 5.11), mainly caused by the fact that there is no longer a unique location in the beam where the difference between the trial time and $t_{sp,true}$ is minimised, but instead there are a number of potential locations - only one of which is the actual scattering position (Figure 5.30). However, due to the design of the pencil beam analysis, it will always select the location along

the centre of the angular beam, even though the actual location is almost always off-centre. The poor reconstruction of d does have a corresponding impact on the reconstruction of θ , but since this reconstructed extremely well in the pencil beam scenario, the absolute performance in the angular beam scenario is still very good (Table 5.12).

It should be noted that the simulated SMELLIE angular beams appear to be significantly wider than those emitted from the real fibres (Figure 5.5), and so the performance of the pencil beam analysis when applied to the real angular beams may in fact be better than these simulations imply. Nevertheless, it is important to develop an analysis that is generalised for a beam of any angular width.

One possible idea for improving the performance of the pencil beam when applied to an angular beam is to use the angular distribution of photons as they are emitted from a particular fibre. These distributions are well-known for all twelve SMELLIE fibres, having been measured to high precision, and so can be used as probability distributions of the photon emission angle. Sampling from this distribution can in principle allow the pencil beam method to account for the angular distribution of photons as they leave the fibre, but this will only be useful in a bulk statistical sense. In order to improve the reconstruction performance at the individual photon level, it is more likely that a new method - one which has been built from the ground up to reconstruct the scattering of photons that have been emitted in a general angular distribution - will be required.

A number of potential sources of uncertainty - both within and outside the SMELLIE system, and applicable to both the pencil and angular beam analyses - have been identified and discussed, but because this work is based purely on Monte Carlo simulation, a quantitative calculation of the overall systematic uncertainties on l_{Ray} , \underline{s} , d and θ cannot be made at this time.

Chapter 6

BiPo Pileup Rejection in the Loaded Phase

‘So you’ve got no signal and no backgrounds ... what’s left?’

— Donal Hill, on the perils of searching for a rare process in a clean detector

6.1 $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ Backgrounds

6.1.1 Backgrounds at SNO+

In the loaded phase of SNO+, for any given $0\nu\beta\beta$ *signal*, a corresponding *background* can be thought of as any single decay or combination of decays that has similar observable characteristics (such as total deposited energy or PMT trigger-time distribution) to those of a $0\nu\beta\beta$ decay, and so could be mistaken for the latter. The major source of background events at SNO+ is the naturally-occurring radioactivity that is intrinsic to all of the materials present in and around the detection medium, and in particular, the isotopes that make up the Uranium and Thorium decay chains, shown in Figures 6.1 and 6.2 respectively.

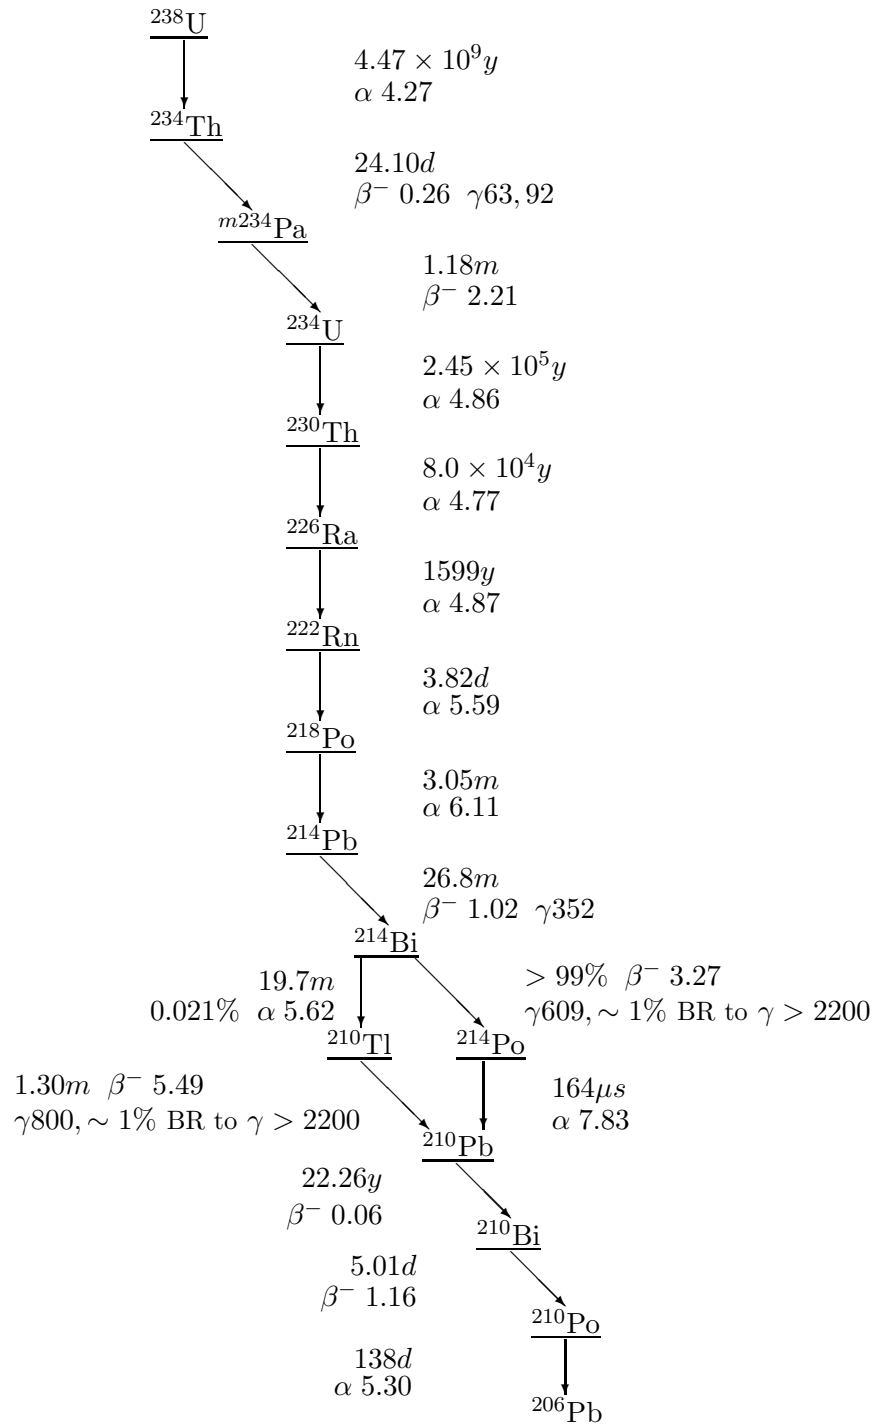


Figure 6.1: The Uranium Chain of radioactive decays [104]. For each isotope, the half-life and decay mode are given, together with the energy of the corresponding emitted particles (the highest single energy in MeV for α decays, the energy spectrum endpoint in MeV for β decays, and the energy in keV for γ emissions).




Figure 6.2: The Thorium Chain of radioactive decays [104]. For each isotope, the half-life and decay mode are given, together with the energy of the corresponding emitted particle (the highest single energy in MeV for α decays, the energy spectrum endpoint in MeV for β decays, and the energy in keV for γ emissions).

There are also several other isotopes that are not part of these decay chains, such as ^{40}K and ^{85}Kr , that are present in sufficient quantities so as to produce background events in the detector.

The backgrounds due to natural radioactivity can be divided into two categories: *external* (originating in the materials used to construct the detector itself, such as the surrounding rocks, and the PMTs and AV) and *internal* (the natural radioactive impurities present in the Te_LABPPO that is used as the detection medium in the loaded phase - see Section 1.2.3 for details regarding this material). The effect of external backgrounds on the detector is reduced substantially by the layers of water shielding between the cavity walls and the PSUP and between the PSUP and the AV, which together absorb most of the radioactivity emanating from the surrounding rocks, PMT glass and other ex-situ components. Internal backgrounds can be reduced through purification of the Te_LABPPO prior to introducing it into the AV.

These hardware-based methods are not a perfect solution - there will still be some external backgrounds that leech into the AV, and some remaining internal background contamination in the Te_LABPPO. However, cuts on both energy and position can be implemented during offline analysis that will further reduce the presence of backgrounds.

The effect of external backgrounds can be reduced through the use of a cut on the reconstructed event radius, i.e. by defining a *fiducial volume* (FV), since any remaining external backgrounds will be localised in the outer region of the AV, whereas the $0\nu\beta\beta$ decays should occur more uniformly throughout the detection medium. The FV is defined as a reconstructed event radius $\leq 3.5\text{m}$ [105], and has been chosen so as to reduce the external ^{208}Tl to a level comparable with the level of ^8B (solar) neutrino ES - this latter background is irreducible during the scintillator phases of SNO+ (due to a lack of directional information, and the fact that they appear point-like and with

a flat energy spectrum), and so effectively defines the minimum possible background rate in the SNO+ detector.

The ^{130}Te $0\nu\beta\beta$ decays have a narrow reconstructed energy range compared to the entire spectrum of possible background energies, as shown on Figure 6.3, and so a cut can be made on this reconstructed energy that will reject the vast majority of backgrounds - both internal and external. The particular $0\nu\beta\beta$ decay energy range is known as the *region of interest* (RoI), and is defined as being between -0.5σ and $+1.5\sigma$ of the mean reconstructed energy of ^{130}Te $0\nu\beta\beta$ decays ($\approx Q_{\beta\beta}$ assuming good energy reconstruction) [105]. (The exact value of σ - the standard deviation of the $0\nu\beta\beta$ reconstructed energy distribution - will depend on the true light yield of the Te_LABPPO cocktail.) An asymmetric RoI is used due to the presence of the ^{130}Te $2\nu\beta\beta$ spectrum, as shown on Figure 6.3. Due to its much steeper slope, the $2\nu\beta\beta$ energy spectrum overlaps quite considerably with that of $0\nu\beta\beta$ below the ^{130}Te $Q_{\beta\beta}$ value of 2.526MeV, and therefore the RoI has been shifted from being symmetric around $Q_{\beta\beta}$ to using more of the spectrum above $Q_{\beta\beta}$, since the only major background in that region is expected to be the ^8B neutrinos.

The overall aim of applying the various background rejection methods in combination with each other is to reduce all other backgrounds to a level that is negligible compared to the irreducible ^8B rate - thereby ensuring that the ^8B neutrino energy spectrum is the dominant remaining background in comparison to the ^{130}Te $0\nu\beta\beta$ signal.

Table 6.1 gives the expected number of events in the detector for a selection of internal and external backgrounds, both before and after applying the RoI and FV cuts, and it can be seen that these two cuts alone do an excellent job of reducing many of the backgrounds to a rate well below that of ^8B - particularly for the external backgrounds, where the isotopes present due to the PMTs would otherwise create a truly overwhelming number of radioactive decays.




Figure 6.3: The reconstructed energy spectra of $^{130}\text{Te } 0\nu\beta\beta$ (black), $^{130}\text{Te } 2\nu\beta\beta$ (red), ^8B (green) and the two BiPo chains: $^{212}\text{BiPo}$ (orange) and $^{214}\text{BiPo}$ (blue) in Te_LABPPO, after application of only the fiducial volume cut as described in the main text. The distributions have been normalised to be consistent with the expected number of events in the region of interest - shown in grey - as given in Table 6.1. The $2\nu\beta\beta$ distribution is significantly larger than the others - peaking at a value of ~ 61000 , hence why its full height has been truncated.

It can be seen however that there are still a few backgrounds with rates comparable to that of ^8B even after applying the RoI and FV cuts. Two of these are the ($^{212}\text{Bi} + ^{212}\text{Po}$) and ($^{214}\text{Bi} + ^{214}\text{Po}$) sub-chains, abbreviated to $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ and originating in the larger Thorium and Uranium chains respectively. These subchains pose a particular problem for the extraction of the $^{130}\text{Te } 0\nu\beta\beta$ signal - not just because the combination of the Bi β and the Po α can have a total energy that overlaps with that of $^{130}\text{Te } 0\nu\beta\beta$ decay, as shown in Figure 6.3 (due to the α undergoing ionisation quenching (see Section 1.2.2) and therefore being reconstructed at a much lower energy than its true value) but also since both of these backgrounds are internal they are less affected by the FV cut, and so remain at a comparatively high level.

