

ENABLING NEUTRINOLESS DOUBLE BETA DECAY IN
THE SNO+ EXPERIMENT
THROUGH THE DEPLOYMENT AND STUDY OF
LIQUID SCINTILLATOR

by

BENJAMIN SUM KI TAM

A thesis submitted to the
Department of Physics, Engineering Physics & Astronomy
in conformity with the requirements for
the degree of Doctor of Philosophy

Queen's University
Kingston, Ontario, Canada

August 2023

Copyright © Benjamin Sum Ki Tam, 2023

Viae ruris, ducitote me domum

Abstract

The search for neutrinoless double beta decay is among the most pressing scientific objectives at this time. If observed, this hypothesised process would not only revolutionise several fields of physics, but our fundamental understanding of the universe itself. The SNO+ experiment is a kilotonne-scale neutrino detector with the primary goal of searching for neutrinoless double beta decay in ^{130}Te , and this thesis recounts the studies and efforts that have enabled this search capability.

A new liquid scintillator was developed for the SNO+ detector to facilitate this search. To accommodate and understand this target medium, the detector hardware was upgraded, and measurements of the scintillator were undertaken to build a precise detector model. A novel chemical process was developed to stably load the liquid scintillator with ^{130}Te , and the hardware to facilitate this process was installed. Due to the rarity of the decay, trace amounts of contamination could cause insurmountable backgrounds. As such, efforts were undertaken to ensure the purity of the liquid scintillator during deployment into the detector. The first analysis of the detector when fully filled with liquid scintillator was performed, providing a characterisation of the scintillator prior to the addition of ^{130}Te .

The purity of the liquid scintillator within SNO+ was verified to be sufficient to perform a competitive search. Furthermore, all hardware and techniques are now in place to proceed with the loading of ^{130}Te into the liquid scintillator. This work has thus placed SNO+ in the position to be the next major global experiment to initiate a search for neutrinoless double beta decay.

Acknowledgements

The efforts that culminated in this thesis belong to all that shared in this journey. Not one moment has passed in the last seven years without the ceaseless support of so many. For this, I hold a debt of gratitude I will always carry.

I first express the greatest possible recognition to my advisor Mark Chen, whose leadership and guidance has shaped my fundamental understanding of science. To have been your student was a privilege surpassed only by being your friend.

Special gratitude is given to Alex Wright, who selflessly mentored me throughout my degree, and with whom my interactions never ceased to be both educational and entertaining. This is extended to Christine Kraus, Aleksandra Bialek, and Erica Caden, who took stewardship of my academic training during my constant stints at SNOLAB.

I am fortunate to have been a member of the Queen's University SNO+ research group, whose close community was the envy of all and made every day of the graduate experience memorable: Art McDonald, Ryan Martin, Peter Skensved, Phil Harvey, Lianpeng Tian, Serena Riccetto, Jooyoung Lee, Debbie Morris, Yan Liu, Ian Lam, Daniel Bartlett, Liz Fletcher, Mark Anderson, Brian Krar, Chloe Lefebvre, Anthony Allega, Jasmine Corning, Jamie Grove, Sabrina Cheng, and Rayhaneh Dehghani. I recognise the many undergraduate summer students who have contributed over the years, but especially acknowledge those I worked with closely: Sandra Cheng, Joanna Krynski, Sam Connolly, and Lily De Loë.

I thank the SNO+ Oxford group for welcoming me early and keeping me company in the waning months of this degree: Steven Biller, Armin Reichold, Jeff Tseng, Will Parker, Josie Paton, Daniel Cookman, Rafael Hunt-Stokes, Cal Hewitt, and Gulliver Milton. Special thanks to Ana Sofia Inácio for ensuring that this thesis has all of the necessary requirements.

The deepest appreciation is given to those in Sudbury whose perseverance underground directly contributed to the work within these pages: Szymon Manecki, Yi-Hsuan "Cindy" Lin, Mark Ward, Matt Depatie, Andy Stripay, Steve Maguire, Shaun Hall, Vincent Albanese, Lina Anselmo with her exceptional chemical technicians, and Paul Larochelle with his tireless team of engineers. This is extended to

the hundreds of individuals at SNOLAB and on the SNO+ experiment with whom I worked alongside both underground and around the world.

They say the real thesis are the friends you made along the way. As that would make for longer document than this, listed here are those who imparted strong impacts on my research. At Queen's: Colin Moore, Matthew Stukel, Hector Hawley Herrera, Dhruv Bisaria, Simran Nerval, Nikhil Aurora, Connor Stone, and Jean-Marie Coquillat. On SNO+: Leon Pickard, Charlie Mills, Lorna Nolan, Michal Rigan, Iwan Morton-Blake, Josh Wang, Katharine Dixon, Adil Hussain, Sammy Valder, and Juliette Deloye.

I thank the Stirling Hall staff for their assistance in all manners throughout the years, with special acknowledgements to Loanne Meldrum, Kyra Funk, Tammie Kerr, Melissa Balson, Julie McDonald, and Samantha Milliard. The same is extended to the faculty, with special acknowledgements to Tony Noble, Ken Clark, Rob Knobel, Marc Dignam, and James Fraser. I also thank the hard-working staff at the McDonald Institute and SNOLAB who always supported my endeavours towards outreach and diversity, and at the Mansion Restaurant & Bar who supported my other extracurricular activities.

I am grateful to my extended family and childhood friends with whom I found myself still always counting on as the years passed. In particular, I am thankful to Tom, Michelle, and Jeff Walton, who have welcomed me into their family.

I give special recognition to my brother Nicholas, whom I never stopped looking up to. My lifelong adoption of his interests led me to this degree, so it is to him that this thesis is dedicated. I pre-emptively appreciate my sister-in-law Tamara for finding all of the linguistic errors herein contained, and thank my nieces Clara and Cordelia for instilling within me a newfound form of joy as I worked towards this degree.

I have the deepest gratitude to my father Ambrose in whose mighty company I shall not now feel ashamed, and my mother Tania whose accomplishment as the only female engineer of her time speaks for itself. Most of all, I thank my parents for the many sacrifices they have made in the defiance of tyranny to give me the opportunity to write this thesis. 勇氣智慧也永不滅。

Finally, I thank my partner Steph Walton, whose patience was *indeed* tested with every page of this thesis. This I have assumed without cause, for her support never wavered throughout the years I pursued this degree. When we met, you assured the quality of my scintillator. Now, you assure the quality of my life.

Statement of Originality

Although the entirety of this thesis was written solely by the Author, the work within is but a part of the long legacy of the global particle physics community, and follows in the hundreds of individuals who have contributed to the SNO and SNO+ experiments.

Chapter 1 provides a literary review and overview of the field, in which the Author aims to motivate the search for $0\nu\beta\beta$ and capture the monumental efforts that led to the current State of the Art. To the best knowledge of the Author, the numerous individuals and collaborations that have contributed to this journey are acknowledged throughout.

Chapter 2 discusses the details surrounding the SNO+ detector. As such, many of the efforts described in Section 2.2 were completed prior to the involvement of the Author, though the Author participated in work related to the commissioning, upgrading, or repairing of the deck (Section 2.2.2), acrylic vessel (Section 2.2.4), cover gas systems (Section 2.2.5), and photomultiplier tubes (Section 2.2.6). The Author also held a notable role in repairs and upgrades to the DAQ hardware described in Section 2.3, serving as one of seven detector electronics experts on the SNO+ collaboration under the leadership of M. Ward.

Under the guidance of P. Skensved and J. Maneira, the Author participated in the construction and commissioning of the detector calibration hardware described in Section 2.4.1, and assisted in the deployment and analysis of the calibration sources described in Section 2.4.2 and 2.4.3. While work on the secondary SNO+ physics goals reviewed in Section 2.5 were led by other members of the collaboration and coordinated by J.R. Wilson, the Author contributed supporting analyses to several studies, most notably the water phase reactor antineutrino measurements discussed in Section 2.5.2.2. Similarly, Chapter 3 reviews the significant long-term efforts undertaken by the collaboration to understand and model the detector backgrounds; the Author contributed to several of these studies at varying capacities under the coordination and leadership of V. Lozza.

Chapter 4 discusses the work related to the development, deployment, and characterisation of the SNO+ scintillator. The development and preliminary measurements of the SNO+ scintillator described in Section 4.2.3 were completed prior to the involvement of the Author under the leadership of M. Chen. However, the Author led

or significantly participated in all aspects related to the Scintillator Fill described in Section 4.3 alongside A. Bialek, P. Larochelle, and L. Pickard. The Author also chiefly led the *ex-situ* measurements described in Section 4.4, as well as preparations for the Enhanced Scintillator Phase described in Section 4.5.

Chapter 5 recounts the efforts towards developing a technique to load a $0\nu\beta\beta$ -capable isotope into the SNO+ scintillator. Prior to the involvement of the Author, the isotope ^{130}Te was selected for use and the loading technique was in early stages of development under the leadership of S. Biller. The Author participated in studies and measurements that resulted in the conception of the diolisation process described in Section 5.1, and performed key studies that led to the understanding of the Te-loaded liquid scintillator properties discussed in Sections 5.2 and 5.3. The Author also contributed to the construction and commissioning of the Te deployment hardware described in Section 5.4.

Finally, the work presented in Chapter 6 was solely performed by the Author, but made extensive use of software and computational tools developed by the SNO+ collaboration. The analysis also greatly benefited from discussions with M. Chen, J.R. Klein, V. Lozza, S. Riccetto, and S. Naugle, and made use of the knowledge gained from supporting studies performed by countless members of the SNO+ collaboration over the years.

Table of Contents

Table of Contents	vii
List of Tables	x
List of Figures	xii
Acronyms	xv
Glossary	xxii
Chapter 1: Concerning Neutrinos	1
1.1 The Standard Model	2
1.2 Neutrinos Beyond the Standard Model	5
1.2.1 Neutrino Oscillations	7
1.2.2 The Massive Neutrino	12
1.3 The Majorana Paradigm	20
1.3.1 The Majorana Mass Mechanism	22
1.3.2 A New Hope	24
1.4 Neutrinoless Double Beta Decay	26
1.4.1 Double Beta Decay	28
1.4.2 The Neutrinoless Alternative	31
1.4.3 The Experimental Methodology	35
1.4.4 The State of the Art	42
Chapter 2: The SNO+ Experiment	47
2.1 The Road to $0\nu\beta\beta$	49
2.2 The Detector Design	51
2.2.1 SNOLAB	52
2.2.2 The Cavity	55
2.2.3 Water Purification Systems	58
2.2.4 Acrylic Vessel	62
2.2.5 Cover Gas Systems	65

2.2.6	Photomultiplier Tubes	66
2.2.7	PMT Support Structure	68
2.3	The Data Acquisition System	70
2.3.1	The Signal Flow	70
2.3.2	The Trigger System	73
2.3.3	The Data Flow	75
2.3.4	Detector Operations	77
2.4	Detector Calibrations	78
2.4.1	Source Deployment Hardware	78
2.4.2	Optical Sources	84
2.4.3	Radioactive sources	87
2.4.4	Electronics Calibrations	92
2.5	Secondary Physics Goals	94
2.5.1	Water Phase Technique	94
2.5.2	Water Phase Results	96
2.5.3	Scintillator Phase Technique	101
2.5.4	Scintillator Phase Prospects	102
Chapter 3: The Background Model		109
3.1	Expected ROI Backgrounds	110
3.1.1	Cosmogenics	112
3.1.2	Solar Neutrinos	116
3.1.3	The Double Beta Background	120
3.1.4	Uranium	121
3.1.5	Thorium	124
3.1.6	Alpha-n Reactions	126
3.1.7	External Backgrounds	127
3.2	ROI Background Model	128
3.3	The Phantom Menace	133
Chapter 4: The Deployment of Liquid Scintillator		134
4.1	Enlightening Scintillators	135
4.1.1	Scintillation Mechanisms	135
4.1.2	Liquid Scintillators	139
4.2	The SNO+ Scintillator	141
4.2.1	Linear Alkylbenzene	141
4.2.2	2,5-diphenyloxazole	144
4.2.3	The SNO+ Scintillator Cocktail	146
4.2.4	Cocktail Performance	147
4.3	The Scintillator Fill	154

4.3.1	LAB Deployment Strategy	155
4.3.2	The Scintillator Purification Plant	158
4.3.3	Scintillator Quality Assurance	170
4.3.4	End of an Era	189
4.4	Scintillator Phase Ex-Situ Measurements	191
4.4.1	Absolute Light Yield	192
4.4.2	Density	193
4.4.3	Absorption	194
4.5	The Enhanced Scintillator Phase	199
Chapter 5: The Loading of Tellurium		203
5.1	The Diolisation Process	205
5.2	Tellurium Stability	206
5.2.1	Amine Stabilisation	208
5.2.2	Thousand Standard Test	211
5.3	The Te Cocktail	213
5.3.1	Material Compatibility	214
5.3.2	Loading Technique	219
5.3.3	Light Yield Quenching	220
5.4	Te Deployment Strategy	223
Chapter 6: The Target Out Analysis		227
6.1	Analysis Techniques	228
6.1.1	Simulations	229
6.1.2	Data Processing	231
6.2	Data Selection	233
6.3	Analysis Infrastructure	237
6.3.1	Kernel Density Estimation	239
6.3.2	Binned Extended Maximum Likelihood Estimation	241
6.3.3	Fit Verification	244
6.4	The Fit	245
6.4.1	High Energy Events	247
6.4.2	2.2 MeV Excess	248
6.5	Improved Fit Results	250
6.6	The Target Out Results	253
6.7	Projected SNO+ $0\nu\beta\beta$ Sensitivity	255
Chapter 7: Concluding Remarks		257
Bibliography		259