Type	Isotope	Events per year	
		Before RoI and FV Cuts	After RoI and FV Cuts
Signal	$0\nu\beta\beta$	~ 245	~ 29
Internal	^{208}Tl	10116	0.004
	^{210}Tl	159	1.27
	$^{212}\text{Bi} + ^{212}\text{Po}$	5665 (non-pileup) 12318 (pileup)	0.00 (non-pileup) 1.74 (pileup)
	$^{214}\text{Bi} + ^{214}\text{Po}$	753506 (non-pileup) 1593 (pileup)	0.00 (non-pileup) 0.54 (pileup)
	$^{234(m)}\text{Pa}$	755258	0.27
External	^{208}Tl	$> 4.4 \times 10^{10}$	2.89
	^{214}Bi	$> 3.7 \times 10^{11}$	0.448
	^8B	1021	7.26

Table 6.1: The expected number of events per year in the Te_LABPPO-filled detector, before and after applying the RoI and FV cuts, for a selection of the most important internal and external backgrounds, with the ^{130}Te $0\nu\beta\beta$ signal included for reference. The pre- and post-cut values for ^{130}Te have been extracted from [90]. All pre-cut values are based on measurements of the masses of each isotope present in the various detector components [106], and the remaining post-cut values are derived from MC simulations [107]. The concepts of “pileup” and “non-pileup” events are detailed Section 6.1.2.

6.1.2 “BiPo Pileup” Events

As previously noted in Section 1.1.2, the information that comprises a single detected event in the SNO+ detector is collected over a trigger window of approximately 400ns width, with the exact value varying slightly between PMT channels. In each window, time “zero” is defined as the trigger time of the individual PMT that caused the sum of all trigger pulses to cross the predefined global trigger’s threshold [108], and the trigger window then spans the time between -180ns and +220ns around this zero-point. If the rate of radioactive background decays is high enough, it is possible for two or more decays to occur within the same 400ns trigger window, creating a *pileup event* in which the information from multiple physical events occurring in the

detector are combined to create a single observed event. This can occur both for multiple independent decays that happen to occur in close temporal proximity, as well as for connected processes in a single decay-chain such as $\text{Bi} \rightarrow \text{Po} \rightarrow \text{Pb}$. In the latter situation, where the first decay in the chain occurs very close to time “zero” in the trigger window, only the “positive-time” 220ns-long section of the window is relevant, since obviously the second decay in the chain cannot occur at a negative time with respect to the first.

Due to their short lifetimes, it is very possible for the Po in either BiPo sub-chain to α decay to its corresponding Pb within 220ns of the preceding Bi β decay. If this happens, the Bi β and Po α will both be emitted - and detected - within the same trigger window, creating a *BiPo pileup event*.

Non-pileup BiPo events - where the Po α is emitted more than 220ns after the Bi β - are fairly straight-forward to identify, as they will be seen as a pair (consecutive or not) of triggered events with the characteristic energies of the β and α . However, pileup events - comprising only one triggered event, as discussed above - look effectively identical to a $^{130}\text{Te } 0\nu\beta\beta$ decay, and therefore can be easily mistaken for signal events if the pileup event’s energy also falls within the RoI. Clearly then, using only an event’s reconstructed energy and position is not enough to identify it as a $^{130}\text{Te } 0\nu\beta\beta$ or a BiPo pileup event, and so a more detailed approach is necessary.

6.1.3 Simulation Considerations

The *coincidence* and *Decay0* Monte-Carlo generators in RAT have been used to simulate the Bi β and Po α decays, with some minor modifications made in order to correctly simulate the individual decays as being part of a single BiPo sub-chain. Left unchanged, Decay0 simulates all radioactive decays of a particular isotope - i.e. both α and β decays of the Bi isotopes, and specific decays cannot be turned off in

the generator itself. However, this would mean that simulation time is wasted generating the ^{212}Bi and ^{214}Bi α decays, which are not used in this analysis. To solve this, the α -decay branches of both Bi isotopes have been turned off in Decay0 by setting their branching fractions to zero in the generator's datafile - meaning that no α 's are simulated, and the simulation time is decreased (or equivalently, more ^{212}Bi and ^{214}Bi β decays can be simulated in the same amount of time). The coincidence generator has been modified to give the Bi β and Po α the same position in the detector, and also to force each Po α decay to only occur *after* the parent Bi β decay (instead of the generator's normal behaviour, which is to locate the second decay at a random position in the detector, and at a random time within a window that spans both before and after the first decay).

The detection medium in all simulations is the material Te_LABPPO, as previously defined in Section 1.2.3, and of the BiPo events only pileup has been simulated - i.e. the time difference between the Bi β and Po α decays, hereafter denoted by δt , has been limited to a maximum value of 220ns for both BiPo sub-chains. The RoI and FV cuts as described in Section 6.1.1 have been applied to both BiPo pileup datasets. Three independently-generated ^{130}Te $0\nu\beta\beta$ datasets have been used: one in conjunction with each BiPo pileup dataset (with both RoI and FV cuts applied), and a third for use outside direct comparisons to BiPo pileup (with the RoI and FV cuts applied when needed, as noted where appropriate). Only the FV cut has been applied to the individual Bi β and Po α decay datasets.

6.1.4 Time Residuals

The *time residual*, t_{res}^i of the i^{th} triggered PMT in an event is given by:

$$t_{res}^i = t_{pmt}^i - t_{flight}^i - t_{ev} \quad (6.1)$$

where t_{pmt}^i is the time at which the i^{th} PMT was triggered, t_{ev} is the reconstructed time of the event, and t_{flight}^i is the transit time of the photon as it travels between the reconstructed event position and the i^{th} PMT. This correction is calculated using the `LightPath` class in RAT [103], which utilises a direct path between the reconstructed event and PMT, travelling through the `Te_LABPPO`, `AV` and `water`, and takes into consideration the refraction at the various surfaces encountered along this path, as well as the different velocities of light in the three materials due to their respective refractive indices.

During normal operation of the detector, the PMTs will be subject to electronic noise and other effects such as flashing and electric field break-down, all of which can cause anomalous PMT triggering. This means that some of the time residuals in any particular physical event may not be truly associated with the event itself, making it difficult to reliably compare time residual distributions on an event-by-event basis. Very early anomalous PMT triggering (noted previously in Section 5.2.1 to be related specifically to the homogeneous background noise in the detector) is a particular problem, since this can imply that an event began emitting light far earlier than expected, perhaps even before the event truly occurred in the detector.

To correct for such anomalous triggering, each event's time residual distribution can be offset by a certain event-specific amount, such that the time residual distributions across all events then have a common starting point that deviates very little from event to event, across any type of decay process. Figure 6.4 shows the distributions of the unaltered 1st, 2nd, 5th and 10th earliest time residuals over many simulated ^{130}Te $0\nu\beta\beta$ decays.




Figure 6.4: The distributions of the 1st (red), 2nd (orange), 5th (green) and 10th (blue) earliest time residuals, as calculated using Equation 6.1, of ^{130}Te $0\nu\beta\beta$ decay events with only the FV cut applied. The corresponding distributions for ^{212}Bi β , ^{214}Bi β , ^{212}Po α and ^{214}Po α decays are identical to within statistical fluctuations.

The earliest time residual can take a very wide range of possible values for different events, from -200ns to 0ns. Therefore, this is not a good common point across different time residual distributions, and neither is the second-earliest time residual, which also can fluctuate by $\sim 175\text{ns}$ between events. However, later time residuals such as the fifth or tenth are clearly much more stable - fluctuating only by a few nanoseconds even across many events, and therefore acting as excellent common points across events. The specific choice of which time residual to use is somewhat arbitrary, so for the purposes of this analysis the tenth-earliest time residual, hereafter denoted by t_{10} , has been used.

In order to anchor the individual events by their t_{10} values, each event's time residual distribution is offset by this value, so that the previous definition of the time residual (Equation 6.1) now becomes:

$$t_{res,off}^i = t_{pmt}^i - t_{flight}^i - t_{ev} - t_{10} \quad (6.2)$$

The offset time residual distributions (normalised to the total number of PMTs in each distribution) for ^{212}Bi β and ^{212}Po α decays in Te_LABPPO are shown in Figure 6.5, and Figure 6.6 shows the same distributions for ^{214}Bi β and ^{214}Po α decays in the same material, with the ^{130}Te $0\nu\beta\beta$ time residual distribution shown on both figures for comparison.

In contrast to the differences between α and β particle scintillation that might be expected given the discussion in Section 1.2.2, there is very little difference between the α and β time residual distributions in Te_LABPPO. This is due to the presence of various solvents that allow the Tellurium to be loaded into the scintillator - these solvents decrease the fraction of luminescence originating from phosphorescence [109], thus reducing the late-light tail that is expected from α particles. There is also no discernible difference between the time residual distributions of ^{212}Bi and ^{214}Bi β 's (and similarly between ^{212}Po and ^{214}Po α 's). This is simply because the only observable difference between these decay processes - the emitted particle's energy - is negated when each distribution is normalised to the total number of triggered PMTs.

Figure 6.7 shows the time residual distributions of three separate $^{212}\text{BiPo}$ pileup events with varying δt , each compared to a single ^{130}Te $0\nu\beta\beta$ time residual distribution. It should be noted that the raw time reconstruction of such pileup events does not always set the Bi β peak to be at $t_{res}^i \simeq 0\text{ns}$, as implied by Figure 6.7. This is a result of how the time reconstruction works - it is optimised for a single timing peak, and so will be biased towards the larger, more dominant peak in a pileup distribution. In most BiPo pileup events, the β peak is larger than that of the α , but if the β peak is smaller, the α peak will be preferentially set at $t_{res}^i \simeq 0\text{ns}$, with the β peak appearing at negative t_{res}^i . However, this situation is naturally rectified with the use of the offset time residual $t_{res,off}^i$, as detailed in Equation 6.2 above.)




Figure 6.5: The offset time residual distributions, as calculated using Equation 6.2 for $^{130}\text{Te } 0\nu\beta\beta$ (black), $^{212}\text{Bi } \beta$ (red) and $^{212}\text{Po } \alpha$ (green) decays in Te_LABPPO, all normalised to the total number of PMTs in each distribution. Only the FV cut has been applied to the $0\nu\beta\beta$ events. There is almost no difference between the distributions, although the α distribution has a very slightly higher and narrower early-time peak.




Figure 6.6: The offset time residual distributions, as calculated using Equation 6.2 for $^{130}\text{Te } 0\nu\beta\beta$ (black), $^{214}\text{Bi } \beta$ (red) and $^{214}\text{Po } \alpha$ (green) decays in Te_LABPPO, all normalised to the total number of PMTs in each distribution. Only the FV cut has been applied to the $0\nu\beta\beta$ events. As with the distributions for ^{212}Bi and ^{212}Po , the only notable difference between the distributions is the slightly higher and narrower early-time peak of the α distribution.