List of Tables

1.1	Techniques and experiments that have contributed to present measurements of 3- ν oscillation parameters	14
1.2	Current best fits for 3- ν neutrino oscillation parameters	17
1.3	Current lower limits for the 3- ν neutrino masses	18
1.4	Commonly considered $0\nu\beta\beta$ candidates	39
1.5	Current results from $0\nu\beta\beta$ experiments	44
2.1	Composition of the SNO+ overburden	53
2.2	UPW Purity Standards	59
2.3	List of considered sources for future energy calibrations	93
2.4	Current best limits on mono- and di-nucleon IND decay modes	98
3.1	Cosmogenic backgrounds as a result of spallation on tellurium nuclei	114
3.2	Predicted background rate in the $0\nu\beta\beta$ ROI	130
3.3	SNO+ sensitivity parameters	132
4.1	Solvent candidates investigated for the SNO+ liquid scintillator	142
4.2	The light yield of solvent candidates	143
4.3	Fluor candidates investigated for the SNO+ liquid scintillator	145
4.4	FRET transfer efficiency between LAB and PPO at different fluor concentrations	150
4.5	Birks' constant for α , β , and p	152
4.6	The Raleigh ratio of known solvents used to deduce the Rayleigh ratio of LAB	154
4.7	Contaminants removed from the scintillator purification plant	160
4.8	QA constraints on density and turbidity of the scintillator	175
4.9	LAB parameters received from CEPESA	177
4.10	Fit parameters used to recover PPO concentration	179
4.11	Cuts used to identify $^{214}\text{BiPo}$ during the scintillator fill	183
4.12	Summary of the scintillator fill QA programme	191
4.13	Scintillator light yield fit parameters	192
4.14	Scintillator density temperature dependence fit parameters	194
4.15	Enhanced SNO+ liquid scintillator composition	201

5.1	Amines considered for stabilising the Te Cocktail	210
5.2	Thousand Sample Test statistics after 73 months	212
5.3	Initial Te cocktail composition	215
5.4	Materials investigated in the MCT	216
5.5	TST Acrylic Compatibility Sample Matrix	217
5.6	Additional Acrylic Compatibility Sample Matrix	218
6.1	Cuts used to identify $^{214}\text{BiPo}$ during the scintillator phase	234
6.2	Live times after each data selection cut	237
6.3	Background MCs generated for the target out fit	239
6.4	Cuts used to identify $^{212}\text{BiPo}$ events during the scintillator phase . . .	251
6.5	High sacrifice cuts used to identify IBD events during the scintillator phase	251
6.6	Target Out 1.5–7 MeV Fit Parameters.	254
6.7	Untagged ROI Events	255

List of Figures

1.1	Standard Model of Particle Physics	5
1.2	Illustration of possible neutrino mass hierarchies	13
1.3	Current tension in measurements of Δm_{21}^2 between solar and reactor experiments	15
1.4	Feynman diagrams of an electron antineutrino being released in a β decay, and later captured in an IBD interaction	27
1.5	Potential mass parabolas for A_{odd} and A_{even} nuclei	30
1.6	Feynman diagrams for $2\nu\beta\beta$ and $0\nu\beta\beta$	33
1.7	“The Lobster Plot”: Effective Majorana mass as a function of lightest neutrino mass	35
1.8	Predictions on $0\nu\beta\beta$ half-lives depending on the isotope and NME model	36
1.9	Illustration of the $2\nu\beta\beta$ and $0\nu\beta\beta$ energy spectra	37
1.10	The Q -value and natural abundances of $0\nu\beta\beta$ candidate isotopes. . .	40
1.11	“The Biller Plot”: The current and projected near-term sensitivities of $0\nu\beta\beta$ experiments	46
2.1	An illustration and photo of the SNO+ detector	48
2.2	Illustrated map of SNOLAB	55
2.3	The overburden provided by underground laboratories around the world	56
2.4	Block Diagram of the UPW Plant	60
2.5	The AV with hold-down and hold-up ropes	64
2.6	The schematic of the PMT, concentrator, and PSUP mounting cell .	67
2.7	Schematic of the electronics and DAQ system	71
2.8	Overview of the calibration deployment systems	79
2.9	Cross-section of a URM and umbilical	81
2.10	Schematic of the UI	83
2.11	Diagram of ELLIE subsystems	86
2.12	Discrimination of α and β hits in the Cherenkov source PMT	92
2.13	^8B solar neutrino spectrum measured during the water phase	99
2.14	Bin-by-bin fit ^8B solar neutrino fit	100
2.15	Simulated scintillator phase spectra for reactor IBD, alpha-n, and geoneutrinos	104

2.16	Sensitivity of SNO+ to Δm_{21}^2 with and without the α, n discriminator	105
2.17	Event directionality recovered from a scintillation event	107
3.1	Summary and expected share of backgrounds in the $0\nu\beta\beta$ ROI	111
3.2	Expected backgrounds in the $0\nu\beta\beta$ ROI	112
3.3	Simulated energy distributions of the decays of isotopes sourced from cosmogenic activation of Te	115
3.4	Expected solar neutrino flux	117
3.5	Radioactive decay chain of ^{238}U	122
3.6	Radioactive chain of ^{232}Th	125
3.7	Expected external events in the ROI as a function of radius	129
4.1	Illustration of the σ bonds that form the molecular structure of benzene	136
4.2	Illustration of the p orbitals that give rise to the π bonds which give scintillating properties to benzene	137
4.3	Jablonski diagram of free π electrons after excitation.	138
4.4	The emission spectrum of LAB.	146
4.5	Structural formulae of LAB and PPO	147
4.6	The light yield of LAB + PPO depending on the concentrations of PPO in LAB	148
4.7	Preliminary temperature-dependent density measurements of the scintillator cocktail	155
4.8	Block diagram summary of the LAB deployment process	156
4.9	Block diagram of the primary purification systems	161
4.10	Block diagram of the secondary distillation systems	165
4.11	Summary of PPO addition and water removal during the recirculation period	169
4.12	Photograph of the scintillator purification plant	171
4.13	Absorption spectra of LAB during scintillator plant commissioning	174
4.14	Absorption spectra for scintillator approved for use in the AV	175
4.15	CAD drawing of SCOUT	181
4.16	Correlating radon increases within the AV with various fill modes	185
4.17	Tagged ^{214}Po within a representative run during the scintillator fill	186
4.18	Tagged ^{214}Po for all runs within April 2020	187
4.19	Tagged ^{214}Po for all runs within July 2020	188
4.20	Measured light yield of the SNO+ scintillator	193
4.21	Measured density of the SNO+ scintillator at various temperatures	195
4.22	1 cm UV-Vis absorbance spectrum of the SNO+ scintillator	196
4.23	10 cm UV-Vis absorbance spectrum, attenuation, and attenuation length of the SNO+ scintillator	198
4.24	Structural formulae of BHT and bis-MSB	199

4.25	Absorption and emission lengths for the components of the enhanced liquid scintillator	200
5.1	Structural formulae of TeA and BD	204
5.2	Structural formulae and NMR spectrum of TeBD molecules	207
5.3	Te cocktail crash due to BD titration	209
5.4	Structural formula for DDA.	210
5.5	Set-up of the Thousand Sample Test	213
5.6	Te cocktail light yield depending on DDA concentration	214
5.7	Te cocktail light yield depending on loading amount and synthesis technique	221
5.8	Attenuation lengths of Te cocktail samples	222
5.9	Photographs of the TeA and TeBD synthesis plants	223
5.10	Block flow diagram for Te deployment systems	224
6.1	Tagged $^{214}\text{BiPo}$ when compared to MC	235
6.2	Run-by-run tagged $^{214}\text{BiPo}$ rate during the scintillator phase	236
6.3	Data-MC offset before and after data cleaning	238
6.4	Illustration of the KDE technique	240
6.5	Verification of the analysis infrastructure	244
6.6	First look at the fitted scintillator phase data	246
6.7	1.5–7 MeV 4-m FV data fitted after removal of tagged $^{214}\text{BiPo}$	247
6.8	The hit-ratio of tagged $^{214}\text{BiPo}$	249
6.9	1.5–7 MeV 4-m FV data fitted after implementation of the hit-ratio cut	250
6.10	1.5–7 MeV 4-m FV data fitted after data cuts and separation of the alpha-n fit	252
6.11	1.5–7 MeV 3.3-m FV data fitted after full data cuts and separation of the alpha-n fit	253

Acronyms

- $0\nu\beta\beta$ Neutrinoless Double Beta (decay).
- $2\nu\beta\beta$ 2-neutrino Double Beta (decay).
- Λ CDM Lambda Cold Dark Matter.
- ν MSM Neutrino Minimal Standard Model.
- ABS** Acrylonitrile Butadiene Styrene.
- ADC** Analogue-to-Digital Converter.
- AECL** Atomic Energy of Canada Limited.
- ALP** Axion-like Particle.
- AMB** Analogue Measurement Board.
- AMELLIE** The ELLIE Attenuation Module.
- ASTM** American Standards for Testing and Materials.
- ATLAS** A Toroidal LHC Apparatus.
- AV** Acrylic Vessel.
- BD** 1,2-butanediol.
- B-E** Bose-Einstein.
- BEST** The Baksan Experiment on Sterile Transitions.
- BHT** Butylated Hydroxytoluene.
- Bis-MSB** 1,4-bis(2-methylstyryl)benzene.
- BPO** 2-(4-biphenyl)-5-phenyloxazole.
- BR** Branching Ratio.

BSM Beyond the Standard Model.

C.L. Confidence Level (%).

CAS RN Chemical Abstracts Service Registry Number.

CC Charged Current.

CJPL China Jinping Underground Laboratory.

CKM Cabibbo-Kobayashi-Maskawa.

CMB Cosmic Microwave Background.

CMOS Complementary Metal-Oxide-Semiconductor Chip.

CMS The Compact Muon Solenoid Experiment.

CoA Certificate of Analysis.

CP Charge-Parity.

CTC Crate Trigger Card.

CUORE The Cryogenic Underground Observatory for Rare Events.

CUPID CUORE with Particle Identification.

CUTE The Cryogenic Underground Test Experiment.

DAMA/LIBRA The Dark Matter Large Sodium Iodide Bulk for Rare Processes Experiment.

DAMA/NAI The Dark Matter NaI Experiment.

DAQ Data Acquisition.

DB Daughterboard, Paddle Card.

DBD Double Beta Decay.

DCB Dodecylbenzene.

DCR Deck Clean Room.

DDA N,N-dimethyldodecylamine.

DEAP The Dark Matter Experiment using Argon Pulse-Shape Discrimination.

DIN Di-isopropylnaphthalene.

DMDA N,N-dimethyldecylamine.

DMTDA Dimethyltetradecylamine.

DOF Degrees of Freedom.

DSNB Diffuse Supernova Neutrino Background.

DT Deuterium-Tritium.

DUNE The Deep Underground Neutrino Experiment.

ECA Electronic Calibrations suite, different from ECAL.

ECAL Electronic Calibrations suite, different from ECA.

EHA 1-ethylhexylamine.

ELLIE Embedded LED/Laser Light Injection Entity.

ESI-MS Electrospray Ionisation Mass Spectrometry.

EXO The Enriched Xenon Observatory.

F-D Fermi-Dirac.

FEA Finite Element Analysis.

FEC Front-End Card, motherboard.

FRET Förster Resonance Energy Transfer.

FV Fiducial Volume.

FWHM Full Width Half Max.

GALLEX The Gallium Experiment.

GCE Grand Canonical Ensemble.

GERDA The Germanium Detector Array.

GLG4sim GenericLAND Geant4 Simulation.

GNO The Gallium Neutrino Observatory.

GT Global Trigger.

HALO The Helium and Lead Observatory.

HDA Hexadecylamine.

HV High Voltage.

IBD Inverse β Decay.

ICP-MS Inductively Coupled Plasma Mass spectrometry.

IH Inverted Hierarchy.

IND Invisible Nucleon Decay.

JUNO The Jiangmen Underground Neutrino Observatory.

K2K The High Energy Accelerator research organisation to Kamioka Experiment.

KamLAND The Kamioka Liquid scintillator Antineutrino Detector Experiment.

Katrin The Karlsruhe Tritium Neutrino experiment.

KDE Kernel Density Estimation.

KMPS Koch Modular Process Systems, LLC.

kTy Kilotonne-Years.

LAB Linear Alkylbenzene.

LAS Linear Alkylbenzene Sulfonate.

LBL Long BaseLine.

LDPE Low Density Polyethelyne.

LEGEND The Large Enriched Germanium for Neutrinoless Double Beta Decay Experiment.

LNGS Laboratori Nazionali del Gran Sasso.

MBL Medium BaseLine.

MC Monte Carlo.

MCT Material Compatibility Test.

MicroBooNE The Micro Booster Neutrino Experiment.

MiniBooNE The Mini Booster Neutrino Experiment.

MINOS The Main Injector Neutrino Oscillation Search Experiment.

MJD The Majorana Demonstrator.

MLE Maximum Likelihood Estimation.

MN 1-methylnaphthalene.

MNS Maki-Nakagawa-Sakata.

MSDS Material Safety Data Sheet.

MTC/A+ Analogue Master Trigger Card.

MTC/D Digital Master Trigger Card.

NC Neutral Current.

NEMO The Neutrino Ettore Majorana Experiment.

NEWS-G New Experiments with Spheres - Gas.

nEXO The next EXO experiment.

NEXT The Neutrino Experiment with Xenon TPC.

NH Normal Hierarchy.

NME Nuclear Matrix Element.

NMR Nuclear Magnetic Resonance.

NMSM New Minimally extended Standard Model.

NO ν A The Neutrinos at the main injector Off-axis ν_e Appearance Experiment.

NPO 2-(1-naphthyl)-5-phenyloxazole.

OA Octylamine.

OWL Outward-looking (PMT).

PC Pseudocumene (1,2,4-trimethylbenzene).

PCH Phenylcyclohexane.

PDF Probability Density Function.

PFA Perfluoroalkoxy Alkane.

PMMA Poly(methyl methacrylate), acrylic.

PMNS Pontecorvo-Maki-Nakagawa-Sakata.

PMT Photomultiplier Tube.

PMTIC PMT Interface Card.

p-p Proton-Proton (solar fusion reaction).

PandaX The Particle and Astrophysical Xenon Experiment.

PP Polypropylene.

PPO 2,5-Diphenyloxazole.

PRS Calimulse[®] Proprietary Sulfonation Process.

PSD Pulse Shape Discrimination or Pressure Safety Device.

PSUP PMT Support Structure.

PTFE Polytetrafluoroethylene.

PTI Photon Technology International.

PXE Phenyl-o-xyleneethane.

QA Quality Assurance.

QPRA Quasiparticle Random Phase Approximation.

RAT Reactor Analysis Tool.

RENO The Reactor Experiment for Neutrino Oscillation.

REPAIR Researching the Effects of the Presence and Absence of Ionizing Radiation.

RO Reverse Osmosis.

ROI Region of Interest.

SAGE The Soviet-American Gallium Experiment.

SCOUT Scintillator Counter Of Uranium and Thorium.

SEMF Semi-Empirical Mass Formula.

SENSEI The Sub-Electron-Noise-Skipper-CCD Experimental Instrument.

SK The Super-Kamiokande Experiment.

SM Standard Model.

SMELLIE The ELLIE Scattering Module.

SNEWS Supernova Early Warning System.

SNO The Sudbury Neutrino Observatory.

SNOD Discriminator Chip.

SNOI Integrator Chip.

SPE Single Photoelectron.

SSM Standard Solar Model.

STEREO The Sterile Reactor Neutrino Oscillations experiment.

SuperCDMS The Super-Cryogenic Dark Matter Search Experiment.

SURFBoard Surface Timing Board.

T2K The Tokai to Kamioka Experiment.

TDA Tetradecylamine.

TeA Telluric Acid, Orthotelluric Acid, $\text{Te}(\text{OH})_6$.

TELLIE The ELLIE Timing Module.

TPC Time Projection Chamber.

TST Thousand Standard Test.

TUBii Trigger Utility Board Mk.II.

UFO Umbilical Flasher Object.

UGBoard Underground Timing Board.

UI Universal Interface.

UPW Ultrapure Water.

URM Umbilical Retrieval Mechanism.

UTC Universal Coordinated Time.

UV-Vis Ultraviolet-visible wavelength absorption spectroscopy.

UVT Ultraviolet Transmitting.

VEV Higgs Vacuum Expectation Value.

XL3 Translator Card (Third Generation).

Glossary

β (beta) A historical term for an electron that is still commonly used.

β (beta) decay A spontaneous decay, where $n \rightarrow p + e^- + \bar{\nu}_e$.

Q value The difference in mass between the parent and daughter nuclei in a reaction, corresponding to the excess energy available for such reaction to occur.

ν MSM Model An extension of the Standard Model which treats neutrinos like all other Dirac particles that acquire their mass through the Higgs mechanism.

0ν Double Beta Decay ($0\nu\beta\beta$) A mode of double beta decay only possible if neutrinos are Majorana particles, $(A, Z) \rightarrow (A, Z + 1) + e^-$.

2ν Double Beta Decay ($2\nu\beta\beta$) The Standard Model double beta decay, $(A, Z) \rightarrow (A, Z + 1) + e^- + \bar{\nu}_e$.

2-component theory of neutrinos The (disproven) Standard Model premise of neutrinos as massless Weyl particles.

3-component neutrinos A model in which there are only three neutrino flavour and three neutrino mass eigenstates.

Acrylic Vessel The central structure of the SNO+ experiment, within which the detector medium is deployed.

Aliphatic Compounds A chemical compound that does not contain benzene structures.

Alpha-n The $^{13}\text{C}(\alpha, n)^{16}\text{O}$ interaction: $\alpha + ^{13}\text{C} \rightarrow ^{16}\text{O} + n$.

Aromatic Molecule Molecules which have a cyclic structure, such as benzene.

Attenuation Length The distance light travels within the scintillator before the original intensity is reduced by a factor of $1/e$. Also known as extinction length.

Birks' Constant An empirically determined constant which relates the light output to the energy loss per unit length for a given particle in a scintillator.

Bottoms The liquid at the bottom of the distillation column where the contaminants collect.

Breakdown A failure in the DAQ system caused by an electrical short.

Cavity The underground cavern within which the SNO+ experiment is located. Separated into the inner cavity between the AV and PSUP, and the outer cavity external to the PSUP.

Channel Threshold The adjustable voltage setting for each PMT.

Cherenkov Source A radioactive deployed source using ^8Li .

Cosmogenics A background caused by extraterrestrial sources, including radioactive isotopes created, or “cosmogenically activated,” by interactions with particles of extraterrestrial origins.

Crash When a component of a chemical mixture precipitates out of solution due to chemical instability.

Deployed Sources Calibration sources that can be manoeuvred directly into the AV (through the UI) or inner cavity (through a guide tube).

Dimer Two individual molecules held together through intermolecular forces.

Dimerisation (of TeA) When the -OH groups from two TeA molecules bind together, causing a molecule with two Te centers.

Diol A chemical compound that contains two hydroxyl (-OH) groups.

Double Beta Decay A decay mode where two simultaneous beta decays occur.

Effective Majorana Mass The mass of the exchanged Majorana electron neutrino in $0\nu\beta\beta$, $|m_{\beta\beta}| = \left| \sum_{i=1}^3 U_{ei}^2 m_i \right|$.

Event A potential particle interaction within SNO+.

Experimental Sensitivity The half-life corresponding to the maximum unobserved $0\nu\beta\beta$ signal hidden beneath backgrounds.

External Backgrounds Non-signal (i.e. not $0\nu\beta\beta$) events that originate from outside the AV but propagate into it.

External Deployment Deployment of a calibration source into the inner cavity through a guide tube.

Fiducial Volume A sub-volume of the detector that is used for analysis. Using a fiducial volume for an analysis is known as “fiducialisation”.

Grey-Disc Simplification An empirical model in which each PMT is considered a flat circle at the opening of the PMT concentrator.

Guide Tubes Cylinders that connect the deck and the inner cavity region of the detector, allowing for external deployments of calibration sources.

In-Window A technique used to discriminate between multiple particle interactions that occur within the same event.

Internal Backgrounds Non-signal (i.e. not $0\nu\beta\beta$) events that originate from within the AV.

Internal Deployment Deployment of a calibration source into the AV through the UI.

Kilotonne-Years A standard measure of detector data-taking time used in the neutrino physics community corresponding to one year of exposure to a one kilotonne target mass.

Laserball A light-diffusing sphere used as an deployed source.

Leptogenesis A hypothesised solution to the matter-antimatter asymmetry of the universe based on the Non-Conservation of Lepton Number.

Light Yield The number of emitted photons for each absorbed unit of energy.

Mass Hierarchy The mass ordering of the neutrino, currently disputed between normal hierarchy where $m_1 < m_2 < m_3$ and inverted hierarchy where $m_3 < m_1 < m_2$.

Mass Splittings (Mass Square Differences) $\Delta m_{ij}^2 = m_j^2 - m_i^2$.

Master Solution LAB loaded with a high concentration solution of PPO, BHT, or bis-MSB).

MSW Effect A correction applied to neutrino oscillations when propagating through matter (“matter oscillations”) instead of a vacuum (“vacuum oscillations”).

Neutrino Oscillations The capability for neutrinos to change flavour eigenstates as they propagate through space-time.

Nhits The number of PMT hits associated with an event.

Oligermisation (of TeA) When the -OH groups from multiple TeA molecules bind together, causing a molecule with multiple Te centers.

Partial Fill Phase An unexpected SNO+ data-taking period caused by a pause in the Scintillator Fill due to the COVID-19 pandemic, when the detector was partially filled with 365 tonnes of liquid scintillator.

Pedestal Baseline values for the charges that can be produced in the DAQ system.

Quantum Yield The fraction of excitation energy in a scintillator that is emitted as light.

Quenching The reduction in light yield due to non-radiative energy loss.

Radiogenics A background caused by natural radioactivity.

Reflux The re-condensation of distilled vapours back into the distillation column.

Run A period of time when the detector was operating under the same condition. Nominally one hour, it is the basic package of data used in SNO+ analyses.

Scintillation Counting An experimental technique in which light emitted from scintillators following their excitation by ionising radiation is detected.

Scintillator A material capable of quickly fluorescing energy as light following excitation by ionising radiation.

Scintillator Cocktail The mixture of chemicals that make up the liquid scintillator.

Scintillator Fill The major SNO+ operation to replace the UPW within the detector with liquid scintillator.

Scintillator Phase The SNO+ data-taking period when the detector was filled with liquid scintillator. Separated into the “initial” and “enhanced” sub-phases.

Secular Equilibrium The situation when the quantity of a radioactive isotope remains constant, which occurs when the production and decay rate of the isotope are equal and determined by a longer-lived parent isotope.

Seesaw Mechanism A hypothesised mechanism to provide Majorana neutrinos with mass.

SNOLAB A Canadian underground laboratory within which the SNO+ experiment is situated.

Solar Neutrino Problem The observed deficiency the solar ν_e flux when compared to solar luminosity measurements.

Spallation A process which occurs when a high energy particle bombards and shatters a nucleus, causing the ejection of follower particles.

Sterile Neutrino A hypothesised right-handed neutrino which interacts with no forces except for gravity.

Stokes’ Shift The difference between the absorption and emission peak of a material.

Survival Probability The probability that a neutrino is measured to have the flavour eigenstate it was created as.

Target Out Analysis A characterisation of the detector backgrounds in the absence of the $0\nu\beta\beta$ target.

Te Cocktail The mixture of chemicals that make up the Te-loaded liquid scintillator.

Tellurium Phase A future SNO+ data-taking period when the detector is filled with liquid scintillator doped with ^{130}Te .

Umbilical A cable which provides all of the connections required to operate a deployed source.

Water Phase The SNO+ data-taking period when the detector was filled with ultrapure water. Separated into the “initial” and “low-background” sub-phases.

Chapter 1

Concerning Neutrinos

Is it not a strange fate that we should suffer so much fear and doubt for so small a thing?

– Bor. 2:10

As far back as the Ancient Greeks, human understanding of the natural world has been underpinned by some form of what we now call the Conservation of Energy [1]. It was therefore of great alarm when early measurements of β decay appeared to violate this pillar of scientific understanding, as the resultant electron was observed to have consistently less energy than the expected Q -value of the decay.

The improbable solution to this mystery was the 1930 postulation by Wolfgang Pauli of “a terrible thing... a particle that cannot be detected” which whisked away the remaining energy [2]. Coined by E. Fermi as the “neutrino”, nearly a century of subsequent study on this particle has irrevocably changed the landscape of physics and human understanding of the universe. Even so, the many questions still unanswered about the neutrino and the field-changing implications of their possible solutions have elevated neutrino physics to among the most pressing scientific frontiers of our time.

1.1 The Standard Model

Elementary particles are the basic building blocks of the universe. Subatomic and indivisible, their properties and interactions with fundamental forces are described in the Standard Model of Particle Physics, a scientific theory that remains largely unchanged since the 1970s [3]. Formed from the union between the understanding of the strong nuclear force by C. Yang and R. Mills [4], violation of parity conservation by C-S. Wu [5], electroweak theory by S.L. Glashow [6] including the subsequent incorporation of the Higgs mechanism [7–9] by A. Salam [10] and S. Weinberg [11], and quantum chromodynamics by H. Fritzsh, H. Leutwyler, and M. Gell-Mann [12], the Standard Model has been consistently successful in passing the most stringent experimental tests yet devised.

Within the Standard Model are twelve spin-1 particles known as vector bosons, named for their obedience to Bose-Einstein (B-E) statistics. Derived from the grand canonical ensemble (GCE), B-E statistics predict the average number of particles \bar{n} that exist within energy state i through the B-E distribution,

$$\bar{n}_i = \frac{g_i}{e^{(\epsilon_i - \mu)/k_B T} + 1}, \quad (1.1)$$

where g is the degeneracy of the energy level, ϵ is the energy of the state, μ is the total chemical potential, k_B is the Boltzmann constant, and T is the temperature. As a gauge field theory, the Standard Model describes these vector bosons as the intermediary particles that mediate interactions between matter fields.

These matter fields are described within the Standard Model as twelve spin- $\frac{1}{2}$ particles. In contrast to the bosons, these spin- $\frac{1}{2}$ particles are known as fermions due

to their allegiance to Fermi-Dirac (F-D) statistics. Also derived from GCE, the F-D distribution takes the similar form of

$$\bar{n}_i = \frac{1}{e^{(\epsilon_i - \mu)/k_{\text{B}}T} + 1}, \quad (1.2)$$

with the notable lack of degeneracy implying the compliance of fermions to the Pauli exclusion principle.

As fermions are spin- $1/2$ particles, their solutions to the Standard Model Lagrangian can be found by solving the Dirac equation [13],

$$(i\gamma^\mu \partial_\mu - m)\psi = 0, \quad (1.3)$$

which itself has solutions in the form of the 4-component Dirac spinor ψ . Using the Weyl “chiral” basis, the Dirac spinor can be written in terms of two 2-component fields,

$$\psi = \begin{pmatrix} \xi \\ \eta \end{pmatrix} \quad (1.4)$$

where ξ and η respectively represent the left and right chiral fields. Within this basis, the γ matrices take the form,

$$\gamma^\mu = \begin{pmatrix} 0 & \sigma^\mu \\ \bar{\sigma}^\mu & 0 \end{pmatrix} \quad \text{and} \quad \gamma^5 = \begin{pmatrix} -I_2 & 0 \\ 0 & I_2 \end{pmatrix}, \quad (1.5)$$

where $\sigma^\mu \equiv (1, \sigma^i)$ and $\bar{\sigma}^\mu \equiv (1, -\sigma^i)$, with σ^i ($i = 1, 2, 3$) being the Pauli matrices. I_2 is a 2×2 identity matrix; as such, application of the chiral operator γ^5 to the Dirac

spinor in the Weyl basis results in,

$$\gamma^5 \begin{pmatrix} \xi \\ \eta \end{pmatrix} = \begin{pmatrix} -\xi \\ \eta \end{pmatrix}. \quad (1.6)$$

This decomposition is a special property of the Weyl basis. Furthermore, Equation 1.3 can be written in terms of Equation 1.4 as

$$i(\partial_0 + \vec{\sigma} \cdot \nabla)\eta = m\xi \quad \text{and} \quad (1.7)$$

$$i(\partial_0 - \vec{\sigma} \cdot \nabla)\xi = m\eta, \quad (1.8)$$

to explicitly show the time-variant coupling between the ξ and η subfields. Particles that can be described by this paradigm are often called “Dirac particles”.

The twelve fermions are all Dirac particles, and fall within two classes: the six quarks which are influenced by the strong force, and the six leptons which are not. Every fermion has a corresponding antiparticle is identical to the original in every aspect if not for an opposite electromagnetic charge. The masses of the fermions are generated through the Higgs coupling, a result of the spontaneous breaking of the electroweak gauge symmetries which give mass to the weak bosons [7–9].

Predicted by the Standard Model, experimental evidence of the existence of the Higgs boson by the ATLAS [14] and Compact Muon Solenoid (CMS) [15] experiments in 2012 marked the crowning achievement of the paradigm, leading many to tout the Standard Model as the most successful scientific theory of all time. However, the Standard Model, as summarised in Figure 1.1, is known to be incomplete through the obvious neglect of gravity and dark matter; the mathematical self-consistency of

the theory itself — with a Lagrangian that requires at least 19 seemingly arbitrary constants — has not yet been confirmed. Yet despite this, the Standard Model is most distinctly challenged by one of its fundamental constituents: the neutrino.




Figure 1.1: The Standard Model of Particle Physics.