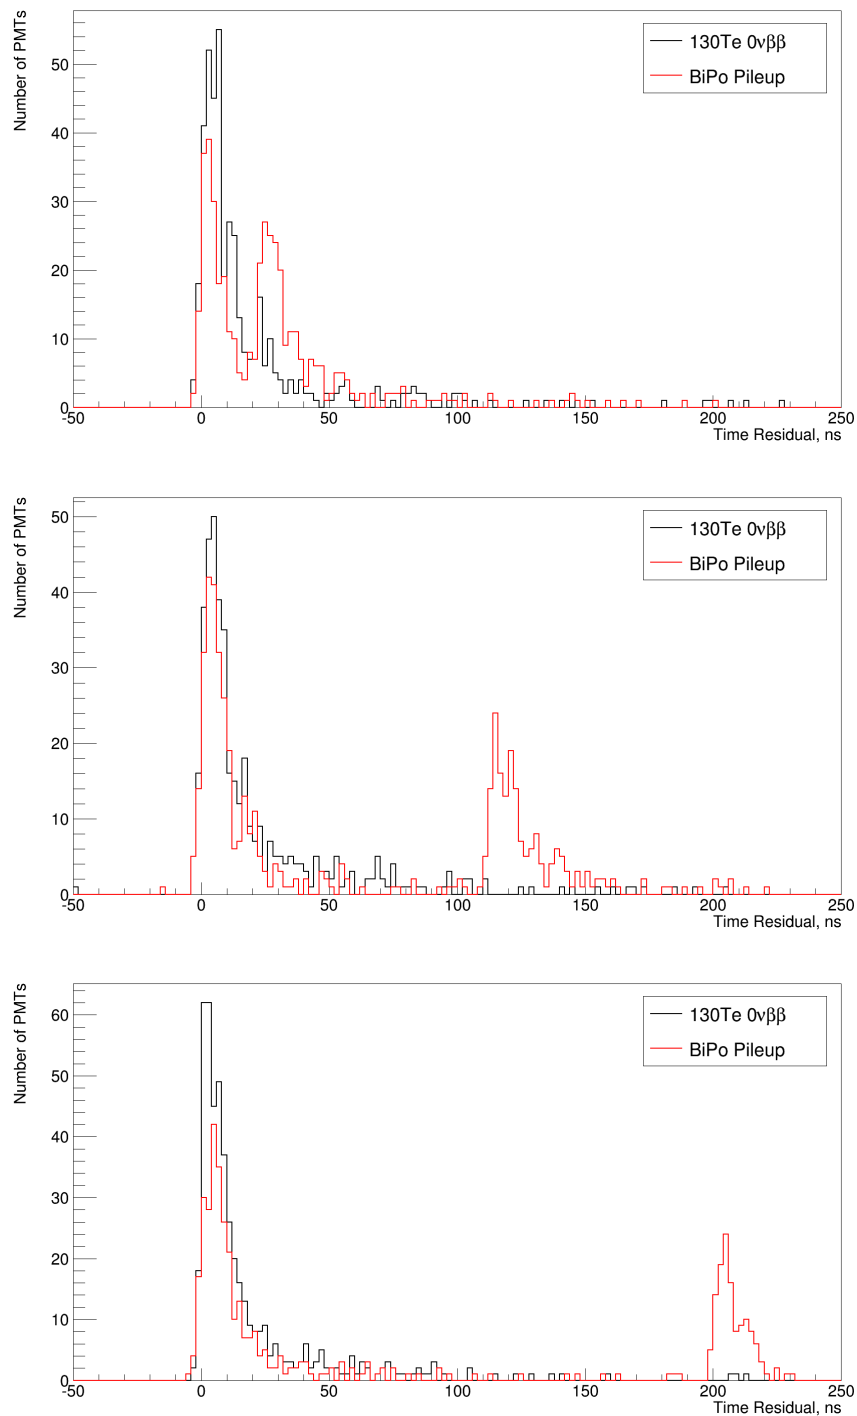


Figure 6.7: The time residual distributions, as calculated using Equation 6.2, for a $^{130}\text{Te } 0\nu\beta\beta$ (black) and a $^{212}\text{BiPo}$ pileup event (red), where the time separation δt between the Bi β (initial red peak) and Po α (second red peak) is $\sim 20\text{ns}$ (top), $\sim 110\text{ns}$ (middle) and $\sim 200\text{ns}$ (bottom). In the case of the largest δt , the time residual distribution is cut off at $\sim 230\text{ns}$ due to the end of the trigger window. Similar distributions are seen for $^{214}\text{BiPo}$ pileup events.

6.2 Cumulative Time Residuals Method

6.2.1 Kolmogorov-Smirnov Γ Parameter

Figure 6.7 shows that as long as δt is large enough, the time residual distribution of a BiPo pileup event is clearly different from that of a $^{130}\text{Te } 0\nu\beta\beta$ event. One method of quantifying this difference is by using a one-sample Kolmogorov-Smirnov (K-S) test, whereby an unknown (potentially BiPo pileup) event's cumulative time residual distribution is compared to a reference cumulative time residual distribution from $^{130}\text{Te } 0\nu\beta\beta$ events, and a K-S statistic is then calculated that quantifies the difference between the distributions.

Figure 6.8 shows the cumulative time residual distribution of the BiPo pileup event shown in Figure 6.7 (middle), together with a reference distribution created from the third independent set of $^{130}\text{Te } 0\nu\beta\beta$ events.




Figure 6.8: The cumulative time residual distribution for a single $^{212}\text{BiPo}$ pileup event with $\delta t \sim 110\text{ns}$ (red - the same event that is shown in Figure 6.7 (middle)), and a reference cumulative time residual distribution from $^{130}\text{Te } 0\nu\beta\beta$ events (black), with both FV and RoI cuts applied.

The reference distribution climbs very quickly after the start of the event, reaching a cumulative fraction very close to 1 within the first 100ns of the event. In contrast, the BiPo pileup event's distribution does not increase nearly as quickly, as the first peak resulting from the Bi β decay contains only $\sim 60\%$ of the total light emanating from the pileup, with the rest of the light from the Po α decay arriving later in the event.

A quantitative measure of the overall difference between the reference and unknown event distributions is given by Γ :

$$\Gamma = \frac{1}{N_{\text{hits}}} \times \sum_b (x_b^{ev} - x_b^{ref})^2 \quad (6.3)$$

where x_b^{ev} and x_b^{ref} are the populations of bin b in the event and reference distributions respectively. (Note that this quantity is not the same as the basic K-S test statistic, which is simply the maximum single value of $x_b^{ev} - x_b^{ref}$ in any bin.) It can be expected that $^{130}\text{Te } 0\nu\beta\beta$ events will have a very small value of Γ , since the reference distribution was created using $^{130}\text{Te } 0\nu\beta\beta$ events and so the differences between bins in the event and reference distributions will be very small, whereas BiPo pileup events will have much larger values of Γ , due to the large deviations from the reference distribution as suggested by Figure 6.8. The distributions of Γ for $^{212}\text{BiPo}$ pileup and $^{214}\text{BiPo}$ pileup events, together with those of $^{130}\text{Te } 0\nu\beta\beta$ events, are shown in Figures 6.9 and 6.10 respectively.

6.2.2 Γ Cut and Remaining Signal vs. Background

A simple cut on Γ , i.e. only accepting events with Γ less than some cut value Γ_{cut} , can be used to reject the vast majority of both $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ pileup events. Table 6.2 gives the percentages of $^{130}\text{Te } 0\nu\beta\beta$ and BiPo pileup events remaining after




Figure 6.9: Distributions of Γ , as calculated using Equation 6.3, for $^{130}\text{Te } 0\nu\beta\beta$ (blue) and $^{212}\text{BiPo}$ pileup (red) events. The $^{130}\text{Te } 0\nu\beta\beta$ events are taken from a different dataset than that used to create the reference distribution shown in Figure 6.8. As expected, $^{130}\text{Te } 0\nu\beta\beta$ events generally have much smaller values of Γ than the BiPo pileup events. The latter distribution extends well beyond the range shown here, hence why there are so few BiPo events compared to $^{130}\text{Te } 0\nu\beta\beta$.




Figure 6.10: Distributions of Γ , as calculated using Equation 6.3, for $^{130}\text{Te } 0\nu\beta\beta$ (blue) and $^{214}\text{BiPo}$ pileup (red) events. Once again, $^{130}\text{Te } 0\nu\beta\beta$ events generally have much smaller values of Γ than the BiPo pileup events, which can have Γ much larger than the range shown here.

such a cut on Γ , for varying placement of Γ_{cut} . Figures 6.11 and 6.12 show the same values in graph-form.

Γ_{cut}	% of Events that have $\Gamma < \Gamma_{cut}$			
	$^{130}\text{Te } 0\nu\beta\beta$	$^{212}\text{BiPo Pileup}$	$^{130}\text{Te } 0\nu\beta\beta$	$^{214}\text{BiPo Pileup}$
0.0001000	82.18	2.62	81.39	1.63
0.0001225	87.28	2.95	86.86	1.92
0.0001450	91.01	3.22	91.10	2.12
0.0001675	93.66	3.60	93.99	2.36
0.0001900	95.88	3.78	95.53	2.52
0.0002125	96.86	3.98	96.56	2.70
0.0002350	97.38	4.21	97.49	2.86
0.0002575	98.04	4.33	98.10	3.02
0.0002800	98.53	4.43	98.59	3.18
0.0003025	98.82	4.54	98.91	3.31
0.0003250	99.12	4.64	99.10	3.48
0.0003475	99.25	4.76	99.45	3.61
0.0003700	99.35	4.91	99.65	3.78
0.0003925	99.54	5.06	99.71	3.89
0.0004150	99.64	5.14	99.74	3.97
0.0004375	99.77	5.19	99.84	4.11
0.0004600	99.84	5.32	99.90	4.19
0.0004825	99.87	5.44	99.97	4.30
0.0005050	99.87	5.54	99.97	4.39
0.0005275	99.87	5.64	99.97	4.45
0.0005500	99.90	5.64	99.97	4.53

Table 6.2: The remaining percentages of $^{130}\text{Te } 0\nu\beta\beta$, $^{212}\text{BiPo}$ pileup and $^{214}\text{BiPo}$ pileup events after accepting only events with $\Gamma < \Gamma_{cut}$, for varying placement of Γ_{cut} . Two separate $^{130}\text{Te } 0\nu\beta\beta$ datasets were used (one for each BiPo type), hence why there are two corresponding columns. In all cases, “100%” refers to all events from each isotope’s original dataset that pass the FV and RoI cuts.

The optimal value of Γ_{cut} will be dependent on a wider optimisation of the acceptable background levels and signal loss, taking into account all sources and rates of background contamination. However, if a 1% $^{130}\text{Te } 0\nu\beta\beta$ signal loss is taken as an example, this corresponds to $\Gamma_{cut} = 0.0003160$ ($^{212}\text{BiPo}$) and 0.0003133 ($^{214}\text{BiPo}$), and a remaining background percentage of 4.60% ($^{212}\text{BiPo}$) and 3.39% ($^{214}\text{BiPo}$). Based on




Figure 6.11: The remaining percentages of $^{130}\text{Te } 0\nu\beta\beta$ (signal) and $^{212}\text{BiPo}$ pileup (background) events after accepting only events with $\Gamma < \Gamma_{cut}$, for varying placement of Γ_{cut} . Each point on the figure is a single value of Γ_{cut} given in Table 6.2.




Figure 6.12: The remaining percentages of $^{130}\text{Te } 0\nu\beta\beta$ (signal) and $^{214}\text{BiPo}$ pileup (background) events after accepting only events with $\Gamma < \Gamma_{cut}$, for varying placement of Γ_{cut} . Each point on the figure is a single value of Γ_{cut} given in Table 6.2.

the expected number of events remaining in the detector after applying the RoI and FV cuts (as shown previously in Table 6.1), these percentages correspond to 0.080 $^{212}\text{BiPo}$ pileup and 0.018 $^{214}\text{BiPo}$ pileup events per year in the RoI and FV.

6.2.3 Post- Γ Cut BiPo Pileup Events

Figures 6.13 and 6.14 show the distributions of δt for $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ pileup events respectively that pass the cut on Γ , when taking $\Gamma_{cut} = 0.0003160$ ($^{212}\text{BiPo}$) and $\Gamma_{cut} = 0.0003133$ ($^{214}\text{BiPo}$), i.e. a 1% $^{130}\text{Te } 0\nu\beta\beta$ signal loss.

When δt is less than $\sim 15\text{ns}$, the individual time residual distribution peaks from the Bi β and Po α are no longer separated. Instead, they merge into a single peak, with an overall distribution identical in shape to that of a $^{130}\text{Te } 0\nu\beta\beta$ decay, therefore making it impossible to distinguish between the signal and pileup background events using time residuals alone.

A similar problem occurs when δt is greater than $\sim 210\text{ns}$, as shown by a small number of $^{214}\text{BiPo}$ pileup events on the far-right of Figure 6.14. In this situation, the light from the Po α arrives too late to be contained within the same trigger window as the light from the Bi β , and so the BiPo pileup time residual distribution is split cleanly between two consecutive windows, the first of which contains the distribution of a lone β particle - identical in shape to that of a $0\nu\beta\beta$ decay, as seen previously in Figure 6.6. Such events, though pileup in truth, therefore become non-pileup when observed in the detector, and as such can be tagged and rejected using a technique that has been developed for non-pileup BiPo events [110], with a rejection efficiency approaching 100% [107].




Figure 6.13: The time separation δt between the Bi β and Po α in $^{212}\text{BiPo}$ pileup events that pass the Γ cut at a 1% ^{130}Te $0\nu\beta\beta$ signal loss.




Figure 6.14: The time separation δt between the Bi β and Po α in $^{214}\text{BiPo}$ pileup events that pass the Γ cut at a 1% ^{130}Te $0\nu\beta\beta$ signal loss. Note the very small peak at $\delta t = 220$ ns.

6.3 Log-Likelihood Difference Method

6.3.1 Log-Likelihood Parameters

It has been shown that the uniformly most powerful (UMP) hypothesis test - that is, the test of a given size with the highest probability of successfully rejecting an incorrect null hypothesis when the alternative hypothesis is correct - is the likelihood ratio test [111]. A log-likelihood difference test is numerically equivalent to this, but results in smaller - and easier to manage - numbers.