1.2 Neutrinos Beyond the Standard Model

The Standard Model includes three flavours of neutrinos: the electron neutrino ν_e , the muon neutrino ν_μ , and the tau neutrino ν_τ , named for their respective associations to the charged leptonic particles e , μ , and τ . Neutrino interactions only occur by means

of the weak force, mediated by the W^\pm and Z^0 vector bosons in charged current (CC) or neutral current (NC) reactions, respectively.

C. Cowan and F. Reines first observed neutrinos^a in 1956 [16], defeating the 26-year-old prophesy of their undetectability through the inverse β decay (IBD) reaction,

$$p + \bar{\nu}_e \rightarrow n + e^+. \quad (1.9)$$

Since this first observation, all neutrinos^b have been measured to have a left-handed helicity, and all antineutrinos have been measured to have a right-handed helicity [5, 17].^c

This asymmetry is incompatible with the Higgs mechanism, which requires non-zero Yukawa couplings of both left- and right-handed fermions to the Higgs field. The simplest reconciliation is to assume that neutrinos simply do not couple to the Higgs field and therefore lack a mass generation mechanism. Such massless neutrinos are compatible with the Dirac equation as Equation 1.7 and Equation 1.8 are, in the case of $m = 0$, reduced to the Weyl equations,

$$\partial_0 \eta = -\vec{\sigma} \cdot \nabla \eta \quad \text{and} \quad (1.10)$$

$$\partial_0 \xi = \vec{\sigma} \cdot \nabla \xi. \quad (1.11)$$

First proposed in 1957 by L. Landau, T. Lee, C. Yang, and A. Salem [18–20], this “two-component theory of neutrinos” with neutrinos as left-handed Weyl particles

^aTechnically, the first neutrino observed is what we now call $\bar{\nu}_e$, an anti-neutrino.

^bOnly electron neutrinos have been measured; the rest are consistent with these measurements.

^cFor a given particle, the helicity is the projection of the particle spin on the direction of momentum while the chirality is an inherent property of the particle which determines whether it transforms in a right- or left-handed representation of the Poincaré group. For massless particles, helicity and chirality coincide.

was perfectly consistent with experimental evidence and predicted maximal parity violation in neutrinos while demonstrating the unity between neutrino chirality and helicity. The Standard Model which followed from this work was therefore fundamentally built under the assumption of massless neutrinos — an assumption vanquished through the observation of neutrino oscillations.

1.2.1 Neutrino Oscillations

The Standard Solar Model (SSM) established by J. Bahcall in the 1960s [21–24] models the interior of the Sun, and can be used to predict the rate of ν_e production from fusion reactions within the Solar core. The solar neutrino flux was first measured later that decade using the Homestake experiment through a CC reaction on ^{37}Cl ,



The flux measured by the Homestake experiment was significantly deficient compared to the predictions of the SSM [25]. For decades, improvements to the SSM [26] and Homestake experiment [27] — supplemented by complementary measurements from the **S**oviet-**A**merican **G**allium **E**xperiment (SAGE) [28], **G**allium **E**xperiment (GALLEX) [29], and Kamiokande-II experiment [30] — reaffirmed this discrepancy, which had come to be known as the “solar neutrino problem”.

1.2.1.1 Vacuum Oscillations

Even prior to the Homestake experiment, B. Pontecorvo had postulated the potential capability of neutrinos to transmute from one form to another as they propagate [31].

This phenomenon, known as “neutrino oscillations,” could account for the solar neutrino problem, as it would imply that some of the electron neutrinos released by the sun had changed form by the time they had reached Earth. This would result in a lower measured flux than expected, as contemporary experiments were exclusively or mainly sensitive to electron flavour neutrinos. Although the initial envisioning of neutrino oscillations as a lepton-sector analogy of $K^0 - \bar{K}^0$ particle-antiparticle oscillations was not consistent with measured fluxes [32], the idea was adapted by Z. Maki, N. Nakagawa and S. Sakata (MNS) to be a lepton-sector version of Cabibbo-Kobayashi-Maskawa (CKM) flavour mixing^d [33].

Within the MNS model, neutrinos would oscillate between flavours if the three neutrino flavour eigenstates $|\nu_{e,\mu,\tau}\rangle$ were merely observable superpositions of three neutrino mass eigenstates $|\nu_{1,2,3}\rangle$,

$$|\nu_\alpha\rangle = \sum_{i=1}^3 U_{\alpha i}^* |\nu_i\rangle \quad (1.13)$$

where $\alpha = e, \mu, \tau$. $U_{\alpha i}^*$ are the corresponding elements within the unitary 3×3 “lepton mixing matrix” proposed by MNS,

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & c_{13}e^{-i\delta_{\text{CP}}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{23}c_{13} \end{pmatrix}, \quad (1.14)$$

^dThough the MNS and CKM mixing theories developed in parallel, the Cabibbo mixing angle was measured first, thereby re-framing modern chronicles of this history as in the case here.

where $c_{ij} = \cos\theta_{ij}$ and $s_{ij} = \sin\theta_{ij}$. Often called the PMNS matrix,^e Equation 1.14 is paramaterised by the three mixing angles $\theta_{12,13,23}$ and includes a Dirac Charge-Parity (CP) violating phase δ_{CP} which allows for neutrino-antineutrino asymmetries. The PMNS matrix is also commonly expressed as,

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.15)$$

Considering that the mass states would be eigenstates of the Hamiltonian \mathcal{H} , solving the time-dependent Schrödinger equation with a plane wave solution of the form,

$$|x(t)\rangle = e^{-iEt}|x\rangle \quad \text{where} \quad E = \sqrt{p^2 + m^2} \quad (1.16)$$

allows for Equation 1.13 to show the flavour state evolution,

$$|\nu_\alpha(\vec{L}, t)\rangle = \sum_\beta \left(\sum_{i=1}^3 U_{\alpha i}^* e^{-i(Et - \vec{p}\vec{L})} U_{\beta i} \right) |\nu_\beta\rangle. \quad (1.17)$$

where $\beta \neq \alpha = e, \mu, \tau$. Under the approximation that $m \ll p$, solving Equation 1.17 for the probability of an α neutrino converting to a β neutrino after time t and displacement \vec{L} is therefore,

$$P(\nu_\alpha \rightarrow \nu_\beta) = |\langle \nu_\beta | \nu_\alpha \rangle(\vec{L}, t)|^2 = \sum_{i=1}^3 \sum_{j=1}^3 U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* \exp\left(-i \frac{\Delta m_{ij}^2 \vec{L}}{2E}\right). \quad (1.18)$$

^eThe addition of P is often made to also recognise B. Pontecorvo as a pioneer of neutrino oscillations.

with mass square differences or “mass splittings” of $\Delta m_{ij}^2 = m_i^2 - m_j^2$. In corollary, the survival probability, or probability that the neutrino will be measured in the original form, is

$$P_{\alpha\alpha} = 1 - P(\nu_\alpha \rightarrow \nu_\beta). \quad (1.19)$$

1.2.1.2 Matter Oscillations

While Equation 1.18 describes the probability of flavour change in neutrinos when propagating through a vacuum, the matter of oscillations are complicated when oscillating through matter. L. Wolfenstein noted in 1978 that neutrinos could undergo forward coherent elastic scattering (ES) with the electrons and neutrons within the medium it was propagating through [34]. S. Mikheev and Y. Smirnov determined that this phenomenon — now known as the Mikheev-Smirnov-Wolfenstein (MSW) effect — can be modelled by using a Hamiltonian amended with an extra term [35–37],

$$\mathcal{H}_{\text{net}} = \mathcal{H} + \mathcal{H}_{\text{Matter}} \quad \text{where} \quad \mathcal{H}_{\text{Matter}}|\nu_\alpha\rangle = V_\alpha|\nu_\alpha\rangle. \quad (1.20)$$

The total potential energy of the α -flavoured neutrino V_α is comprised of contributions from CC and NC reactions,

$$V_{\text{CC}} = \sqrt{2}G_F N_e \quad \text{and} \quad (1.21)$$

$$V_{\text{NC}} = -\frac{1}{2}\sqrt{2}G_F N_n, \quad (1.22)$$

where N_e and N_n are respectively the electron and neutron number densities of the medium, and G_F is the Fermi constant,. The total potential can then be written as,

$$V_\alpha = V_{\text{CC}}\delta_e + V_{\text{NC}} = \sqrt{2}G_F \left(N_e\delta_e - \frac{1}{2}N_n \right) \quad (1.23)$$

where the delta function δ_e reflects the exclusivity of CC reactions with ν_e particles.

Solving the time-dependent Schrödinger equation with the expanded matter-inclusive Hamiltonian amends the survival probability from the vacuum case seen in Equation 1.18 to,

$$P(\nu_\alpha \rightarrow \nu_\beta)(x) = \sin^2(2\theta_{\text{eff}})\sin^2\left(\frac{\Delta m_{\text{eff}}^2 x}{4E}\right). \quad (1.24)$$

Δm_{eff} and θ_{eff} are the effective mass difference and effective mixing angle^f between flavours α and β , where

$$\Delta m_{\text{eff}} = 2E\sqrt{\left(V_{\text{CC}} - \frac{\Delta m^2}{2E}\cos 2\theta\right)^2 + \left(\frac{\Delta m^2}{2E}\right)^2 \sin^2 2\theta} \quad \text{and} \quad (1.25)$$

$$\sin 2\theta_{\text{eff}} = \frac{\frac{\Delta m^2}{2E}\sin 2\theta}{\sqrt{\left(V_{\text{CC}} - \frac{\Delta m^2}{2E}\cos 2\theta\right)^2 + \left(\frac{\Delta m^2}{2E}\right)^2 \sin^2 2\theta}}. \quad (1.26)$$

Nearly half a century since the establishment of the Solar Neutrino Problem, PMNS-type neutrino oscillations with the MSW effect were first observed to be consistent with measurements of atmospheric neutrinos from the **Super-Kamiokande** (SK) experiment in 1998 [38], then directly measured and confirmed by the **Sudbury Neutrino Observatory** (SNO) in 2002 [39]. As seen in Equation 1.18 and Equation 1.24, oscillation-driven neutrino flavour transitions would require a non-zero mass

^fA consequence of Equation 1.24 is the expectation of a resonance angle at $\sin\theta_{\text{eff}} = \pi/4$. This implies a maximal amount of flavour mixing when the electron number density is $N_e = \frac{\sqrt{2}}{2G_F} \frac{\Delta m^2}{2E} \cos 2\theta$.

splitting and therefore non-zero neutrino masses. As such, the experimental results from SK and SNO not only resolved the solar neutrino problem, but provided the first measurement of physics Beyond the Standard Model (BSM) by unequivocally proving that neutrinos must have mass.

1.2.2 The Massive Neutrino

Despite being the first to sail beyond the horizon of BSM physics, fundamental questions posed by the newfound existence of massive neutrinos have remained largely unanswered. Indeed, nearly twenty years after uncovering evidence of neutrino masses, the masses themselves are not well measured. It is even unclear which of the three neutrinos $\nu_{1,2,3}$ with masses $m_{1,2,3}$ that superimpose to form the three flavour states $\nu_{e,\mu,\tau}$ is the lightest or heaviest. Modern measurements have not yet determined if the neutrino mass hierarchy is $m_1 < m_2 < m_3$ (“Normal Hierarchy,” NH) or $m_3 < m_1 < m_2$ (“Inverted Hierarchy,” IH)[§] since only two of the required three mass splittings are known, as shown in Figure 1.2.

1.2.2.1 The Known Parameters

The persistent unknowns surrounding the neutrino have not been due to lack of pursuit. Numerous experiments employing a variety of neutrino sources and detection techniques have measured many of the neutrino oscillation parameters within the 3- ν paradigm, as summarised in Table 1.1.

[§]The mass hierarchy is also occasionally referred to as the mass ordering, with “Normal Ordering” (NO) and “Inverted Ordering” (IO).

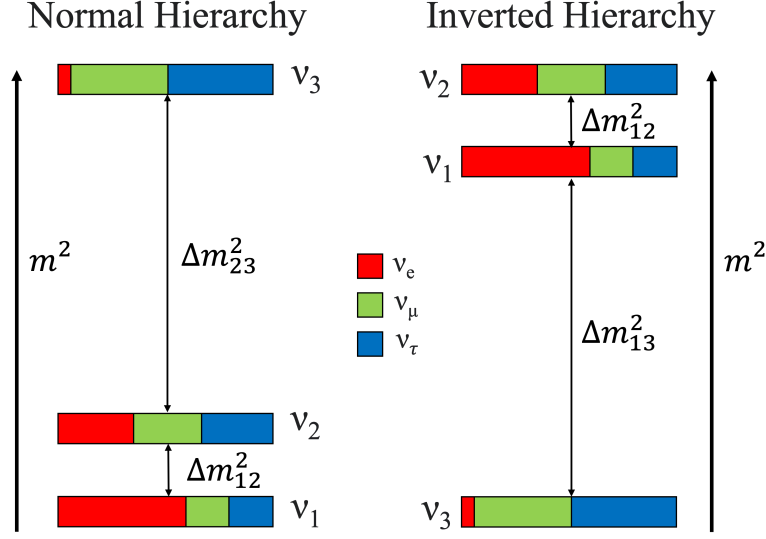


Figure 1.2: An illustration of the possible neutrino mass hierarchies. On the left, Normal Hierarchy where $m_1 < m_2 < m_3$. On the right, Inverted Hierarchy where $m_3 < m_1 < m_2$. The problem arises as only two mass splittings are known, allowing for the two possible configurations.

Experiments which have contributed to the understanding of oscillation parameters $\theta_{12,13}$ and Δm_{21}^2 by observing solar neutrinos include Homestake [25], GALLEX and the **G**allium **N**eutrino **O**bservatory (GNO) [40,41], SAGE [42], SK [43], SNO [44], and Borexino [45]. In doing so, these “solar experiments” have also constrained the sign of Δm_{21}^2 to be positive (i.e. $m_1 < m_2$) through observations of the MSW effect, as the term $V_{CC} - \frac{\Delta m^2}{2E} \cos 2\theta$ in Equation 1.26 demonstrates that the effective mixing angle differs between Δm_{ij}^2 and $\Delta m_{ji}^2 = -\Delta m_{ij}^2$.

Experiments observing atmospheric neutrinos such as SK [46] and the IceCube Neutrino Observatory [47] have contributed towards measurements of $\theta_{13,23}$ and $|\Delta m_{31,32}^2|$, along with medium baseline (MBL) reactor experiments where $L/E < 100$ km/MeV such as the Daya Bay Reactor Neutrino Experiment [48], the **R**eactor **E**xperiment

Table 1.1: The techniques and experiments that have contributed to present measurements of 3- ν oscillation parameters.

Technique	Experiments	Parameters
Solar Neutrinos	Homestake, SAGE, GALLEX/GNO, SK, SNO, Borexino	$\theta_{12,13}, \Delta m_{21}^2$
Atmospheric Neutrinos	SK, IceCube	$\theta_{13,23}, \Delta m_{31,32}^2 , \delta_{\text{CP}}$
MBL Reactors	Daya Bay, RENO, Double Chooz	$\theta_{13}, \Delta m_{31,32}^2 $
LBL Reactors	KamLAND	$\Delta m_{21}^2, \theta_{12,13}$
LBL Accelerators	K2K, MINOS, T2K, NO ν A	$\theta_{13,23}, \Delta m_{31,32}^2 , \delta_{\text{CP}}$

for Neutrino Oscillation (RENO) [49], and Double Chooz [50]. The **K**amioka **L**iquid scintillator **A**ntineutrino **D**etector (KamLAND), a long baseline (LBL) reactor experiment, has also quantified Δm_{21}^2 and $\theta_{12,13}$ through measurements of the $\bar{\nu}_e$ survival probability [51]. However, their measurements of Δm_{21}^2 are presently in tension with those of SK and SNO at 1.5σ , as can be seen in Figure 1.3 [43].

The LBL accelerator experiments such as the **K**EK^h **t**o **K**amioka (K2K) experiment [52], the **M**ain **I**njector **N**eutrino **O**scillation **S**earch (MINOS) and MINOS+ [53], the **T**okai **t**o **K**amioka (T2K) experiment [54], and the **N**eutrinos at the **M**ain **I**njector **O**ff-axis ν_e **A**ppearance (NO ν A) experiments [55] have also measured $\theta_{13,23}$ and $|\Delta m_{31,32}^2|$. However, the sign of $|\Delta m_{31,32}^2|$ has yet to be determined — achieving this would definitively determine the neutrino mass hierarchy, as $m_2 > m_1$

^hKEK stands for the High Energy Accelerator research organisation when transliterated from Japanese.




Figure 1.3: The current 1.5σ tension in the measurement of Δm^2_{21} between solar neutrino and reactor neutrino experiments. In green are the 1σ , 2σ , and 3σ confidence levels for the combined measurement from SK and SNO, in blue the measurements from KamLAND, and in red the combined result. Figure adapted from [43].

and $|\Delta m^2_{31,32}| > |\Delta m^2_{21}|$ is already measured to a sufficient precision. The mass hierarchy could nevertheless still be resolved through more precise measurements of $|\Delta m^2_{21}|$ and $|\Delta m^2_{32}|$, as the oscillation signal slightly differs between NH and IH; statistical analyses of current measurements already prefer NH at a $2.7 - 3.5\sigma$ level [56]. Improving precision of oscillation parameters towards this end is the main objective

of upcoming MBL reactor experiments such as the **Jiangmen Underground Neutrino Observatory** (JUNO) [57].

Meanwhile, the atmospheric and LBL accelerators are also capable of constraining δ_{CP} through measurements of $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation channels. However, present best constraints from T2K and NO ν A are in tension at the 90% confidence level (C.L.) within the NH paradigm [55, 58]. Further investigation of these oscillation channels are the main objective of next-generation LBL accelerator experiments, including Hyper-Kamiokande [59] and the **Deep Underground Neutrino Experiment** (DUNE) [60].

The extraction of the oscillation parameters from experimental measurements is challenging, as doing so requires numerous *a priori* assumptions which differ depending on the mass hierarchy and choice of phenomenological models. A popular global analysis of oscillation data has been performed by the NuFIT Collaboration utilising the New Minimally extended Standard Model (NMSM), a paradigm that extends the Standard Model in the simplest way to achieve consistency with dark energy, non-baryonic dark matter, neutrino masses, baryon asymmetry, and cosmic inflation [61]. The most recent oscillation parameters using NuFit 5.2 (November 2022) which include recent atmospheric neutrino data from SK can be found in Table 1.2 [62]. Analyses with differences in Standard Model extensions, datasets, and solar metallicities have also been performed to recover slightly different oscillation parameters [63].

Table 1.2: Current best fits for 3- ν neutrino oscillation parameters utilising NuFit 5.2, a global analysis utilising the NMSM model by the NuFit collaboration. Data retrieved from [62].

Parameter	NH Fit $\pm 1\sigma$	IH Fit $\pm 1\sigma$
$\sin^2(\theta_{12})$	$0.303^{+0.012}_{-0.011}$	$0.303^{+0.012}_{-0.011}$
$\sin^2(\theta_{23})$	$0.572^{+0.018}_{-0.023}$	$0.578^{+0.016}_{-0.021}$
$\sin^2(\theta_{13})$	$0.02203^{+0.00056}_{-0.00059}$	$0.02219^{+0.00060}_{-0.00057}$
$\delta_{\text{CP}} [^\circ]$	197^{+42}_{-25}	286^{+27}_{-32}
$\Delta m_{21}^2 [\times 10^{-5} \text{ eV}^2]$	$7.41^{+0.21}_{-0.20}$	$7.41^{+0.21}_{-0.20}$
$\Delta m_{31}^2 [\times 10^{-3} \text{ eV}^2]$	$2.511^{+0.028}_{-0.027}$	$-2.424^{+0.032}_{-0.025}$
$\Delta m_{32}^2 [\times 10^{-3} \text{ eV}^2]$	$2.437^{+0.028}_{-0.027}$	$-2.498^{+0.032}_{-0.025}$

1.2.2.2 The Neutrino Mass Scale

Although the absolute mass scale of the neutrino is not known, fixing the lightest neutrino massⁱ to zero provides a lower constraint on neutrino masses. Using oscillation parameters from the NuFit 5.2 global analysis in Table 1.2, the calculated lower bound neutrino masses are summarised on Table 1.3. The **K**arlsruhe **T**ritium **N**eutrino (KATRIN) experiment has also constrained the mass of $\bar{\nu}_e < 0.11 \text{ eV}$ (90% C.L.) through measurements of the β decay endpoint in tritium [64].

The upper end of the mass scale is currently constrained by model-dependent cosmological parameters. Using measurements of cosmic microwave background anisotropies (CMB) and baryon acoustic oscillation by the Planck observatory, the sum of the neutrino masses has been constrained to $\sum_{i=1}^3 \nu_i < 0.12 \text{ eV}$ (95% C.L.) if assuming the Λ CDM model [65]. Over $\mathcal{O}(10^6)$ times lighter than the other elementary particles,

ⁱThe lightest neutrino being m_1 for NH and m_3 for IH.

Table 1.3: Current lower limits for 3- ν neutrino masses parameters calculated utilising NuFit 5.2, a global analysis utilising the NMSM model by the NuFit collaboration.

Neutrino	NH (meV)	IH (meV)
m_1	0	$49.23^{+0.25}_{-0.32}$
m_2	$8.61^{+0.12}_{-0.12}$	$49.98^{+0.25}_{-0.32}$
m_3	$50.11^{+0.28}_{-0.27}$	0

the relative smallness of the neutrino masses — although not phenomenologically excluded — are among many of the motivations for improving our understanding of neutrino mass generation mechanisms.

1.2.2.3 The Dirac Mass Mechanism

As the masslessness of the neutrino was the cornerstone of Weyl neutrino paradigm of Equation 1.10 and Equation 1.11, the two-component theory of neutrinos is no longer consistent with experimental evidence. Its defeat resurrects the question of how neutrino masses are generated. The simplest mass generation mechanism would be to assume that neutrinos, like any other fermions, act like Dirac particles and acquire their masses through the Higgs mechanism. Within this “Neutrino Minimal Standard Model” (ν MMSM)^j [66], the neutrino component of the diagonalised Higgs-lepton Yukawa Lagrangian takes the form

$$\mathcal{L}_{HL} \supset \mathcal{L}_D = - \left[\sum_{i=1}^3 m_D \bar{\nu}_{i,L} \nu_{i,R} \right] + \text{H.c} \quad (1.27)$$

^j ν MMSM differs from the NMSM used in the NuFit analysis of Table 1.2 in that the former only includes neutrino masses.

where $i = 1, 2, 3$. The neutrino would then have a Dirac mass of m_D ,

$$m_{Di} = \frac{y_i^\nu v}{\sqrt{2}}, \quad (1.28)$$

proportional to the neutrino Yukawa coupling y_i^ν and Higgs Vacuum Expectation Value (VEV) of $v \sim 246$ GeV. As can be seen in Equation 1.27, the existence of both left- and right-handed neutrinos are required for this mass generation mechanism.

Although the ν MSM paradigm is a relatively clean extension, the model does not address the fundamental problem that initially led to the assumption of neutrino masslessness: how only left-handed neutrinos and right-handed antineutrinos have ever been measured. To be compatible with both these experimental observations and ν MSM, right-handed neutrino fields would not only need to exist, but have electric charge $q = 0$, weak isospin $I_3 = 0$, and weak hypercharge $Y_W = 0$ to not interact with any of the Standard Model fundamental forces. Indeed, attempts at measuring the effects of these “sterile” neutrinos by the Planck [67, 68], IceCube [69], **Mini Booster Neutrino Experiment** (MiniBooNE) and MicroBooNE [70], the **Baksan Experiment on Sterile Transitions** (BEST) [71], and the **Sterile Reactor Neutrino Oscillations** (STEREO) [72] experiments have found generally disfavoured but conflicting evidence for their existence. Furthermore, the value of y_i^ν required to be consistent with the observable neutrinos alone would be $\mathcal{O}(10^{-12})$ GeV, over $\mathcal{O}(10^{10})$ times smaller than the Yukawa couplings to other elementary particles. This anomalous smallness and the lack of evidence for sterile neutrinos have prompted many to wonder if there instead exists a more elegant solution to the mysteries brought forth by the massive neutrino.

1.3 The Majorana Paradigm

The most sophisticated alternative for the Dirac mass mechanism is based on a reformulation of the Dirac spinor seen in Equation 1.4. First postulated by E. Majorana [73], the Dirac spinor can be reduced from four to two degrees of freedom (DOF) by constructing a spinor where the two-component right chiral field η can be rewritten as functions of the two-component left chiral field ξ :

$$\psi_M = \begin{pmatrix} \xi \\ i\sigma_2\xi^* \end{pmatrix}. \quad (1.29)$$

This “Majorana spinor” features the Pauli matrix,

$$\sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad (1.30)$$

and satisfies the Dirac equation of Equation 1.3. Indeed, by substituting the Majorana spinor of Equation 1.29 into the Dirac equation, the two-component Majorana equation of motion can be recovered as,

$$\sigma^\mu \delta_\mu \xi + m\sigma_2\xi^* = 0. \quad (1.31)$$

Particles described within this Majorana paradigm are known as “Majorana particles”.

The most interesting feature of the Majorana spinor is elucidated when a charge conjugation operator is applied,

$$\psi_M^c \equiv i\gamma^2\psi_M^* = \psi_M. \quad (1.32)$$

As seen in Equation 1.32, the charge conjugation operator has no effect on the Majorana spinor. As this is the operation used to transform a particle into its antiparticle, Majorana particles are therefore their own antiparticles.^k

Thus far, there is no evidence for the existence of any Majorana particles, and the conditions for their actuality are stringent. For particles of gauge charge q interacting with gauge field A_μ , the equations of motion for a particle described with ψ and its antiparticle ψ^c are, respectively,

$$(i\gamma^\mu(\delta_\mu + iqA_\mu) - m)\psi = 0 \quad \text{and} \quad (1.33)$$

$$(i\gamma^\mu(\delta_\mu - iqA_\mu) - m)\psi^c = 0. \quad (1.34)$$

In the Majorana case of $\psi = \psi^c$, these equations are contradictory unless all gauge charges are zero. Furthermore, the particles must be spin- $1/2$ fermions, such that they can be described by the four-component Dirac spinor, which the Majorana spinor is reduced from. Finally, the particle must have mass, as there would otherwise be no distinction between the Dirac and Majorana paradigms. As they are now known to possess mass, the neutrino now stands as the only known particle to fulfil these criteria.

^kIn contrast, particles utilising the four-component Dirac spinor have distinct particles and antiparticles, since applying a charge conjugation operator to Equation 1.3 results in $\psi^c \equiv i\gamma^2\psi^* \neq \psi$.

1.3.1 The Majorana Mass Mechanism

If neutrinos are indeed Majorana particles, the Majorana particle-antiparticle equivalence of Equation 1.32 can be applied to Equation 1.27 to rewrite the neutrino mass term as,

$$\mathcal{L}_M = - \left[\sum_{i=1}^3 \left(\frac{1}{2} m_L \bar{\nu}_{i,L} (\nu_{i,L})^C + \frac{1}{2} m_R (\bar{\nu}_{i,R})^C \nu_{i,R} \right) \right] + \text{H.c.}, \quad (1.35)$$

where m_L and m_R are the left- and right-handed Majorana masses. These Majorana masses can be generated by coupling to only a single handedness, and would provide an alternative to the Dirac mass mechanism of Section 1.2.2.3.

By combining the Dirac mass term of Equation 1.27 and the Majorana mass term of Equation 1.35, a general Lorentz-invariant mass term for a Majorana neutrino can be constructed. Simplified to a single neutrino generation for demonstrational simplicity,¹ this combined mass term is,

$$\mathcal{L}_{D+M} = \left[-\frac{1}{2} m_L \bar{\nu}_L (\nu_L)^C - m_D \bar{\nu}_L \nu_R - \frac{1}{2} m_R (\bar{\nu}_R)^C \nu_R \right] + \text{H.c.} \quad (1.36)$$

This combined mass term can be re-written in Matrix form as,

$$\mathcal{L}_{D+M} = -\frac{1}{2} \begin{pmatrix} (\bar{\nu}_L)^C & \bar{\nu}_R \end{pmatrix} \begin{pmatrix} m_L & m_D \\ m_D & m_R \end{pmatrix} \begin{pmatrix} \nu_L \\ (\nu_R)^C \end{pmatrix} + \text{H.c.} \quad (1.37)$$

Under a model known as the ‘‘Type-I Seesaw Mechanism’’ [74], the following assumptions are then made:

- There are no left-handed Majorana mass terms (i.e. $m_L = 0$).

¹Generalising this to three neutrino generations requires a diagonalisation of 3×3 mass matrices and three dimensional neutrino vectors, introducing a technical complication while leaving the principle unchanged.

- The Dirac mass m_D is generated through the usual Higgs mechanism and is therefore of $\mathcal{O}(10)$ GeV a similar scale to the charged leptons and quarks.
- The right-handed Dirac mass will be larger than that of the electroweak $\mathcal{O}(100)$ GeV scale, and therefore $m_R \gg m_D$.