The log-likelihood of an unknown event being a $^{130}\text{Te } 0\nu\beta\beta$ decay based only on the shape of its time residual distribution is given by:

$$\mathcal{L}_{\text{Te}} = \sum_{i=1}^N \ln (N \times P_{\text{Te}} (t_{\text{res}}^i)) \quad (6.4)$$

and similarly, the log-likelihood of the same event being a BiPo pileup with a true $\beta - \alpha$ time separation δt_{true} is:

$$\mathcal{L}_{\text{BiPo}} = \sum_{i=1}^N \ln ([(N - A) \times P_{\text{Bi}} (t_{\text{res}}^i)] + [A \times P_{\text{Po}} (t_{\text{res}}^i - \delta t_{\text{true}})]) \quad (6.5)$$

where N is the number of triggered PMTs in the event which have a time residual that falls within the same trigger window as the event start time, A is the mean number of triggered PMTs produced by a single Po α decay and which are triggered in the time between the start of the α -peak of the time residual distribution (i.e. event start time + δt_{true}) and the end of the trigger window, and $P_X(t)$ is the probability of getting the time residual t in a particular decay process ($X = \text{Te}, \text{Bi}$ or Po for $^{130}\text{Te } 0\nu\beta\beta$, ^{212}Bi and ^{214}Bi , or ^{212}Po and ^{214}Po α decays respectively). The time residual

probability distributions $P_X(t)$ are those previously shown in Figures 6.5 and 6.6.

When dealing with an unknown event, δt_{true} in Equation 6.5 will of course also be unknown, but can be found by trialling different values of δt in place of δt_{true} , with the condition that the overall value of \mathcal{L}_{BiPo} will be maximised when $\delta t = \delta t_{true}$. This maximisation procedure is performed by one of the *optimisers* that are available in RAT. Some specific examples are *Minuit2* [112], *Powell*, *Simulated Annealing* [113], *Grid* [114] and *MetaDriveCorrectSeed* [115], but in general the optimisers all take one or more free parameters (in this case, δt) which are trialled within a specified range (here, between 0 and 220ns) with a particular step-size between consecutive trial values (fixed at 4ns for this optimisation, but in many optimisers this is a dynamic value), in order to maximise the value of a quantity that is dependent on the parameters (\mathcal{L}_{BiPo}).

The step-size of 4ns has been chosen as the best compromise between speed and accuracy of the δt optimisation. At best, the optimisers can reliably calculate δt only to within 4 or 5ns of δt_{true} , i.e. the maximisation of \mathcal{L}_{BiPo} occurs when $\delta t = \delta t_{opt} \approx \delta t_{true}$. This occurs even when using a step-size of 1ns, and so there is little reason to use a step-size smaller than this error. Using a larger step-size, and therefore requiring fewer steps to cover the entire range, also increases the speed of the optimisation. Figures 6.15 and 6.16 show the value of δt_{opt} compared to δt_{true} as calculated using the Grid optimiser, for $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ pileup events respectively, along with $^{130}\text{Te } 0\nu\beta\beta$ decays.

There is very good correlation between δt_{opt} and δt_{true} over most of the latter's range, for both types of BiPo pileup events. The optimisation begins to lose performance when $\delta t_{true} \geq 190\text{ns}$ - this is close to the end of the trigger window, and so the poorer optimisation is likely due to the same effect which allowed a small number of BiPo pileup events to pass the Γ cut in the Cumulative Time Residuals method - that




Figure 6.15: The value of δt_{opt} , i.e. the trial value of δt which maximises $\mathcal{L}_{\text{BiPo}}$ as calculated by the Grid optimiser, compared to δt_{true} for $^{212}\text{BiPo}$ pileup (red) and $^{130}\text{Te } 0\nu\beta\beta$ (blue) events.




Figure 6.16: The value of δt_{opt} , i.e. the trial value of δt which maximises $\mathcal{L}_{\text{BiPo}}$ as calculated by the Grid optimiser, compared to δt_{true} for $^{214}\text{BiPo}$ pileup (red) and $^{130}\text{Te } 0\nu\beta\beta$ (blue) events.

is, in events with $\delta t_{true} \geq 190\text{ns}$, the second time residual peak from the Po α decay falls into the next trigger window, and so the overall distribution looks like a non-pileup β decay. However, the optimisation is clearly still good enough to distinguish most of even these BiPo pileup events from ^{130}Te $0\nu\beta\beta$ decays, which can be thought of as having $\delta t_{true} = 0\text{ns}$, and for which the optimisation works extremely well. Only BiPo pileup events with $\delta t_{true} \geq 210\text{ns}$ completely fail to optimise, resulting in δt_{opt} close to 0ns .

The quantitative value of interest is the difference (on a per-event basis) between the two log-likelihood values:

$$\Delta\mathcal{L} = \mathcal{L}_{\text{BiPo}} - \mathcal{L}_{\text{Te}} \quad (6.6)$$

BiPo pileup events are expected to have large positive values of $\Delta\mathcal{L}$, since $\mathcal{L}_{\text{BiPo}} > \mathcal{L}_{\text{Te}}$ for such events, and correspondingly ^{130}Te $0\nu\beta\beta$ events should have $\mathcal{L}_{\text{Te}} > \mathcal{L}_{\text{BiPo}}$ and therefore negative or very small values of $\Delta\mathcal{L}$. Figures 6.17 and 6.18 show the distributions of $\Delta\mathcal{L}$ for ^{130}Te $0\nu\beta\beta$ events, together with $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ pileup events respectively.

It should be noted that the following definition of $\Delta\mathcal{L}$ is also technically valid:

$$\Delta\mathcal{L} = \mathcal{L}_{\text{Te}} - \mathcal{L}_{\text{BiPo}}$$

but the definition given in Equation 6.6 has been chosen so that ^{130}Te $0\nu\beta\beta$ events generally have smaller values of the parameter $\Delta\mathcal{L}$ than BiPo pileup events, in the same way that ^{130}Te $0\nu\beta\beta$ events have smaller values of the parameter Γ when using the Cumulative Time Residual method discussed previously.




Figure 6.17: Distributions of $\Delta\mathcal{L}$, as calculated using Equation 6.6, for $^{130}\text{Te } 0\nu\beta\beta$ (blue) and $^{212}\text{BiPo}$ pileup (red) events. As expected, $^{130}\text{Te } 0\nu\beta\beta$ events generally have smaller values of $\Delta\mathcal{L}$ than the BiPo pileup events. The BiPo distribution extends well beyond $\Delta\mathcal{L} = 500$, hence why there are so few $^{212}\text{BiPo}$ events in the displayed range of $\Delta\mathcal{L}$.




Figure 6.18: Distributions of $\Delta\mathcal{L}$, as calculated using Equation 6.6, for $^{130}\text{Te } 0\nu\beta\beta$ (blue) and $^{214}\text{BiPo}$ pileup (red) events. Once again, $^{130}\text{Te } 0\nu\beta\beta$ events generally have smaller values of $\Delta\mathcal{L}$ than the BiPo pileup events, and the BiPo distribution extends well beyond the displayed range of $\Delta\mathcal{L}$.

6.3.2 $\Delta\mathcal{L}$ Cut and Remaining Signal vs. Background

As with the parameter Γ in the Cumulative Time Residuals method previously, accepting only events with $\Delta\mathcal{L}$ less than some cut value $\Delta\mathcal{L}_{cut}$ will reject the vast majority of both $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ pileup events. Table 6.3 give the percentages of $^{130}\text{Te } 0\nu\beta\beta$ and BiPo pileup events remaining after such a cut, for varying placement of $\Delta\mathcal{L}_{cut}$. Figures 6.19 and 6.20 show the same values in graph-form.

$\Delta\mathcal{L}_{cut}$	% of Events that have $\Delta\mathcal{L} < \Delta\mathcal{L}_{cut}$			
	$^{130}\text{Te } 0\nu\beta\beta$	$^{212}\text{BiPo}$ Pileup	$^{130}\text{Te } 0\nu\beta\beta$	$^{214}\text{BiPo}$ Pileup
-2	81.20	1.44	39.81	0.36
-1	86.99	1.71	49.97	0.48
0	92.15	1.89	61.25	0.65
1	95.32	2.24	71.01	0.90
2	96.96	2.49	77.99	1.13
3	98.17	2.87	84.16	1.35
4	98.89	3.20	89.07	1.63
5	99.31	3.33	93.38	1.87
6	99.71	3.55	96.02	2.09
7	99.80	3.63	97.59	2.32
8	99.84	3.78	98.46	2.58
9	99.93	3.83	99.10	2.85
10	99.93	3.88	99.42	3.14
11	99.93	4.03	99.65	3.41
12	99.67	4.13	99.84	3.66
13	99.67	4.21	99.90	3.94
14	99.67	4.28	99.94	4.25
15	100.0	4.33	100.0	4.54
16	100.0	4.43	100.0	4.79
17	100.0	4.54	100.0	5.08

Table 6.3: The remaining percentages of $^{130}\text{Te } 0\nu\beta\beta$, $^{212}\text{BiPo}$ pileup and $^{214}\text{BiPo}$ pileup events after accepting only events with $\Delta\mathcal{L} < \Delta\mathcal{L}_{cut}$, for varying placement of $\Delta\mathcal{L}_{cut}$. Two separate $^{130}\text{Te } 0\nu\beta\beta$ datasets were used (one for each BiPo type), hence why there are two corresponding columns. In all cases, “100%” refers to all events from each isotope’s original dataset that pass the FV and RoI cuts.

Once again taking an acceptable ^{130}Te signal loss of 1% as an example - corresponding




Figure 6.19: The remaining percentages of $^{130}\text{Te } 0\nu\beta\beta$ (signal) and $^{212}\text{BiPo}$ pileup (background) events after accepting only events with $\Delta\mathcal{L} < \Delta\mathcal{L}_{cut}$, for varying placement of $\Delta\mathcal{L}_{cut}$. Each point on the figure is a single value of $\Delta\mathcal{L}_{cut}$ given in Table 6.3.




Figure 6.20: The remaining percentages of $^{130}\text{Te } 0\nu\beta\beta$ (signal) and $^{214}\text{BiPo}$ pileup (background) events after accepting only events with $\Delta\mathcal{L} < \Delta\mathcal{L}_{cut}$, for varying placement of $\Delta\mathcal{L}_{cut}$. Each point on the figure is a single value of $\Delta\mathcal{L}_{cut}$ given in Table 6.3.

to $\Delta\mathcal{L}_{cut} = 4.263$ ($^{212}\text{BiPo}$) and 8.844 ($^{214}\text{BiPo}$) - leaves 3.24% of $^{212}\text{BiPo}$ and 2.81% of $^{214}\text{BiPo}$ pileup events remaining in the RoI and FV, or 0.056 and 0.015 events per year respectively, based on the values given in Table 6.1.

6.3.3 Post- $\Delta\mathcal{L}$ Cut BiPo Pileup Events

Figures 6.21 and 6.22 show the distributions of δt_{true} for $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ pileup events respectively that pass the cut on $\Delta\mathcal{L}$, when allowing for a 1% $^{130}\text{Te } 0\nu\beta\beta$ signal loss.

The Log-Likelihood Difference method suffers from the same basic shortcomings as the Cumulative Time Residuals method: when δt is either very small or very large, the BiPo pileup time residual distribution becomes indistinguishable in shape from that of a $^{130}\text{Te } 0\nu\beta\beta$ event. However, the Log-Likelihood Difference method is able to work correctly at slightly lower δt than the Cumulative Time Residuals method - 10ns for the former compared to 15ns in the latter.

6.4 Systematic Uncertainties

Some of the sources of systematic uncertainty in both the Cumulative Time Residuals and Log-Likelihood Differences methods - such as a non-ideal detector's geometry, and the errors in the PMT timing calibration - have previously been discussed in Section 5.6, and will impact these analyses in a similar fashion.

The overall noise rate of the detector will also contribute an uncertainty, since PMTs that were triggered due to noise are still included in the time residual distribution of a single event, and consequently will affect the shape of the distribution, and the comparisons of it with the reference PDF(s) of each method. Additionally, if the noise rate is much higher than expected, there will be a larger number of early triggered PMTs, and so the 10th earliest time residual may no longer be a consistent value




Figure 6.21: The time separation δt_{true} between the Bi β and Po α in $^{212}\text{BiPo}$ pileup events that pass the $\Delta\mathcal{L}$ cut at a 1% ^{130}Te $0\nu\beta\beta$ signal loss.




Figure 6.22: The time separation δt_{true} between the Bi β and Po α in $^{214}\text{BiPo}$ pileup events that pass the $\Delta\mathcal{L}$ cut at a 1% ^{130}Te $0\nu\beta\beta$ signal loss.

across different events (as in Section 6.1.4 above) - if not corrected for, this will result in some events' time residual distributions being incorrectly offset in comparison to the reference time residual PDFs in the Log-Likelihood Difference method.