Within these assumptions, the central 2×2 mass matrix^m in Equation 1.37 can be simplified and diagonalised to rewrite the Lagrangian as,

$$\mathcal{L}_{D+M} = -\frac{1}{2} \begin{pmatrix} (\bar{\nu}_L)^C & \bar{\nu}_R \end{pmatrix} \begin{pmatrix} m_D^2/m_R & 0 \\ 0 & m_R \end{pmatrix} \begin{pmatrix} \nu_L \\ (\nu_R)^C \end{pmatrix} + \text{H.c.} \quad (1.38)$$

Within the Type-I Seesaw Mechanism, the usual left-handed neutrinos would have a mass of,

$$m_{\nu,L} = m_D^2/m_R. \quad (1.39)$$

While right-handed neutrinos are still sterile, they are unconstrained by Standard Model gauge transformations and can have masses well above normal Standard Model scales. Indeed, potential masses even at Grand Unification scales of $m_{\nu,R} \sim 10^{15}$ GeV are consistent with experimentally observed mass limits of $m_{\nu,L} < 0.01$ MeV [65], providing strong theoretical interest in Majorana neutrinos and the Type-I Seesaw Mechanism.

^mThis would be a 6×6 matrix if considering all three generations of neutrinos.

1.3.2 A New Hope

Beyond the conciliatory energy scales absent in the Dirac mass mechanism - compelling enough by itself - there are further diverse reasons to hope that neutrinos are Majorana particles:

- Within the Lagrangian, the Majorana masses would be generated by the Weinberg operator - the only dimension-five operator consistent with the gauge symmetries of the Standard Model - and is thus already expected to be the first sign of this form of high-scale physics [75].
- The reduction of the neutrino field from four to two DOF is consistent with experimental observations, as there have only ever been two distinctly measurable neutrinos.
- The existence of Majorana neutrinos motivates searches for other Majorana particles, including dark matter candidates [76].
- Majorana neutrinos may account for the existence of the universe itself through “Leptogenesis,” a hypothesised solution to the ongoing mystery of matter-antimatter asymmetry in the universe. If neutrinos are indeed Majorana particles, heavy right-handed neutrinos frozen out in the early universe would decay to leptons and anti-leptons. These decays may be asymmetric due to CP violation and give rise to an excess of leptons, ultimately resulting in the observed baryon asymmetry of the universe [77].

Of course, the introduction of Majorana neutrinos would not come without theoretical complications. The Majorana mass term seen in Equation 1.35 is not gauge

invariant, and hence the phase of the Majorana fields cannot be freely redefined. Consequently, the PMNS lepton mixing matrix of Equation 1.14 must be extended to account for this through the addition of two CP-violating Majorana phase terms, ϕ_1 and ϕ_2 [78]:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\phi_1} & 0 & 0 \\ 0 & e^{i\phi_2} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.40)$$

Despite this complication, the Majorana extension cancels out upon being squared. Thus, it does not interfere with the neutrino survival probability in Equation 1.18, thereby remaining consistent with experimental measurements of neutrino oscillations [79]. The Majorana mass mechanism would also change the lepton number by ± 2 .ⁿ However, as lepton number is an accidental symmetry and not intrinsic to the Standard Model, this violation of the Conservation of Lepton Number is not strictly disallowed. Nevertheless, new BSM physics would be required as the mass terms cannot be generated in a gauge invariant way, instead requiring a Higgs Triplet with $I_3 = 1$ and $Y_W = -2$ [80].

Evidence of the Majorana nature of neutrinos would be a landmark discovery, with far-reaching implications across multiple fields of physics. Thus, with field-changing impacts apparent and known theoretical challenges already at bay, the possibility of neutrinos as Majorana particles remains the most exciting solution to the various questions still posed by the neutrino. Yet for all the revolutionary consequences, the most cogent reason to desire the Majorana paradigm is more practical: there exists a

ⁿLepton number conservation is an accidental symmetry wherein the total sum of leptons (each with $L = 1$) and anti-leptons (each with $L = -1$) must be the same before and after each interaction.

means to probe for it. This technique, known as “Neutrinoless Double Beta Decay,” is not only within our current limits of understanding, but may prove the simplest way to measure the other ill-charted properties of the neutrino from the absolute neutrino masses to the neutrino mass hierarchy.

1.4 Neutrinoless Double Beta Decay

With the board now set by theoretical motivations for Majorana particles, the pieces must be moved through experimentation. At first glance, such an investigation need not be challenging, as antineutrino sources from β decays in the form of,

$$n \rightarrow p + e^{-} + \bar{\nu}_e \quad (1.41)$$

are routinely measured by IBD reactions in the form of,

$$p + \bar{\nu}_e \rightarrow n + e^{+}. \quad (1.42)$$

As seen in the Figure 1.4(a), this Standard Model interaction expectedly conserves lepton number ($\Delta L = 0$). Within the Majorana paradigm where the antineutrino and neutrino are the same particle, it is conceivable that the same antineutrino source could induce the β^{-} IBD reaction instead:

$$n + \nu_e \rightarrow p + e^{-}, \quad (1.43)$$

This is seen in Figure 1.4(b), including the signature Majorana lepton number violation of $\Delta L = 2$.

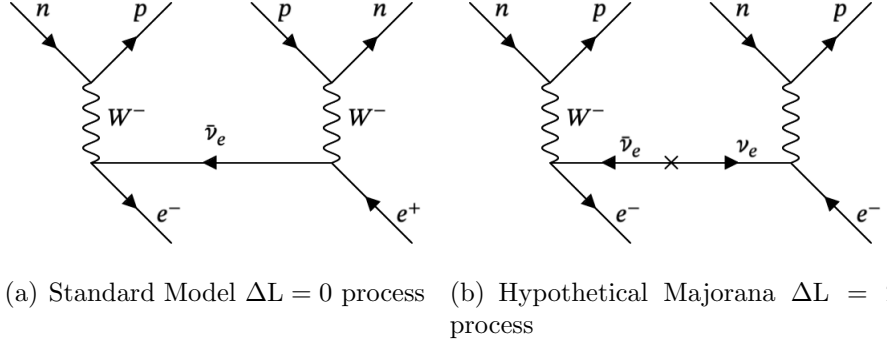


Figure 1.4: Feynman diagrams of an electron antineutrino being released in a β decay, and later captured in an IBD interaction. Time moves from left to right, and the Standard Model process shown in (a) is well studied. If neutrinos are Majorana particles, there is an additional process shown in (b) that is possible.

Within this hypothetical Majorana interaction, the “anti”-neutrino produced in the source must be right-chiral (u_R) and right-helicity (u_\downarrow) while the neutrino which induces the IBD reaction must be left-chiral (u_L) while maintaining helicity (u_\downarrow). Since u_L decomposes into helicity states $u_{\uparrow\downarrow}$ in the form [81],

$$u_L = \frac{1}{2}(1 + \kappa)u_\uparrow + \frac{1}{2}(1 - \kappa)u_\downarrow, \quad (1.44)$$

where neutrinos of $m < 0.1$ eV have,

$$\kappa = \frac{p}{E + m} \sim \left(1 - \frac{m}{E}\right) \lesssim 1, \quad (1.45)$$

the $u_{L,\downarrow}$ interaction is suppressed by a factor of (m^2/E^2) .^o This helicity suppression poses a problem in typical neutrino experiments, as the energies involved in both reactor and accelerator experiments exceed $\mathcal{O}(10^6)$ eV. Thus, the Standard Model

^oFrom Fermi’s Second Golden Rule, the rate is proportional to the square of perturbation.

$\Delta L = 0$ process — dominant by a factor of at least $\mathcal{O}(10^{14})$ — would be impenetrable with modern detection capabilities. A nuclear process that inherently suppresses the $\Delta L = 0$ mode would be required for such an observation to be feasible.

1.4.1 Double Beta Decay

First developed by E. Fermi [82], the β^\pm decay processes in which

$$p \rightarrow n + e^+ + \nu_e \quad \text{and} \quad (1.46)$$

$$n \rightarrow p + e^- + \bar{\nu}_e \quad (1.47)$$

are now well understood to respectively occur in nuclei in the form of,

$$(A, Z) \rightarrow (A, Z - 1) + e^+ + \nu_e \quad \text{and} \quad (1.48)$$

$$(A, Z) \rightarrow (A, Z + 1) + e^- + \bar{\nu}_e. \quad (1.49)$$

This process spontaneously occurs in nuclei with an isobar of lower mass.

The mass difference between the parent and daughter nuclei is not simply the mass difference in raw constituent masses of the nuclei involved, as the binding energy E_B of the nucleus must be accounted for:

$$m_{tot} = Zm_p + Nm_n - E_B. \quad (1.50)$$

Within the Semi-Empirical Mass Formula (SEMF) first proposed by C. Weizäcker, this binding energy is approximated to [83, 84],

$$E_B = a_V A - a_S A^{2/3} - a_C \frac{Z(Z-1)}{A^{1/3}} - a_A \frac{(N-Z)^2}{A} + \delta(A, Z), \quad (1.51)$$

and is affected by empirically determined terms related to the strong force through volume a_V and surface area a_S , electrostatic repulsion a_C , and the asymmetry of the nucleus a_A . The binding energy is also affected by the effect of spin coupling, proportional to a pairing term,

$$\delta(A, Z) = \begin{cases} +a_P A^{k_p} & \text{for even } Z, N \text{ (even } A) \\ 0 & \text{if odd } A \\ -a_P A^{k_p} & \text{for odd } Z, N \text{ (even } A) \end{cases} \quad (1.52)$$

where k_p is an experimentally constrained energy-binding term of $-3/4 < k_p < -1/2$.

Plotting the SEMF mass for a set of isobars yields a parabolic function; spontaneous decays are driven from higher mass nuclei to the stable nucleus of minimum mass. As can be seen in Fig.1.5 (left), the decay chain is straightforward in the case of odd A when the pairing term has no effect.

In the case of even A , there are a few dozen naturally occurring isotopes where decays to the neighbouring isobar of $Z \pm 1$ are energetically disallowed due to the pairing term.^P Consequently, as seen in Figure 1.5, a situation could arise where a single β decay is forbidden, but two simultaneous β decays are not. First noted by

^PThere also exists double β decay isotopes for which the intermediate state transition is energetically allowed but spin forbidden.




Figure 1.5: Mass parabolas for (left) a nucleus of A_{odd} and (right) a nucleus of A_{even} . Nuclei will undergo spontaneous decays until they reach the isobar with the lowest mass. In A_{even} nuclei, beta decay channels may be energetically disallowed, allowing for the ascendancy of (red) a double β decay process.

M. Goeppert-Mayer [85], both processes in this “two-neutrino double beta” ($2\nu\beta\beta$) decay would happen simultaneously within the nucleus of the heavier isobar as,

$$(A, Z + 2) \rightarrow (A, Z) + 2e^+ + 2\nu_e \quad \text{or} \quad (1.53)$$

$$(A, Z - 2) \rightarrow (A, Z) + 2e^- + 2\bar{\nu}_e. \quad (1.54)$$

The half-life of the decay can be found from utilising Fermi’s Second Golden Rule to be,

$$(\mathcal{T}_{1/2}^{2\nu})^{-1} = G_{2\nu} |\mathcal{M}_{2\nu}|^2. \quad (1.55)$$

The phase space factor G is determined by integrating over the six-dimensional space of momentum-position for each of the four leptons emitted in the decay. The Nuclear

Matrix Element (NME) \mathcal{M} is the probability amplitude for the decay containing only Gamow-Teller components. Determining the NME requires a model of how every particle in the nucleus transitions from initial to final states - non-trivial for $2\nu\beta\beta$ decays with multiple intermediate states and high-nucleon isotopes involved. There are numerous different approaches used to estimate $\mathcal{M}_{2\nu}$, with the most popular being the Quasiparticle Random Phase Approximation (QRPA) [86].

Depending on the isotope, experimentally-consistent NME formalisms place the half-life of $2\nu\beta\beta$ decays in the range of $\mathcal{T}_{1/2}^{2\nu} \sim \mathcal{O}(10^{19-22})$ years [87], making the process among the rarest in the Standard Model.⁹ Indeed, $2\nu\beta\beta$ decays have been experimentally observed in only twelve of the 35 naturally occurring isotopes capable of the process [88]. Ongoing precision measurements of $2\nu\beta\beta$ decays are excellent experimental examinations of NME models. However, experiments capable of measuring $2\nu\beta\beta$ decays typically have a more Majorana objective in mind.

1.4.2 The Neutrinoless Alternative

$2\nu\beta\beta$ decays in the form of Equation 1.53 and Equation 1.54 are Standard Model processes with a Feynman diagram shown in Figure 1.6(a). However, if neutrinos are Majorana particles, then there is an additional channel where the $u_{R\downarrow}$ “anti”-neutrino created by the first β decay immediately induces the second β decay as a $u_{L\downarrow}$ neutrino, as seen in Fig.1.6(b). This “neutrinoless double beta” ($0\nu\beta\beta$) decay channel, only

⁹Noting that from measurements made by Planck, the universe itself is only $13.787 \pm 0.020 \times 10^9$ years old according to the Λ CDM model [65].

accessible through the exchange of a Majorana neutrino, would take the form

$$(A, Z - 2) \rightarrow (A, Z) + 2e^- \quad \text{or} \quad (1.56)$$

$$(A, Z + 2) \rightarrow (A, Z) + 2e^+. \quad (1.57)$$

Along with the signature Majorana lepton number violation of $\Delta L = 2$, $0\nu\beta\beta$ is reminiscent of the β decay \rightarrow IBD process shown in Figure 1.4(b). However, the $\Delta L = 0$ Standard Model case shown in Figure 1.4(a) is energetically forbidden in double β decays and therefore naturally suppressed.^r Thus, unlike the Majorana search modes using reactors or accelerators, the search for $0\nu\beta\beta$ is free of the otherwise overbearing $\Delta L = 0$ signal.

That is not to say that the search for $0\nu\beta\beta$ is wholly unobstructed. As in the case for $2\nu\beta\beta$, the half-life of the $0\nu\beta\beta$ process can be determined through Fermi's Second Golden Rule to be,

$$(\mathcal{T}_{1/2}^{0\nu})^{-1} = G_{0\nu} |\mathcal{M}_{0\nu}|^2 \langle m_{\beta\beta} \rangle^2. \quad (1.58)$$

As in the 2ν case, determining $\mathcal{M}_{0\nu}$ is an active area of research, with several models and techniques still under investigation.^s Even though the $0\nu\beta\beta$ decay channel has superior phase space to the $2\nu\beta\beta$ (having only two final leptons instead of four), the $[\nu_L W^- e^-]$ vertex in Fig.1.6(b) still suffers from the same helicity suppression factor

^rThe $\Delta L = 0$ process in the form of what is seen in Figure 1.4(a) would require the neutrino from $(A, Z - 2) \rightarrow (A, Z + 1) + e^- + \bar{\nu}_e$ to initiate the second stage, $(A, Z - 1) + \bar{\nu}_e \rightarrow (A, Z + 2) + e^+$. Even if there was sufficient energy to account for the excess leptons, the condition of double beta decay (as demonstrated in Fig.1.5) is predicated on how the initial process is energetically forbidden to begin with.

^sExperimental validation of $2\nu\beta\beta$ NME formalisms can provide evidence towards $0\nu\beta\beta$ models, though uncertainties will exist as the $0\nu\beta\beta$ NME will contain Fermi components in addition to Gamow-Teller components.




Figure 1.6: Feynman diagrams for (a) the Standard Model $2\nu\beta\beta$ decay, and (b) the hypothetical $0\nu\beta\beta$ decay which could occur if neutrinos are Majorana particles.

of $(m/E)^2$, with the relevant energies still in similar scales of $m \sim \mathcal{O}(0.1)$ eV and $E \sim \mathcal{O}(10^6)$ eV.

This helicity suppression is absorbed into the “effective Majorana mass” $m_{\beta\beta}$, an inherent, model-independent component of the NME that can be extracted [89]. $m_{\beta\beta}$ is the mass of the exchanged Majorana electron neutrino as a superposition of mass

eigenstates, stated as [90],

$$|m_{\beta\beta}| = \left| \sum_{i=1}^3 U_{ei}^2 m_i \right|. \quad (1.59)$$

With U_{ei} being the entries of the PMNS matrix with Majorana extensions shown in Equation 1.40 projected onto the electron flavour state, $m_{\beta\beta}$ contains the oscillation parameters,

$$|m_{\beta\beta}| = |c_{12}^2 c_{13}^2 m_1 + s_{12}^2 c_{13}^2 m_2 e^{i\phi_1} + s_{13}^2 m_3 e^{i\phi_2}|. \quad (1.60)$$

This link between the $0\nu\beta\beta$ decay and PMNS matrix grants the search for neutrinoless double beta decay a powerful capability to probe fundamental neutrino properties beyond confirming the existence of the Majorana paradigm. As demonstrated in Fig.1.7, recovering the Majorana mass would explore both the absolute neutrino mass scales and the neutrino hierarchy problem.

Utilising existing limits on the oscillation parameters from Table 1.2, the expected $0\nu\beta\beta$ half-life can be calculated from Equation 1.58 to be $\mathcal{T}_{1/2}^{0\nu} \sim \mathcal{O}(10^{25-28})$ years depending on the isotope chosen. As can be seen in Figure 1.8, the half-life expectation show variability when different approaches are used to estimate the NME, even within the same isotope. As with any nuclear decay, the half-life would be the observable experimental quantity in any measurement of $0\nu\beta\beta$. The challenge thus falls to constructing an experiment capable of detecting such rare decays.




Figure 1.7: The effective Majorana mass as a function of lightest neutrino mass, including the Majorana phases ϕ_1 and ϕ_2 . The various shades correspond to 1 – 5 σ uncertainties on the oscillation parameters used. Colloquially known as the “lobster plot”. Also shown on the plot are current sensitivity limits, as well as sensitivity projections for future, larger experiments with tonne or hundred-tonne quantities of $0\nu\beta\beta$ isotope. Figure adapted from [90].

1.4.3 The Experimental Methodology

In principle, the fundamental approach to performing a $0\nu\beta\beta$ decay search is straightforward. While observing $\beta\beta$ decays, the electrons^t that are emitted from $2\nu\beta\beta$ will yield a continuous energy spectrum, as the neutrinos will carry away a portion of the

^tThese would be positrons in the case of $\beta^+\beta^+$ decays, of which none have ever been observed as they are suppressed further by $\epsilon\epsilon$ and $\epsilon\beta$ decay modes. Only the $\beta^-\beta^-$ case will be henceforth considered, though the distinction is trivial to the methodology.




Figure 1.8: Predictions of the half-life of $0\nu\beta\beta$ depending on the isotope and NME model. (R)QRPA, LSSM, EDF, PHFB, and IBM-2(I) are some potential NME models still being considered. Figure taken from [81].

available energy. It follows that in the case of $0\nu\beta\beta$, the energy of the electrons will exactly sum to the Q -value of the decay, as demonstrated in Figure 1.9. Thus, a statistically significant measurement of this signature peak at the Q -value of the decay would confirm the Majorana nature of the neutrino. In fortuitous poetry, the missing Q -value β decay peak which began the neutrino saga may yet be its deliverance.

The execution of such a search is unfortunately not so elementary. The rarity of the decay significantly complicates attempts at detection, as even trace amounts of backgrounds in the energy region of interest (ROI) around the Q -value of the decay can engulf the $0\nu\beta\beta$ signal. Therefore, the chief experimental challenge behind every $0\nu\beta\beta$ search is the reduction, removal, or discrimination against these background




Figure 1.9: An illustration of the predicted summed energy spectrum of two electrons in both $2\nu\beta\beta$ and $0\nu\beta\beta$ decay channels. While the neutrinos will purloin energy in the 2ν channel, they are (definitionally) not present in the 0ν channel, compelling the two electrons to carry away the full Q -value of the decay.

sources through shielding, purification, or analytical techniques. While most background sources originating from cosmic rays or environmental radioactivity can be mitigated in this way, there is a contribution from the $0\nu\beta\beta$ isotope itself that cannot: the $2\nu\beta\beta$ decay.

Considering their difference in half-lives, the $0\nu\beta\beta$ decay of a particular isotope may still be $> \mathcal{O}(10^3)$ times less frequent than the corresponding $2\nu\beta\beta$ decay from the same isotope. Furthermore, limitations on detector energy resolution will smear out the distinctive $0\nu\beta\beta$ decay peak at the Q -value. Being so much greater in rate, the endpoint of the $2\nu\beta\beta$ spectrum will overlap with the smeared $0\nu\beta\beta$ distribution, thereby creating an irremovable background.

Luckily, as seen in Equation 1.55 and Equation 1.58, the half-life of the decays scale inversely with the phase space, which in turn scale as,

$$\begin{aligned} G^{2\nu} &\propto Q^{-11} \\ G^{0\nu} &\propto Q^{-5}. \end{aligned} \tag{1.61}$$

Since this could be interpreted as,

$$\frac{T_{1/2}^{2\nu}}{T_{1/2}^{0\nu}} \propto Q^6, \tag{1.62}$$

isotopes with higher Q -values will have improved $0\nu\beta\beta$ search viability through mitigation of the $2\nu\beta\beta$ background. Coupled with the fact that lower energies are more beset by ambient background signals from both radiogenic sources, only the few isotopes with $Q_{\beta\beta} > 2 \text{ MeV}$ are typically considered viable for a $0\nu\beta\beta$ search. The other major factor limiting isotope viability is natural abundance of the isotope, as isotopes with low natural abundance must undergo potentially cost-prohibitive enrichment.^u The Q -values and natural abundances of commonly considered $0\nu\beta\beta$ candidate isotopes are summarised in Table 1.4.

As seen in Figure 1.10, there is unfortunately no single isotope that provides an obvious advantage in both Q -value and natural abundance. $0\nu\beta\beta$ experiments must therefore turn to other factors to improve their search capabilities. From the

^uAt the time of writing, this is increasingly exacerbated as geopolitical factors become progressively alarming.

Table 1.4: The eleven most commonly considered $0\nu\beta\beta$ candidates, with Q -values from [91] and natural abundance from [92].

Decay Transition	Q -value (MeV)	Natural Abundance (%)
${}^{48}_{20}\text{Ca} \rightarrow {}^{48}_{22}\text{Ti}$	4.261	0.187
${}^{76}_{32}\text{Ge} \rightarrow {}^{76}_{34}\text{Se}$	2.039	7.8
${}^{82}_{34}\text{Se} \rightarrow {}^{82}_{36}\text{Kr}$	2.998	9.2
${}^{96}_{40}\text{Zr} \rightarrow {}^{96}_{42}\text{Mo}$	3.348	2.8
${}^{100}_{42}\text{Mo} \rightarrow {}^{100}_{44}\text{Ru}$	3.035	9.6
${}^{116}_{48}\text{Cd} \rightarrow {}^{116}_{50}\text{Sn}$	2.813	7.5
${}^{130}_{52}\text{Te} \rightarrow {}^{130}_{54}\text{Xe}$	2.527	34.08
${}^{136}_{54}\text{Xe} \rightarrow {}^{136}_{56}\text{Ba}$	2.459	8.9
${}^{150}_{60}\text{Nd} \rightarrow {}^{150}_{62}\text{Sm}$	3.371	5.6

radioactive decay law, the $0\nu\beta\beta$ half-life can be evaluated as,

$$\mathcal{T}_{1/2}^{0\nu} = \ln(2)\epsilon \frac{N}{N_{\beta\beta}} T \quad (1.63)$$

where ϵ is the detection efficiency, T is the livetime of the experiment, and $N_{\beta\beta}$ are the number of detected $0\nu\beta\beta$ events. N corresponds to the number of $0\nu\beta\beta$ capable nuclei within the experiment and can be calculated with,

$$N = \left(\frac{xaN_A M}{A} \right), \quad (1.64)$$

where x is the number of $\beta\beta$ atoms within each molecule, a is the natural abundance, N_A is the Avogadro number, M is the total mass of the source within the experiment,




Figure 1.10: The Q -value and natural abundances of $0\nu\beta\beta$ candidate isotopes. Those most favourable for $0\nu\beta\beta$ searches are filled in. Also shown are the most energetic particles emitted from the decay chains of ^{238}U and ^{232}Th , the dominant radiogenic sources in the $\mathcal{O}(1)$ MeV energy region as explained in Chapter 3.

and A is the molecular weight of the source. If a $0\nu\beta\beta$ signal is observed, the half-life can then be calculated from Equation 1.63 with an uncertainty $\sigma_{\mathcal{T}}$ proportional to the Poisson fluctuation on the number of decays,

$$\frac{\sigma_{\mathcal{T}}}{\mathcal{T}_{1/2}^{0\nu}} = \frac{\sqrt{N_{\beta\beta}}}{N_{\beta\beta}}. \quad (1.65)$$

In the case where there is no apparent signal, $N_{\beta\beta}$ can be rewritten as,

$$n_B = n_\sigma \sqrt{b\Delta_E T M}, \quad (1.66)$$

corresponding to the maximum number of $0\nu\beta\beta$ decay counts hidden beneath the background at an n_σ - σ C.L.^v Here, b is the rate of background counts (quoted as events/year/keV/kg) and Δ_E is the width of the ROI. Combining Equation 1.63, Equation 1.65, and Equation 1.66 results in,

$$\mathcal{S}_{1/2}^{0\nu} = \ln(2)\epsilon \frac{1}{n_\sigma} \frac{x a N_A}{A} \sqrt{\frac{M T}{b \Delta_E}}. \quad (1.67)$$

Known as the “experimental sensitivity”, $\mathcal{S}_{1/2}^{0\nu}$ is the half-life corresponding to the maximum possible $0\nu\beta\beta$ signal hidden beneath the background and is the standard factor of merit used to compare $0\nu\beta\beta$ experiments.

From Equation 1.67, it is apparent that the sensitivity of a $0\nu\beta\beta$ experiment for a given isotope — and therefore their detection potential — are dependent on these detector design choices:

- the detector efficiency ϵ and energy resolution Δ_E , primarily affected by the detection technology used,
- the mass of the detector target M ,
- the fraction of M that is the $0\nu\beta\beta$ isotope x ,
- the level of enrichment used to increase a above natural levels, and

^v n - σ being the number n_σ of standard deviations σ .

- the mitigation of backgrounds b through shielding, purification, or analysis techniques.

There has yet to be found a practical^w detector design that excels at all qualities, as current techniques that improve one aspect may negatively affect others. However, this does allow for a deep breadth of variety in detector designs intended for $0\nu\beta\beta$ decay searches.

1.4.4 The State of the Art

Although there has been hitherto no credible signal of $0\nu\beta\beta$ decay,^x the possibility of discovery around any corner — and the field-changing implications such a measurement would bring — has excited fierce global competition. Due to the rarity of the decay, all $0\nu\beta\beta$ experiments struggle to find ways to mitigate detector backgrounds. Regardless, the experiments themselves differ widely in detector design, each attempting to press advantages to increase their sensitivities in the race to find $0\nu\beta\beta$ decay.

Some experiments have detector designs based on exceptional energy resolution, such as the **C**ryogenic **U**nderground **O**bservatory for **R**are **E**vents (CUORE), which uses TeO₂ crystals cooled to 7 mK as low temperature bolometers [97].^y However, the sensitivity of CUORE is limited by background rejection capabilities due to unreliable particle identification [99].

^wPractical meaning affordable in most cases.

^xSome members of the Heidelberg-Moscow experiment have claimed a discovery since 2001 [93]. Their result is widely disputed by the community [94] (including other members of the Heidelberg-Moscow collaboration [95]), and outright excluded to high confidence by subsequent experiments [96].

^yThe CUORE collaboration claim that their experiment is the coolest m³ in the universe — believable, since the CMB itself is 2.726 K [98].

While CUORE utilises the high natural abundance of ^{130}Te , other experiments that concentrate on high energy resolution use extreme enrichment and purification. These experiments include the **M**ajorana **D**emonstrator (MJD), which uses p-type point contact detectors with 88%-enriched ^{76}Ge [100], and the **G**ermanium **D**etector **A**rray (GERDA) experiment, which used 86%-enriched ^{76}Ge crystal diodes within a liquid argon cryostat [101]. Both exhibit superior energy resolution to CUORE; GERDA was further boosted by powerful background rejection techniques such as Pulse Shape Discrimination (PSD) and an active argon medium to achieve exceptionally low background levels in their ROI. However, such experiments are limited by mass due to the cost of their enrichment.