The properties of the Te_LABPPO loaded scintillator are also important to consider when identifying sources of uncertainties. Not only will there be some error on the time residual calculation due to uncertainties on the velocity of light in Te_LABPPO (by way of its refractive index), but the cocktail's light yield and quenching properties are directly related to the quantity A in Equation 6.5 - and therefore any uncertainty in the measurements of these properties will propagate through to the calculations of $\mathcal{L}_{\text{BiPo}}$, $\Delta\mathcal{L}$ and the expected remaining signal and background fractions for a given $\Delta\mathcal{L}_{\text{cut}}$.

The detector's overall performance - the energy, position and time resolution - is the source of one of the most important uncertainties in any analysis performed at SNO+. The energy resolution dictates how effective the RoI cut (discussed previously in Section 6.1.1) will be, and thus the properties of any post-RoI cut events used in the methods detailed in this analysis. The calculation of the time residuals requires the reconstructed position and time of each event, and so the reconstruction performance of these quantities introduces an error into the time residual, and everything that follows from there.

6.5 Conclusion

Two of the major backgrounds that will be found in the same energy region of interest and spatial fiducial volume as the ^{130}Te $0\nu\beta\beta$ decay signal in the loaded phase of SNO+ are the ($^{212}\text{Bi} + ^{212}\text{Po}$) and ($^{214}\text{Bi} + ^{214}\text{Po}$) decay sub-chains. Due to the very short life-times of both Po isotopes, it is possible for these sub-chains to fully occur

within a single 400ns-wide trigger window, resulting in a BiPo pileup event which could be mistaken for a ^{130}Te $0\nu\beta\beta$ decay event. However, two methods have been developed can discriminate between these pileup backgrounds and the signal.

The Cumulative Time Residuals method is the simpler of the two methods to run in analysis, requiring only that the shape of the cumulative time residual distribution of ^{130}Te $0\nu\beta\beta$ decays be known in advance. The exact rejection factor will depend on the level of accepted signal loss that is chosen, but at a 1% signal loss this method can reject $\sim 95\%$ of both types of BiPo pileup event.

The Log-Likelihood Difference method gives a slightly higher BiPo pileup rejection factor of $\sim 97\%$ for the same 1% level of signal loss, and in addition can also return the exact value of δt - the time difference between the emission of the Bi β and the Po α - to a fairly high precision on the order of a few nanoseconds. However, this method does require more prior knowledge: the time residual distributions of ^{130}Te $0\nu\beta\beta$, Bi β and Po α decays separately, as well as the mean number of triggered PMTs in a single Po α decay.

It should be noted that neither of these methods rely upon the signal and background specifically being $0\nu\beta\beta$ or BiPo pileup respectively - the same principles can be applied in order to distinguish between any potential signal and background decays, or even between two different types of background, as long as they have sufficiently different time residual distributions.

Appendix A

Rayleigh Scattering Formulation

The following is a derivation of the Rayleigh scattering formulation from basic electro-dynamical principles, rather than as a special case of Mie Theory.

Consider a dielectric sphere, i.e. the scattering particle, of refractive index n_{par} and radius a located in a uniform electric field of magnitude E_0 . This initial setup is based on fulfillment of the Rayleigh criteria given in Equations 3.11 and 3.12. The magnitude of the induced dipole moment on the sphere is given by [116]:

$$\begin{aligned} P_0 &= x |\underline{E}| \\ &= 4\pi a^3 \epsilon_0 \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right) E_0 \end{aligned} \tag{A.1}$$

where x is known as the *electric polarisability* of the scattering particle.

The scattered irradiance, i.e. the energy received per unit time per unit area, at a distance r from the scattering particle (equivalent to the intensity of scattered light at r) is given by Equation A.2 for incident photon polarisation \underline{p} perpendicular to the scattering plane, and Equation A.3 for \underline{p} parallel to the scattering plane [117].

$$|\langle S_{sca} \rangle|_{\perp} = I_{\perp}(r, \theta, \phi) = \frac{\mu_0 \omega^4 P_0^2}{32\pi^2 c r^2} \quad (\text{A.2})$$

$$|\langle S_{sca} \rangle|_{\parallel} = I_{\parallel}(r, \theta, \phi) = \frac{\mu_0 \omega^4 P_0^2}{32\pi^2 c r^2} \cos^2 \theta \quad (\text{A.3})$$

Substituting Equation A.1 into these, and then simplifying using Equation 3.4, gives:

$$I_{\perp}(r, \theta, \phi) = \frac{I_0}{r^2} \times \frac{16\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 \quad (\text{A.4})$$

$$I_{\parallel}(r, \theta, \phi) = \frac{I_0}{r^2} \times \frac{16\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 \cos^2 \theta \quad (\text{A.5})$$

Comparing these to Equations 3.5 and 3.6 respectively shows that the differential cross sections for Rayleigh scattering of photons with incident polarisation \underline{p} perpendicular and parallel to the scattering plane are respectively given by:

$$\sigma'_{\perp}(\theta, \phi) = \frac{16\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 \quad (\text{A.6})$$

$$\sigma'_{\parallel}(\theta, \phi) = \frac{16\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 \cos^2 \theta \quad (\text{A.7})$$

These differential cross sections exhibit the well-known λ^{-4} dependence of Rayleigh scattering, as well as the $\cos^2 \theta$ distribution of scattered light when \underline{p} is parallel to the scattering plane. When \underline{p} is perpendicular to the scattering plane however, there is no dependence of the scattered light intensity on θ . There is also no dependence in either case on the second scattering angle ϕ .

Using Equations 3.7 and 3.8, the differential cross section and intensity of Rayleigh scattered unpolarised light are respectively given by:

$$\sigma'_{sca}(\theta, \phi) = \frac{8\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 (1 + \cos^2\theta) \quad (\text{A.8})$$

$$I_{sca}(r, \theta, \phi) = \frac{I_0}{r^2} \frac{8\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 (1 + \cos^2\theta) \quad (\text{A.9})$$

As noted in Equations 3.18 and 3.19, the Rayleigh scattering length l_{Ray} is defined by:

$$\begin{aligned} l_{Ray} &= \langle d \rangle \\ &= \left(\iint \beta(r, \theta, \phi) d\Omega \right)^{-1} \end{aligned} \quad (\text{A.10})$$

where the volume scattering function is given by [76]:

$$\beta(r, \theta, \phi) = \frac{I_{sca}(r, \theta, \phi)}{I_0} \frac{r^2}{V} \quad (\text{A.11})$$

and the *Rayleigh ratio* R is equal to the volume scattering function at a scattering angle of $\theta = 90^\circ$ [77]:

$$\begin{aligned} R &= \beta(\theta = 90^\circ) \\ &= \frac{I_{sca}(\theta = 90^\circ)}{I_0} \frac{r^2}{V} \end{aligned} \quad (\text{A.12})$$

Rearranging Equation A.12 for $\frac{r^2}{V}$ and substituting this and Equation A.9 back into Equation A.11 yields the volume scattering function for Rayleigh scattering from single particles:

$$\beta_{sca}(r, \theta, \phi) = \frac{1}{V} \frac{8\pi^4 a^6}{\lambda^4} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 (1 + \cos^2 \theta) \quad (\text{A.13})$$

and, using Equation A.10, the Rayleigh scattering length for single particles is therefore:

$$l_{Ray} = \left(\frac{128\pi^5 a^6}{3\lambda^4} \frac{1}{V} \left(\frac{n_{par}^2 - 1}{n_{par}^2 + 2} \right)^2 \right)^{-1} \quad (\text{A.14})$$

Appendix B

Extra Results: Section 5.2.4

The tables in this appendix show the optimised values of the four free parameters that maximise N (given by Equation 5.4), along with the corresponding values of purity and efficiency, for a single value of n_γ , but varying fibre / wavelength combinations. For each combination, the specific contributions to the acceptance region from the various PMT classifications are also given.

A discussion of these results is given in the main text of Section 5.2.4.

	Fibre ID	FS007	FS107	FS207	FS025	FS125	FS225
	C_{bottom}	241	242	242	241	242	242
Optimised	m_{bottom}	11	11	11	11	11	11
Parameters	C_{top}	220	220	215	220	220	215
	m_{top}	51	50	46	51	49	47
	N	215.080	216.701	180.300	218.741	215.831	185.489
Results	Purity	0.774	0.781	0.764	0.777	0.784	0.764
	Efficiency	0.069	0.068	0.060	0.070	0.066	0.063
	Beamsplit	0 (0.00)	0 (0.00)	12 (0.09)	0 (0.00)	0 (0.00)	11 (0.08)
Number	AV-reflected	230 (1.32)	216 (1.28)	219 (1.66)	271 (1.53)	176 (1.08)	201 (1.45)
(Percentage)	Concentrator	100 (0.57)	95 (0.56)	75 (0.57)	113 (0.64)	95 (0.58)	55 (0.40)
of PMTs	-reflected						
	Cavity	617 (3.55)	564 (3.35)	530 (4.03)	603 (3.40)	525 (3.22)	547 (3.94)
in	-scattered						
Acceptance	InnerAV	13473	13160 (78.11)	10050 (76.39)	13772 (77.65)	12807 (78.44)	10607 (76.44)
Region	-scattered	(77.44)					
	Multiple	2027 (11.65)	1953 (11.59)	1584 (12.04)	2016 (11.37)	1857 (11.37)	1756 (12.65)
	-effect						
	Noise	950 (5.46)	860 (5.10)	687 (5.22)	961 (5.42)	868 (5.32)	700 (5.04)

Table B.1: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs. Each simulation contained 5000 laser head pulses with $n_\gamma = 1500$, emitted from the given fibre at a wavelength of 375nm.

	Fibre ID	FS137	FS037	FS237	FS255	FS055	FS155
Optimised Parameters	C_{bottom}	241	242	242	240	242	242
	m_{bottom}	11	11	11	12	11	11
	C_{top}	220	219	218	220	216	218
	m_{top}	51	49	47	50	50	47
	N	217.038	212.932	196.772	218.235	210.640	202.546
Results	Purity	0.774	0.786	0.771	0.786	0.786	0.780
	Efficiency	0.069	0.065	0.064	0.066	0.064	0.066
Number (Percentage) of PMTs in Acceptance Region	Beamspot	0 (0.00)	0 (0.00)	14 (0.09)	0 (0.00)	16 (0.10)	23 (0.16)
	AV-reflected	273 (1.54)	125 (0.80)	239 (1.61)	192 (1.16)	208 (1.35)	197 (1.33)
	Concentrator -reflected	120 (0.68)	71 (0.45)	82 (0.55)	88 (0.53)	105 (0.68)	82 (0.55)
	Cavity -scattered	649 (3.66)	494 (3.15)	530 (3.56)	578 (3.50)	450 (2.92)	483 (3.27)
	InnerAV -scattered	13737 (77.42)	12328 (78.62)	11472 (77.14)	12990 (78.57)	12115 (78.55)	11538 (78.05)
	Multiple -effect	1963 (11.06)	1825 (11.64)	1790 (12.04)	1798 (10.88)	1686 (10.93)	1712 (11.58)
	Noise	1001 (5.64)	828 (5.28)	744 (5.00)	887 (5.37)	843 (5.47)	748 (5.06)

Table B.2: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs. Each simulation contained 5000 laser head pulses with $n_\gamma = 1500$, emitted from the given fibre at a wavelength of 375nm.

	Fibre ID	FS007	FS107	FS207	FS025	FS125	FS225
	C_{bottom}	241	242	242	242	242	242
Optimised	m_{bottom}	11	11	11	11	11	11
Parameters	C_{top}	220	220	215	220	220	216
	m_{top}	51	49	46	51	50	46
	N	199.700	193.890	166.77	191.601	197.342	167.464
Results	Purity	0.781	0.783	0.771	0.775	0.786	0.768
	Efficiency	0.072	0.069	0.063	0.069	0.070	0.064
	Beamsport	0 (0.00)	0 (0.00)	13 (0.12)	0 (0.00)	0 (0.00)	26 (0.24)
Number	AV-reflected	226 (1.58)	197 (1.48)	167 (1.56)	198 (1.44)	177 (1.31)	178 (1.61)
(Percentage)	Concentrator	89 (0.62)	91 (0.69)	56 (0.52)	100 (0.73)	95 (0.70)	77 (0.70)
of PMTs	-reflected						
	Cavity	479 (3.35)	407 (3.07)	400 (3.74)	400 (2.91)	394 (2.91)	403 (3.65)
in	-scattered						
Acceptance	InnerAV	11160	10398 (78.33)	8253 (77.12)	10661 (77.50)	10629 (78.56)	8483 (76.78)
Region	-scattered	(78.13)					
	Multiple	1432 (10.03)	1314 (9.90)	1143 (10.68)	1455 (10.58)	1359 (10.04)	1243 (11.25)
	-effect						
	Noise	897 (6.28)	867 (6.53)	670 (6.26)	943 (6.85)	876 (6.47)	639 (5.78)

Table B.3: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs. Each simulation contained 5000 laser head pulses with $n_\gamma = 1500$, emitted from the given fibre at a wavelength of 405nm.

	Fibre ID	FS137	FS037	FS237	FS255	FS055	FS155
	c_{bottom}	242	241	242	242	242	242
Optimised	m_{bottom}	11	11	11	11	11	11
Parameters	c_{top}	220	218	217	220	216	218
	m_{top}	51	49	48	50	49	47
	N	193.598	191.019	178.621	200.593	186.844	183.027
Results	Purity	0.778	0.780	0.771	0.790	0.786	0.778
	Efficiency	0.069	0.070	0.068	0.069	0.064	0.069
	Beamspot	0 (0.00)	0 (0.00)	2 (0.02)	0 (0.00)	1 (0.01)	28 (0.23)
Number	AV-reflected	192 (1.40)	182 (1.38)	219 (1.78)	214 (1.58)	198 (1.64)	170 (1.39)
(Percentage)	Concentrator	107 (0.78)	93 (0.70)	81 (0.66)	82 (0.61)	88 (0.73)	77 (0.63)
of PMTs	-reflected						
in	Cavity	397 (2.89)	461 (3.49)	386 (3.14)	378 (2.79)	379 (3.14)	395 (3.22)
Acceptance	-scattered						
Region	InnerAV	10676	10292 (78.00)	9477 (77.10)	10705 (78.99)	9485 (78.64)	9535 (77.84)
	-scattered	(77.83)					
	Multiple	1414 (10.31)	1296 (9.82)	1367 (11.12)	1333 (9.84)	1144 (9.48)	1302 (10.63)
	-effect						
	Noise	931 (6.79)	867 (6.57)	760 (6.18)	841 (6.21)	767 (6.36)	742 (6.06)

Table B.4: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs. Each simulation contained 5000 laser head pulses with $n_\gamma = 1500$, emitted from the given fibre at a wavelength of 405nm.

	Fibre ID	FS007	FS107	FS207	FS025	FS125	FS225
Optimised Parameters	c_{bottom}	242	240	242	242	242	242
	m_{bottom}	11	12	11	11	11	11
	c_{top}	220	220	213	220	220	214
	m_{top}	51	49	48	50	49	48
	N	142.619	146.300	123.112	140.192	143.366	127.265
Results	Purity	0.750	0.759	0.729	0.747	0.753	0.739
	Efficiency	0.069	0.071	0.067	0.068	0.070	0.068
Number (Percentage) of PMTs in Acceptance Region	Beamsplit	0 (0.00)	0 (0.00)	2 (0.03)	0 (0.00)	0 (0.00)	6 (0.08)
	AV-reflected	187 (2.07)	156 (1.74)	228 (2.95)	216 (2.43)	186 (2.08)	194 (2.51)
	Concentrator -reflected	65 (0.72)	54 (0.60)	50 (0.65)	54 (0.61)	56 (0.63)	38 (0.49)
	Cavity -scattered	271 (3.00)	302 (3.37)	297 (3.84)	261 (2.94)	280 (3.14)	278 (3.59)
	InnerAV -scattered	6768 (75.03)	6801 (75.89)	5639 (72.88)	6641 (74.74)	6726 (75.34)	5717 (73.91)
	Multiple -effect	828 (9.18)	802 (8.95)	836 (10.81)	808 (9.09)	783 (8.77)	836 (10.81)
	Noise	901 (9.99)	847 (9.45)	685 (8.85)	905 (10.19)	896 (10.04)	666 (8.61)

Table B.5: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs. Each simulation contained 5000 laser head pulses with $n_\gamma = 1500$, emitted from the given fibre at a wavelength of 440nm.

	Fibre ID	FS137	FS037	FS237	FS255	FS055	FS155
	c_{bottom}	240	242	242	242	242	242
Optimised	m_{bottom}	12	11	11	11	11	11
Parameters	c_{top}	220	219	217	220	217	218
	m_{top}	51	49	48	51	48	47
	N	143.716	142.347	133.887	138.861	138.183	131.771
Results	Purity	0.748	0.758	0.746	0.738	0.757	0.742
	Efficiency	0.071	0.068	0.068	0.069	0.065	0.069
	Beamsplit	0 (0.00)	8 (0.09)	0 (0.00)	0 (0.00)	5 (0.06)	15 (0.18)
Number	AV-reflected	136 (1.46)	136 (1.59)	195 (2.39)	270 (2.91)	200 (2.46)	165 (2.02)
(Percentage)	Concentrator	63 (0.68)	49 (0.57)	42 (0.51)	59 (0.63)	64 (0.79)	60 (0.74)
of PMTs	-reflected						
in	Cavity	305 (3.28)	305 (3.58)	270 (3.31)	297 (3.20)	235 (2.90)	260 (3.19)
Acceptance	-scattered						
Region	InnerAV	6955 (74.81)	6467 (75.81)	6090 (74.64)	6855 (73.77)	6141 (75.67)	6050 (74.16)
	-scattered						
	Multiple	894 (9.62)	759 (8.90)	820 (10.05)	874 (9.41)	755 (9.30)	841 (10.31)
	-effect						
	Noise	944 (10.15)	807 (9.46)	742 (9.09)	937 (10.08)	716 (8.82)	767 (9.40)

Table B.6: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs. Each simulation contained 5000 laser head pulses with $n_\gamma = 1500$, emitted from the given fibre at a wavelength of 440nm.

	Fibre ID	FS007	FS107	FS207	FS025	FS125	FS225
Optimised Parameters	c_{bottom}	238	241	241	238	239	241
	m_{bottom}	12	11	11	12	12	11
	c_{top}	220	220	214	220	220	217
	m_{top}	50	49	46	50	49	43
Results	N	80.004	79.771	71.454	80.285	78.343	71.248
	Purity	0.661	0.672	0.663	0.660	0.663	0.665
	Efficiency	0.061	0.059	0.056	0.062	0.059	0.055
Number (Percentage) of PMTs in Acceptance Region	Beamsplit	0 (0.00)	0 (0.00)	7 (0.18)	0 (0.00)	0 (0.00)	4 (0.10)
	AV-reflected	146 (2.94)	162 (3.52)	142 (3.62)	142 (2.82)	126 (2.68)	162 (4.22)
	Concentrator -reflected	18 (0.36)	32 (0.69)	23 (0.59)	26 (0.52)	26 (0.55)	17 (0.44)
	Cavity -scattered	163 (3.28)	117 (2.53)	151 (3.85)	188 (3.74)	153 (3.25)	131 (3.41)
	InnerAV -scattered	3286 (66.08)	3108 (67.19)	2599 (66.27)	3319 (66.01)	3118 (66.31)	2554 (66.53)
	Multiple -effect	390 (7.84)	350 (7.57)	361 (9.20)	400 (7.96)	414 (8.80)	318 (8.28)
	Noise	970 (19.51)	857 (18.53)	639 (16.29)	953 (18.95)	865 (18.40)	653 (17.01)

Table B.7: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs. Each simulation contained 5000 laser head pulses with $n_\gamma = 1500$, emitted from the given fibre at a wavelength of 500nm.

	Fibre ID	FS137	FS037	FS237	FS255	FS055	FS155
	c_{bottom}	239	241	241	238	241	241
Optimised	m_{bottom}	12	11	11	12	11	11
Parameters	c_{top}	220	219	217	220	212	218
	m_{top}	50	48	47	49	50	45
	N	79.837	80.402	73.643	79.750	74.084	74.752
Results	Purity	0.663	0.683	0.659	0.666	0.668	0.673
	Efficiency	0.060	0.058	0.057	0.059	0.053	0.057
	Beamsplit	0 (0.00)	3 (0.07)	4 (0.09)	0 (0.00)	0 (0.00)	1 (0.02)
Number	AV-reflected	130 (2.66)	116 (2.64)	159 (3.73)	175 (3.66)	148 (3.63)	179 (4.42)
(Percentage)	Concentrator	36 (0.74)	32 (0.73)	25 (0.59)	23 (0.48)	22 (0.54)	20 (0.49)
of PMTs	-reflected						
	Cavity	163 (3.33)	137 (3.12)	125 (2.93)	149 (3.11)	121 (2.97)	120 (2.97)
in	-scattered						
Acceptance	InnerAV	3243 (66.28)	3003 (68.28)	2811 (65.86)	3187 (66.62)	2722 (66.85)	2721 (67.25)
Region	-scattered						
	Multiple	392 (8.01)	343 (7.80)	376 (8.81)	361 (7.55)	307 (7.54)	303 (7.49)
	-effect						
	Noise	929 (18.99)	764 (17.37)	768 (17.99)	889 (18.58)	752 (18.47)	702 (17.35)

Table B.8: The optimised free parameters of the acceptance cuts given in Equations 5.2 and 5.3 that bound the acceptance region, the resulting purity and efficiency values, and the corresponding contributions to the acceptance region from the various optical processes as well as noise PMTs. Each simulation contained 5000 laser head pulses with $n_\gamma = 1500$, emitted from the given fibre at a wavelength of 500nm.

Appendix C

Extra Results: Section 5.4.2 - 5.4.4

The following tables show the reconstruction performance of the pencil beam analysis across all fibre / wavelength combinations of the simulated SMELLIE system in two reconstruction scenarios:

1. using the true position of incidence on the PMT and the true path time of the scattered photon
4. using the position at the centre of the PMT face and the calculated path time from Equation 5.13

Tables C.1 and C.2 show the bias and resolution respectively for the reconstruction of the scattering position \underline{s} , where the bias is calculated using Equation 5.16, and the resolution by Equation 5.17. Tables C.3 and C.4 give the results for the reconstruction of d - the distance travelled by a photon before undergoing scattering. Finally, the bias and resolution of the reconstruction of the scattering angle θ are shown in Tables C.5 and C.6 respectively. For d and θ , the bias and resolution are found by taking the mean and standard deviation respectively of a Gaussian function fitted to the distributions of $(d_{reco} - d_{true})$ and $(\theta_{reco} - \theta_{true})$.

A discussion of these results is given in the main text of Sections 5.4.2 to 5.4.4.

	375nm	405nm	440nm	500nm
FS007	0.373 (223.528)	0.472 (439.659)	0.350 (30.537)	0.012 (408.038)
FS107	0.343 (421.789)	0.460 (34.362)	0.337 (32.240)	0.038 (14.570)
FS207	0.282 (36.970)	0.497 (38.293)	0.344 (36.787)	0.046 (25.975)
FS025	0.382 (23.583)	0.468 (35.345)	0.321 (464.296)	0.020 (17.174)
FS125	0.340 (33.457)	0.498 (37.798)	0.371 (37.742)	0.014 (30.921)
FS225	0.247 (42.137)	0.464 (199.871)	0.350 (43.054)	0.036 (290.396)
FS137	0.621 (44.360)	0.475 (46.439)	0.316 (51.774)	0.016 (35.461)
FS037	0.330 (78.207)	0.494 (80.761)	0.378 (75.020)	0.005 (52.255)
FS237	0.292 (62.179)	0.494 (75.869)	0.346 (66.946)	0.026 (56.783)
FS255	0.385 (49.848)	0.443 (307.368)	0.326 (44.936)	0.025 (33.601)
FS055	0.290 (42.979)	0.468 (52.504)	2.421 (37.218)	0.043 (34.608)
FS155	0.601 (72.784)	0.669 (77.051)	0.362 (73.537)	0.081 (313.193)

Table C.1: The bias (in mm), calculated using Equation 5.16, of \underline{s}_{reco} (the reconstructed scattering position) of accepted-scattered photons emitted in a pencil beam, in reconstruction scenarios 1 (unbracketed) and 4 (bracketed) as defined in the main text.

	375nm	405nm	440nm	500nm
FS007	1.004 (288.351)	1.029 (293.265)	1.312 (282.263)	1.275 (297.547)
FS107	0.969 (289.558)	0.973 (270.289)	1.296 (284.990)	1.291 (281.098)
FS207	0.959 (278.980)	0.927 (276.844)	1.222 (280.109)	1.195 (279.538)
FS025	1.000 (291.712)	1.027 (279.318)	1.331 (292.976)	1.340 (287.455)
FS125	1.021 (292.011)	1.002 (282.078)	1.329 (292.621)	1.326 (288.931)
FS225	0.902 (269.169)	0.900 (264.032)	1.198 (262.076)	1.136 (238.473)
FS137	0.854 (298.270)	1.007 (278.554)	1.377 (289.389)	1.300 (276.084)
FS037	0.997 (295.288)	0.975 (286.653)	1.303 (286.079)	1.264 (285.721)
FS237	0.974 (278.808)	0.949 (266.277)	1.255 (274.380)	1.237 (268.848)
FS255	0.990 (272.100)	0.972 (268.070)	1.289 (270.684)	1.309 (274.585)
FS055	0.944 (258.953)	0.909 (258.110)	16.874 (257.347)	1.186 (252.828)
FS155	0.870 (282.200)	0.838 (275.409)	1.195 (278.260)	1.173 (251.467)

Table C.2: The resolution (in mm), calculated using Equation 5.17, of \underline{s}_{reco} (the reconstructed scattering position) of accepted-scattered photons emitted in a pencil beam, in reconstruction scenarios 1 (unbracketed) and 4 (bracketed) as defined in the main text.

	375nm	405nm	440nm	500nm
FS007	-0.252 (-222.214)	0.349 (-442.528)	0.287 (25.084)	-0.045 (-419.394)
FS107	-0.218 (-447.424)	0.375 (28.560)	0.320 (29.334)	0.055 (5.771)
FS207	0.120 (31.676)	0.748 (33.596)	0.763 (28.053)	0.478 (14.024)
FS025	-0.258 (15.013)	0.354 (28.824)	0.245 (-450.470)	-0.010 (8.169)
FS125	-0.206 (23.752)	0.379 (28.309)	0.341 (27.948)	0.058 (25.988)
FS225	0.123 (40.386)	0.733 (-197.025)	0.779 (43.192)	0.520 (-288.925)
FS137	-0.254 (36.482)	0.350 (40.726)	0.250 (44.291)	0.002 (26.177)
FS037	-0.165 (67.382)	0.445 (66.576)	0.354 (63.245)	0.103 (45.452)
FS237	-0.058 (52.414)	0.553 (70.791)	0.540 (61.601)	0.265 (45.396)
FS255	-0.268 (43.476)	0.347 (-315.491)	0.295 (40.192)	0.017 (27.275)
FS055	-0.141 (37.832)	0.457 (50.616)	0.365 (37.115)	0.152 (33.469)
FS155	-0.005 (59.127)	0.594 (65.792)	0.589 (64.312)	0.322 (-286.768)

Table C.3: The bias (in mm) of d_{reco} (the reconstructed distance travelled before a photon undergoes Rayleigh scattering) of accepted-scattered photons emitted in a pencil beam, in reconstruction scenarios 1 (unbracketed) and 4 (bracketed) as defined in the main text.

	375nm	405nm	440nm	500nm
FS007	0.625 (272.077)	0.744 (279.989)	1.034 (256.560)	1.010 (296.115)
FS107	0.646 (294.402)	0.702 (253.170)	1.108 (264.315)	1.102 (252.708)
FS207	0.946 (258.814)	0.944 (252.343)	1.284 (255.680)	1.237 (245.436)
FS025	0.644 (266.588)	0.742 (257.174)	1.055 (266.915)	1.068 (259.885)
FS125	0.655 (259.203)	0.716 (249.586)	1.058 (259.130)	1.085 (257.388)
FS225	0.986 (261.319)	0.918 (258.186)	1.289 (257.375)	1.232 (231.695)
FS137	0.642 (266.365)	0.716 (248.561)	1.102 (249.070)	1.091 (242.885)
FS037	0.651 (264.502)	0.703 (254.428)	1.051 (250.986)	1.077 (242.268)
FS237	0.781 (263.799)	0.762 (250.636)	1.148 (254.661)	1.132 (237.327)
FS255	0.637 (256.221)	0.735 (261.097)	1.087 (252.589)	1.119 (240.989)
FS055	0.621 (249.913)	0.704 (251.227)	0.996 (246.915)	1.026 (241.720)
FS155	0.790 (253.785)	0.792 (251.857)	1.130 (252.288)	1.094 (227.748)

Table C.4: The resolution (in mm) of d_{reco} (the reconstructed distance travelled before a photon undergoes Rayleigh scattering) of accepted-scattered photons emitted in a pencil beam, in reconstruction scenarios 1 (unbracketed) and 4 (bracketed) as defined in the main text.

	375nm	405nm	440nm	500nm
FS007	-0.027 (-1.594)	0.165 (-3.212)	0.163 (0.195)	0.084 (-3.134)
FS107	-0.075 (-3.298)	0.142 (0.256)	0.143 (0.242)	0.036 (0.067)
FS207	-0.085 (0.219)	0.134 (0.227)	0.094 (0.243)	0.001 (0.220)
FS025	-0.022 (0.196)	0.146 (0.221)	0.165 (-3.547)	0.102 (0.153)
FS125	-0.056 (0.204)	0.136 (0.243)	0.154 (0.267)	0.051 (0.127)
FS225	-0.116 (0.381)	0.112 (-1.385)	0.095 (0.363)	-0.001 (-2.217)
FS137	-0.020 (0.283)	0.173 (0.294)	0.154 (0.264)	0.106 (0.250)
FS037	-0.104 (0.480)	0.126 (0.522)	0.120 (0.489)	0.024 (0.384)
FS237	-0.118 (0.437)	0.084 (0.501)	0.073 (0.470)	-0.039 (0.353)
FS255	-0.020 (0.357)	0.163 (-2.308)	0.161 (0.340)	0.108 (0.267)
FS055	-0.055 (0.354)	0.151 (0.427)	0.136 (0.337)	0.056 (0.318)
FS155	-0.090 (0.501)	0.122 (0.537)	0.111 (0.491)	0.038 (-2.564)

Table C.5: The bias (in $^\circ$) of θ_{reco} (the reconstructed Rayleigh scattering angle θ) of accepted-scattered photons emitted in a pencil beam, in reconstruction scenarios 1 (unbracketed) and 4 (bracketed) as defined in the main text.

	375nm	405nm	440nm	500nm
FS007	0.696 (2.579)	0.669 (2.753)	0.665 (2.192)	0.690 (2.946)
FS107	0.677 (2.919)	0.692 (2.176)	0.718 (2.217)	0.668 (2.229)
FS207	0.750 (2.179)	0.756 (2.120)	0.791 (2.183)	0.823 (2.345)
FS025	0.685 (2.289)	0.636 (2.167)	0.666 (3.025)	0.728 (2.303)
FS125	0.701 (2.208)	0.703 (2.139)	0.703 (2.241)	0.737 (2.306)
FS225	0.766 (2.326)	0.724 (2.407)	0.767 (2.232)	0.749 (2.676)
FS137	0.674 (2.322)	0.679 (2.109)	0.662 (2.168)	0.715 (2.117)
FS037	0.710 (2.181)	0.731 (2.193)	0.734 (2.235)	0.742 (2.304)
FS237	0.691 (2.241)	0.750 (2.176)	0.760 (2.254)	0.749 (2.269)
FS255	0.658 (2.236)	0.656 (2.534)	0.674 (2.255)	0.726 (2.379)
FS055	0.693 (2.141)	0.689 (2.182)	0.690 (2.211)	0.713 (2.049)
FS155	0.679 (2.214)	0.755 (2.163)	0.763 (2.241)	0.764 (2.868)

Table C.6: The resolution (in $^\circ$) of θ_{reco} (the reconstructed Rayleigh scattering angle θ) of accepted-scattered photons emitted in a pencil beam, in reconstruction scenarios 1 (unbracketed) and 4 (bracketed) as defined in the main text.

Bibliography

- [1] J. Boger et al., *The Sudbury Neutrino Observatory*, Nucl. Instrum. Methods Phys. Res., Sect. A **449** (2000) 172–207.
- [2] A. Hallin, M. Hedayatipour, and K. Singh, *The Expected Performance and Motion of the Upward and Downward Ropes in the SNO+ Rope Net*, SNO+ DocDB **1872** (2013) [Version 1].
- [3] N. Fatemighomi, *Cover Gas Commissioning Report*, SNO+ DocDB **2918** (2015) [Version 1].
- [4] G. Doucas et al., *Light Concentrators for the Sudbury Neutrino Observatory*, Nucl. Instrum. Methods Phys. Res., Sect. A **370** (1996) 579–596.
- [5] A. Mastbaum, *Electronics Overview and Status*, SNO+ DocDB **1339** (2012) [Version 1].
- [6] A. Mastbaum, *Electronics Update*, SNO+ DocDB **2073** (2013) [Version 1].
- [7] J. B. Birks et al., *The theory and practice of scintillation counting*. Pergamon Press, 1964.
- [8] G. Laustriat, *The luminescence decay of organic scintillators*, Molecular Crystals **4** (1968) 127–145.
- [9] M. Yeh, *Radiopure Metal-doped Liquid Scintillator*, . (Presentation at the Low Radioactivity Techniques Conference 2015).

- [10] L. Segui, *Scintillator purification and development in the SNO+ Experiment*, .
(Presentation at the IoP HEPP & APP Conference 2015).
- [11] I. T. Coulter, *Modelling and reconstruction of events in SNO+ related to future searches for lepton and baryon number violation*. PhD thesis, 2013.
- [12] M. Berglund and M. E. Wieser, *Isotopic compositions of the elements 2009*,
Pure Appl. Chem. **83** (2011) 397–410.
- [13] J. Chadwick, [*unknown title*], Verhandlungen der Deutschen Physikalischen
Gesellschaft **16** (1914) 383.
- [14] C. D. Ellis and W. A. Wooster, *The Average Energy of Disintegration of Radium E*,
Proceedings of the Royal Society **117** (1927) 109–123.
- [15] L. Meitner and W. Orthmann, *About an absolute Determination of the Power of the Primary β -Radiation from Radium E*,
Zeitschrift fur Physik **60** (1930) 143–155.
- [16] W. Pauli, *Open letter to the group of radioactive people at the Gauverein meeting in Tübingen*, 1930.
<http://microboone-docdb.fnal.gov/cgi-bin/RetrieveFile?docid=953;filename=pauli%20letter1930.pdf>.
- [17] J. Chadwick, *Possible Existence of a Neutron*, Nature **129** (1932) 312.
- [18] E. Fermi, *Attempt of a Theory of β -Radiation*, Zeitschrift fur Physik **88**
(1934) 161–177.
- [19] G. W. Rodeback and J. S. Allen, *Neutrino Recoils Following the Capture of Orbital Electrons in A^{37}* ,
Phys. Rev. **86** (1952) 446–450.
- [20] F. Reine and C. L. Cowan Jr., *Detection of the Free Neutrino*, Phys. Rev. **92**
(1953) 830–831.

- [21] C. L. Cowan Jr et al., *Detection of the Free Neutrino: a Confirmation*, Science **124** (1956) 103–104.
- [22] C. M. G. Lattes, G. P. S. Occhialini, and C. F. Powell, *Observations on the Tracks of Slow Mesons in Photographic Emulsions*, Nature **160** (1947) 453–456.
- [23] G. Danby et al., *Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos*, Phys. Rev. Lett. **9** (1962) 36–49.
- [24] J. K. Bienlein et al., *Spark chamber study of high-energy neutrino interactions*, Physics Letters **13** (1964) 80–86.
- [25] M. Goldhaber, L. Grodzins, and A. W. Sunyar, *Helicity of Neutrinos*, Phys. Rev. **109** (1958) 1015–1017.
- [26] M. L. Perl et al., *Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation*, Phys. Rev. Lett. **35** (1975) 1489–1491.
- [27] D. Decamp et al., *A precise determination of the number of families with light neutrinos and of the Z boson partial widths*, Phys. Lett. B **235** (1990) 399–411.
- [28] K. Kodama et al., *Observation of tau Neutrino interactions*, Phys. Lett. B **504** (2001) 218–224.
- [29] H. Bethe and R. Peierls, *The Neutrino*, Nature **133** (1934) 532.
- [30] K. Zuber, *Neutrino Physics (2nd Edition)*. CRC Press, 2012.
- [31] M. Schwartz, *Feasibility of Using High-Energy Neutrinos to Study the Weak Interactions*, Phys. Rev. Lett. **4** (1960) 306–307.
- [32] A. Gando et al., *Partial radiogenic heat model for Earth revealed by geoneutrino measurements*, Nature Geoscience **4** (2011) 647–651.

- [33] G. Bellini et al., *Observation of geo-neutrinos*, Phys. Lett. B **687** (2010) 299–304.
- [34] S. A. Colgate and R. H. White, *The Hydrodynamic Behavior of Supernovae Explosions*, The Astrophysical Journal **143** (1966) 626–675.
- [35] Y. Fukuda et al., *Measurement of a small atmospheric ν_μ/ν_e ratio*, Phys. Lett. B **433** (1998) 9–18.
- [36] Y. Fukuda et al., *Evidence for Oscillation of Atmospheric Neutrinos*, Phys. Rev. Lett. **81** (1998) 1562–1567.
- [37] S. Fukuda et al., *Tau Neutrinos Favored over Sterile Neutrinos in Atmospheric Muon Neutrino Oscillations*, Phys. Rev. Lett. **85** (2000) 3999–4003.
- [38] J. N. Bahcall, *Solar Neutrinos I - Theoretical*, Phys. Rev. Lett. **12** (1964) 300–302.
- [39] R. Davis Jr., *Solar Neutrinos II - Experimental*, Phys. Rev. Lett. **12** (1964) 303–305.
- [40] B. T. Cleveland et al., *Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector*, The Astrophysical Journal **496** (1998) 505–526.
- [41] Q. R. Ahmad et al., *Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **87** (2002) 071301.
- [42] S. Fukuda et al., *Solar ^8B and hep Neutrino Measurement from 1258 Days of Super-Kamiokande Data*, Phys. Rev. Lett. **86** (2001) 5651–5655.

- [43] Q. R. Ahmad et al., *Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **89** (2002) 011301.
- [44] Z. Maki, M. Nakagawa, and S. Sakata, *Remarks on the Unified Model of Elementary Particles*, Progress of Theoretical Physics **28** (1962) 870–880.
- [45] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2011.
- [46] K. A. Olive (Particle Data Group) et al., *Review of Particle Physics*, Chinese Physics C **38** (2014) 090001.
- [47] J. L. Hewett et al., *Fundamental Physics at the Intensity Frontier*, arXiv **1205.2671 [hep-ex]** (2012) .
- [48] S. Nagy et al., *On the Q -value of the tritium β -decay*, Europhys. Lett. **74** (2006) 404–410.
- [49] E. Majorana, *Symmetric theory of the electron and the positron*, Il Nuovo Cimento **14** (1937) 171–184.
- [50] M. Goeppert-Mayer, *Double-Beta Disintegration*, Phys. Rev. **48** (1935) 512–516.
- [51] G. Racah, *On Symmetry between Particles and Antiparticles*, Il Nuovo Cimento **14** (1937) 322–328.
- [52] W. H. Furry, *On Transition Probabilities in Double Beta-Disintegration*, Phys. Rev. **56** (1939) 1184–1193.
- [53] S. M. Bilenky and C. Giunti, *Neutrinoless Double-Beta Decay: A Brief Review*, Modern Physics Letters A **27** (2012) 1230015.

- [54] J. Suhonen and O. Civitarese, *Weak-interaction and nuclear-structure aspects of nuclear double beta decay*, Physics Reports **300** (1998) 123–214.
- [55] S. R. Elliott and J. Engel, *Double-beta decay*, J. Phys. G: Nucl. Part. Phys. **30** (2004) R183–R215.
- [56] H. V. Klapdor-Kleingrothaus and I. V. Krivosheina, *The Evidence for the Observation of $0\nu\beta\beta$: The Identification of $0\nu\beta\beta$ Events from the Full Spectra*, Modern Physics Letters A **21** (2006) 1547–1566.
- [57] B. Majorovits, *Phase II Upgrade of the GERDA Experiment for the Search of Neutrinoless Double Beta Decay*, Physics Procedia **61** (2015) 254–259.
- [58] P. Guzowski, *Status of the SuperNEMO experiment*, . (Presentation at the IoP HEPP & APP Conference 2014).
- [59] D. R. Artusa et al., *Searching for Neutrinoless Double-Beta Decay of ^{130}Te with CUORE*, Advances in High Energy Physics **2015** (2015) 879871.
- [60] I. Nasteva, *Neutrinoless double beta decay search with the NEMO 3 experiment*, AIP Conference Proceedings **1078** (2008) 332–334.
- [61] M. Agostini et al., *Results on Neutrinoless Double- β Decay of ^{76}Ge from Phase I of the GERDA Experiment*, Phys. Rev. Lett. **111** (2013) 122503.
- [62] R. Arnold et al., *Search for neutrinoless double-beta decay of ^{100}Mo with the NEMO-3 detector*, Phys. Rev. D: Part. Fields **89** (2014) 111101.
- [63] E. Andreotti et al., *^{130}Te neutrinoless double-beta decay with CUORICINO*, Astroparticle Physics **34** (2011) 822–831.
- [64] A. Gando et al., *Limit on Neutrinoless $\beta\beta$ Decay of ^{136}Xe from the First Phase of KamLAND-Zen and Comparison with the Positive Claim in ^{76}Ge* , Phys. Rev. Lett. **110** (2013) 062502.

- [65] S. Rahaman et al., *Double-beta decay Q values of ^{116}Cd and ^{130}Te* , Phys. Lett. B **703** (2011) 412–416.
- [66] M. Kerker, *The Scattering of Light and other Electromagnetic radiation*. Academic Press, 1969.
- [67] C. Raman, *A new radiation*, Indian J. Phys. **2** (1928) 387–398.
- [68] G. Mie, *Contributions to the Optics of Turbid Media, particularly of Colloidal Metal Solutions*, Annalen der Physik **330** (1908) 377–445.
- [69] H.C. van de Hulst, *Light Scattering by Small Particles*. Dover Publications, 1981.
- [70] Lord Rayleigh, F.R.S., *On the transmission of light through an atmosphere containing small particles in suspension, and on the origin of the blue of the sky*, Philosophical Magazine Series 5 **47** (1899) 375–384.
- [71] C. F. Bohren and D. R. Huffman, *Absorption and Scattering of Light by Small Particles*. John Wiley & Sons, 1983.
- [72] W. Wiscombe, *Improved Mie scattering algorithms*, Appl. Opt. **19** (1980) 1505–1509.
- [73] G. Grehan and G. Gouesbet, *Mie theory calculations: new progress, with emphasis on particle sizing*, Appl. Opt. **18** (1979) 3489–3493.
- [74] A. Gogoi, A. Choudhury, and G. A. Ahmed, *Mie scattering computation of spherical particles with very large size parameters using an improved program with variable speed and accuracy*, J. Mod. Opt. **57** (2010) 2192–2202.
- [75] X. Zhou et al., *Rayleigh scattering of linear alkylbenzene in large liquid scintillator detectors*, arXiv **1504.00987 [physics.ins-det]** (2015) .

- [76] N. G. Jerlov, *Optical Oceanography*. Elsevier, 1968.
- [77] A. Morel, *Optical aspects of oceanography [collection of the papers presented at the Symposium on Optical Aspects of Oceanography, held at the Institute of Physical Oceanography in Copenhagen, 19th-23rd June 1972]*. Academic Press, 1974.
- [78] R. Strutt, *The Light Scattered by Gases: Its Polarisation and Intensity*, Proceedings of the Royal Society **95** (1918) 155–176.
- [79] M. V. Smoluchowski, *Molecular kinetic theory of gases in critical opalescence state, as well as some related phenomena*, Annalen der Physik **330** (1908) 205–226.
- [80] A. Einstein, *Theory of opalescence of homogeneous fluids and liquid mixtures near the critical state*, Annalen der Physik **338** (1910) 1275–1298.
- [81] J. F. Eykman, *Refractometric Searches (continued)*, Recueil des Travaux Chimiques des Pays-Bas **14** (1895) 185–202.
- [82] J. Cabannes, *Relationship between the degree of Polarisation and the intensity of light Scattered by anisotropic molecules. New determination of the Avogadro constant*, Journal de Physique et Le Radium **1** (1920) 129–142.
- [83] W. Kaye and J. McDaniel, *Low-Angle Laser Light Scattering - Rayleigh Factors and Depolarization Ratios*, Appl. Opt. **13** (1974) 1934–1937.
- [84] M. Wurm et al., *Optical scattering lengths in large liquid-scintillator neutrino detectors*, Rev. Sci. Instrum. **81** (2010) 053301.
- [85] Q. Liu et al., *Rayleigh scattering and depolarization ratio in linear alkylbenzene*, arXiv **1504.01001 [physics.ins-det]** (2015) .

- [86] G. S. Kell, *Density, Thermal Expansivity, and Compressibility of Liquid Water from 0° to 150° C*, J. Chem. Eng. Data **20** (1975) 97–105.
- [87] M. Chen and I. T. Coulter, *Optical properties in RAT*, SNO+ DocDB **1664** (2012) [Version 1].
- [88] National Instruments, *Bus-Powered M Series Multifunction DAQ for USB*, 2014.
<http://sine.ni.com/ds/app/doc/p/id/ds-9/lang/en>.
- [89] PicoQuant GmbH, *LDH Series Picosecond Pulsed Diode Laser Heads*, 2014.
<http://www.picoquant.com/products/category/picosecond-pulsed-sources/ldh-series-picosecond-pulsed-diode-laser-heads>.
- [90] S. Biller, *SNO+ with Tellurium*, Physics Procedia **61** (2015) 205–210.
- [91] PicoQuant GmbH, *PDL 828 "Sepia II"*, 2014.
<http://www.picoquant.com/products/category/picosecond-pulsed-driver/pdl-828-sepia-ii-computer-controlled-multichannel-picosecond-diode-laser-driver>.
- [92] Corning, *Corning® InfiniCor® 50µm Optical Fibers Data Sheet*, 2008.
<http://www.princetel.com/datasheets/Infenicor50.pdf>.
- [93] Corning, *Corning Multimode Fiber Performance*, 2007.
<http://www.corning.com/docs/opticalfiber/co4249.pdf>.
- [94] K. Clark, *ELLIE Fibre Radioactivity*, SNO+ DocDB **1223** (2011) [Version 1].
- [95] G. Keiser, *Optical Fiber Communications (3rd Edition)*. McGraw-Hill, 2000.
- [96] Laser Components (UK) Ltd, *Fiber optic Switch Operation Manual*, [private communication] (2009) .

- [97] K. Clark, *SMELLIE Plate and Collimator Designs*, [private communication] (2011) .
- [98] DuPont, *DuPont Delrin acetal resin Molding Guide*, 2006.
http://plastics.dupont.com/plastics/pdflit/americas/delrin/DELFIN_Mold_Guide_11_06.pdf.
- [99] National Instruments, *What is Emitter Coupled Logic (ECL)?*, 2014.
<http://digital.ni.com/public.nsf/allkb/BDAC410DD4A7D98B86256D040065BBD1>.
- [100] J. R. Wilson, *ELLIE Generator*, SNO+ DocDB **2314** (2014) [Version 1].
- [101] S. Langrock, *Position Verification of SMELLIE in RAT*, SNO+ DocDB **2000** (2013) [Version 3].
- [102] S. Langrock, *SMELLIE dark runs March 2014*, SNO+ DocDB **2483** (2014) [Version 1].
- [103] R. Stainforth, *Light Paths in RAT*, SNO+ DocDB **1700** (2013) [Version 6].
- [104] H. M. O’Keeffe, *Low Energy Background in the NCD Phase of the Sudbury Neutrino Observatory*. PhD thesis, University of Oxford, 2008.
- [105] V. Lozza et al., *Neutrinoless Double Beta Decay Sensitivity for SNO+ Phase I (Draft)*, SNO+ DocDB **3152** (2015) [Version 1].
- [106] M. Chen, V. Lozza, and H. M. O’Keeffe, *Expected Radioactive Backgrounds in SNO+*, SNO+ DocDB **507** (2015) [Version 27].
- [107] A. Mastbaum, *Double-Beta White Paper Sensitivity Plots*, SNO+ DocDB **2593** (2014) [Version 8].

- [108] G. O. Gann, *Dealing with Time in RAT*, SNO+ DocDB **481** (2010) [Version 2].
- [109] S. Grullon, *Light Yield and Scintillation Timing Profiles of Te-loaded Liquid Scintillator for SNO+*, . (Poster at Neutrino 2014).
- [110] M. Seddighin, *214Bi Rejection*, SNO+ DocDB **1547** (2012) [Version 4].
- [111] J. Neyman and E. S. Pearson, *On the Problem of the Most Efficient Tests of Statistical Hypotheses*, Philosophical Transactions of the Royal Society of London A **231** (1933) 289–337.
- [112] The ROOT Math Library Team, *Minuit2 Minimization Package*, 2014. seal.web.cern.ch/seal/MathLibs/Minuit2/html/index.html.
- [113] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C (Second Edition)*. Cambridge University Press, 2002.
- [114] E. Marzec, *Grid Optimiser*, SNO+ DocDB **2839** (2014) [Version 2].
- [115] M. Mottram, *Water Fitter Updates*, SNO+ DocDB **2517** (2014) [Version 1].
- [116] R. K. Wangsness, *Electromagnetic Fields (2nd Edition)*. John Wiley & Sons, 1986.
- [117] J. D. Jackson, *Classical Electrodynamics (2nd Edition)*. Wiley, 1975.