Even more expensive than the usage of ^{76}Ge is that of ^{136}Xe , as the global supply and production of the isotope is limited. However, the easier enrichment of ^{136}Xe allows the isotope to stay competitive. Experiments using ^{136}Xe include the **E**nriched **X**enon **O**bservatory (EXO-200), which deployed 200 kg liquid Xe enriched to 80.6% within a Time Projection Chamber (TPC) [102,103], and the **K**amLAND **Z**ero **N**eutrino **D**ouble-Beta Decay Experiment (KamLAND-Zen 400), which deployed a balloon containing 381 kg (later increased to 745 kg in KamLAND-Zen 800) of 91%-enriched Xe within the KamLAND experiment [104].

Beyond these three main isotopes of ^{130}Te , ^{76}Ge , and ^{136}Xe , $0\nu\beta\beta$ decay has been searched for in a multitude of other isotopes by other experiments. The successor to CUORE, the **C**UORE **U**ppgrade with **P**article **I**dentification (CUPID) experiment, has already run small-scale demonstrators with ^{82}Se (CUPID-0) [105] and ^{100}Mo (CUPID-Mo) [106]. The **N**eutrino **E**ttore **M**ajorana **O**bservatory (NEMO-3) has

performed a variety of searches utilising external $0\nu\beta\beta$ sources [107], publishing limits on ^{82}Se , ^{100}Mo , ^{96}Zr , ^{150}Nd , ^{116}Cd , and ^{48}Ca [108–113]. Observations of ^{116}Cd by the Aurora experiment [114] and ^{48}Ca by ELEGANT VI [115] have also been performed. A summary of these experiments, and their latest published sensitivities, is shown in Table 1.5.

Table 1.5: Current results from former and ongoing $0\nu\beta\beta$ experiments. Highlighted are experiments presently holding the best published sensitivity in their respective isotope.

Experiment	Isotope	Sensitivity ($\times 10^{25}$ Years), 90% C.L.
CUORE	^{130}Te	≥ 2.2 [99]
MJD	^{76}Ge	≥ 8.3 [116]
GERDA	^{76}Ge	≥ 18 [96]
EXO-200	^{136}Xe	≥ 3.5 [117]
KamLAND-Zen	^{136}Xe	≥ 23 [104]
CUPID-0	^{82}Se	≥ 0.35 [118]
NEMO-3 (Se)	^{82}Se	≥ 0.025 [108]
CUPID-Mo	^{100}Mo	≥ 0.18 [119]
NEMO-3 (Mo)	^{100}Mo	≥ 0.11 [109]
NEMO-3 (Zr)	^{96}Zr	≥ 0.00092 [110]
NEMO-3 (Nd)	^{150}Nd	≥ 0.0020 [111]
NEMO-3 (Cd)	^{116}Cd	≥ 0.010 [112]
Aurora	^{116}Cd	≥ 0.022 [114]
NEMO-3 (Ca)	^{48}Ca	≥ 0.0020 [113]
ELEGANT-VI	^{48}Ca	≥ 0.058 [115]

Of the present and former experiments, none exceed the one-tonne isotope milestone which the scientific community has defined as The Next Generation of $0\nu\beta\beta$ experiments. As the hunt for $0\nu\beta\beta$ intensifies, most collaborations have begun preparing for this next leg of the race. The CUPID experiment, successor to CUORE, will switch from ^{130}Te to ^{100}Mo enriched crystals, and has already published results using their demonstrators CUPID-0 and CUPID-Mo [105, 106]. The GERDA and Majorana Demonstrator experiments will unite to form the **L**arge **E**nriched **G**ermanium for **N**eutrinoless **D**ouble **B**eta **D**ecay (LEGEND) experiment, aiming to use 1000 kg enriched ^{76}Ge -enriched detectors; their 200 kg demonstrator, LEGEND-200, is currently under development [120]. The EXO-200 experiment is aiming to scale up their liquid Xe TPC technology in the **n**ext **EXO** (nEXO) experiment [121]. nEXO will compete in their isotopic arena against the upcoming gaseous Xe TPC experiments, the **N**eutrino **E**xperiment with **X**enon **T**PC (NEXT) Experiment [122] and the **P**article **a**nd **A**strophysical **X**enon (PandaX-III) Experiment [123]. The NEMO-3 experiment will be succeeded by the SuperNEMO experiment [124], for which a demonstrator is now operational [125]. The projected sensitivities of major current and future planned $0\nu\beta\beta$ experiments can be seen in Figure 1.11.

Despite the tight competition and prodigious engagement in the global scientific community, all of the tonne-scale experiments are still in early stages of development and years from coming to fruition.

All save one.

Not a successor to an earlier $0\nu\beta\beta$ experiment but to SNO itself, the SNO+ experiment aims to observe $0\nu\beta\beta$ through an initial 3.9 tonne deployment of ^{130}Te . The detector hardware is prepared, tested, and understood. A new experimental




Figure 1.11: The current bounds (solid), projected near-term sensitivities (dashed), and Next Generation sensitivities (dotted) of $0\nu\beta\beta$ experiments, depending on the NME formalism. Also shown are regions compatible with IH and NH. Colloquially known as the “Biller Plot,” this figure is adapted from [126] by the original author.

medium capable of carrying the isotope has been developed, deployed, and studied. The SNO+ experiment now stands as the only new challenger poised to make a measurement of $0\nu\beta\beta$ decay in the immediate future, thereby potentially resolving many of the pressing mysteries launched by its predecessor.

Chapter 2

The SNO+ Experiment

Very bright was that sword when it was made whole again; the light of the sun shone redly in it, and the light of the moon shone cold, and its edge was hard and keen.

– Tol. 2:3

The SNO+ experiment is a particle detector built with the primary objective of searching for $0\nu\beta\beta$. The heart of the detector is a central 6-m radius spherical Acrylic Vessel (AV) filled with 904L of fluid detection medium. Particle interactions (“events”) within the AV are observed by 9362 photomultiplier tubes (PMTs) mounted on a 8.89-m geodesic steel PMT Support structure (PSUP) that encapsulates the AV. The entire structure is located in a cavity excavated 2092 ± 6 m underground; the entire cavity external to the AV, including the space between the AV and PSUP, is filled with ultrapure water (UPW). An illustration and photo of the detector is seen in Figure 2.1.

As the successor to SNO, the SNO+ experiment has inherited many of the infrastructure and design principles that enabled the highly successful SNO physics




Figure 2.1: (Left) An illustration of the SNO+ detector, showing the central AV surrounded by the PSUP which houses a majority of the PMTs. (Right) A photo of the SNO+ detector taken from one of the Nikon D5000 cameras equipped with a fisheye lens permanently installed on the PSUP. This photo is taken from near the bottom of the detector looking upwards.

programme — and ultimately led to the 2015 Nobel Prize in Physics. The SNO+ AV, PSUP, and PMTs are the very same used in SNO and remain installed at their original location in what is now the SNOLAB facility near Sudbury, Canada.^a

The central SNO+ $0\nu\beta\beta$ strategy is to gain advantage through deployment of large masses of the $0\nu\beta\beta$ isotope. To this end, it can be seen from Figure 1.10 that the obvious isotope to choose was ^{130}Te . With a natural abundance of 34.08%, ^{130}Te could be cheaply procured in large quantities relative to other isotopes as no enrichment is required. An initial deployment of 3.9 tonnes of $^{\text{nat}}\text{Te}$ is planned, corresponding to

^aThe organisation that is now known as SNOLAB and the assorted laboratory space beyond the original SNO cavity only came to be in the waning days of the SNO experiment as an expansion to support the SNO+ and the Dark Matter Experiment using Argon Pulse-shape Discrimination (DEAP-1) experiments.

1300 kg ^{130}Te . This makes SNO+ not just the only new tonne-scale $0\nu\beta\beta$ experiment on the horizon, but the only currently planned experiment using ^{130}Te .

2.1 The Road to $0\nu\beta\beta$

Being primarily concerned with the ^8B solar neutrinos with an energy spectrum extending well into $\mathcal{O}(10)$ MeV, the SNO detector had an effective electron kinetic energy threshold of $E > 3.5$ MeV [127], while ^{130}Te has an endpoint energy of [128],

$$Q_{\beta\beta}^{^{130}\text{Te}} = 2527.518 \pm 0.013 \text{ keV}. \quad (2.1)$$

Therefore, major detector improvements were necessary to access these lower energies, the most significant of which being the deployment of liquid scintillator as the detector medium.

With existing liquid scintillators insufficiently compatible with the acrylic that makes up the AV, a new scintillator was developed for the SNO+ experiment based on Linear Alkylbenzene (LAB) doped with 2.2 g/L 2,5-diphenyloxazole (PPO).^b Additional enhancements would later be made to further increase both the stability and light yield of the liquid scintillator. As the deployment of ^{130}Te within liquid scintillator had never before been attempted, a new loading method had to be designed. Further discussion on the SNO+ scintillator and Te-loading methodology is found in Chapters 4 and 5, respectively.

With the precious SNO detector medium of heavy water ($^2\text{H}_2\text{O}$, D_2O) returned to Atomic Energy of Canada Limited (AECL) and the detector emptied, the road

^bSolute concentrations written as g/L henceforth mean $g_{\text{solute}}/L_{\text{LAB}}$ unless otherwise specified.

to $0\nu\beta\beta$ in SNO+ was laid out in three titular phases based on the deployed target medium.

Water phase First, the SNO+ detector was refilled with 904 tonnes of ultrapure water (H_2O , UPW) as the target medium. Acting as a water Cherenkov detector, the primary goals of the water phase were to measure external background sources, tune the detector response, achieve stable running of all detector systems, and test the calibration systems.

Scintillator phase The water was then replaced with 792 tonnes of liquid scintillator, thereby increasing the light yield to access lower energy processes at a higher resolution. The scintillator phase began on April 29, 2022 and is ongoing at the time of this writing. The primary objective of this period is to quantify the backgrounds of the deployed scintillator prior to the insertion of the $0\nu\beta\beta$ isotope.

Tellurium phase Finally, the scintillator would be loaded with ^{130}Te to enable the search for $0\nu\beta\beta$. The initial loading amount of 0.5%^c $^{\text{nat}}\text{Te}$ will be deployed, with higher ^{130}Te concentrations of up to 3% already planned.

Each of the three phases would also be facilitated by a commissioning period: the water fill, scintillator fill, and tellurium fill.

The water phase commenced following the water fill on 4 May 2017, and lasted until 16 July 2019, when the scintillator fill period began. The scintillator fill was paused between 21 March 2020 and 26 October 2020, as the onset of the COVID-19 pandemic

^cIsotope concentrations written as % henceforth mean % of total mass.

led to both a major restriction in access to the SNOLAB underground facility and a shortage of trained personnel. Prior to this interruption, SNO+ was filled with 365 tonnes of LAB doped with 0.6 g/L PPO. This unexpected period afforded an adventitious opportunity to understand and calibrate the detector while filled with scintillator. The scintillator phase commenced following the end of the scintillator fill on 29 April 2022, and is currently ongoing.

2.2 The Detector Design

As demonstrated in Equation 1.67, the sensitivity of SNO+ scales as,

$$\mathcal{S}_{1/2}^{0\nu} \propto \sqrt{\frac{M}{b}}. \quad (2.2)$$

While the SNO+ strategy is to maximise M , the chief experimental challenge — as with all $0\nu\beta\beta$ experiments — is to suppress the number of background events b .

Minimising radiogenic and cosmogenic backgrounds sourced from outside of the AV are achieved through shielding. Meanwhile, radiogenic backgrounds from sources within the AV are primarily mitigated through ensuring the high purity of the target medium. In both cases, utilising and maintaining high purity materials used in all components of the experiment are of the utmost importance. A discussion of the expected SNO+ backgrounds can be found in Chapter 3, with the purification of the scintillator and tellurium discussed in Chapter 4 and Chapter 5, respectively.

Mitigating backgrounds in this way was also the primary consideration behind the design of the SNO experiment, which required low neutron rates to measure NC neutrino interactions. Although the SNO+ $0\nu\beta\beta$ search has purity requirements that

are $\mathcal{O}(2)$ more stringent, many of the same fundamental SNO design principles for mitigating backgrounds could still be used [129–131].

2.2.1 SNOLAB

Among the most severe potential backgrounds are those of cosmic origin. As the Earth careens through the universe at a rate of 190–350 km/s,^d the atmosphere is bombarded by high energy cosmic rays. Consisting of $\sim 95\%$ protons, $\sim 5\%$ He nuclei, and small amounts of heavier nuclei, these particles interact with atmospheric nuclei and produce a multitude of high energy decay products primarily consisting of π^\pm and K^\pm mesons [132]. The daughter products then decay to,

$$\begin{aligned}
 \pi^+ &\rightarrow \mu^+ + \nu_\mu \\
 \pi^- &\rightarrow \mu^- + \bar{\nu}_\mu \\
 K^+ &\rightarrow \mu^+ + \nu \\
 K^- &\rightarrow \mu^- + \bar{\nu} \\
 K^\pm &\rightarrow \pi^\pm + \pi^0.
 \end{aligned}
 \tag{2.3}$$

The decay products themselves interact with the atmosphere, thereby causing a cascading chain reaction called an “air shower”. These air shower particles — mostly γ , e^\pm , p , n , and μ^\pm — litter the surface of the Earth with a vertical intensity of $I = (0.564 \pm 0.072) \text{ cm}^{-2} \text{ sr}^{-1} \text{ min}^{-1}$ at sea level [133], creating an overwhelming background across the entire energy spectrum. In particular, high energy muons strike the surface of the Earth at a vertical intensity of $I_\mu \sim 0.01 \text{ cm}^{-2} \text{ sr}^{-1} \text{ min}^{-1}$ [134], corresponding to an average muon flux of $\Phi_\mu = \mathcal{O}(10^{-2}) \text{ cm}^{-2} \text{ s}^{-1}$. Upon being struck

^dThis depends on the annually periodic alignment between our solar and galactic orbits.

by one of these high energy muons, a nucleus shatters and ejects numerous follower particles in a process known as spallation.

As there is no man-made shielding thick or absorbent enough to withstand such a barrage of cosmic ray particles, the crust of the Earth is used to break the tide. Thus, the SNO+ experiment is located 2092 ± 6 m below the local surface at $46^{\circ}28'30''$ N ($56^{\circ}33'$ magnetic north) $81^{\circ}12'04''$ W. The region itself is 309 m above sea level, with variations of up to ± 50 m in the local 5 km radius. Historically an area of large-scale mining operations, the granitic crust is well studied and comprised nearly entirely of norite rock, with $\langle Z^2/A \rangle = 6.01$ and a chemical composition shown in Table 2.1 [135]. The rock density directly above the detector has been measured to vary between 2.8 – 2.9 g/cm³ depending on the depth, with an inferred average rock density of 2.83 ± 0.05 g/cm³.

Table 2.1: The composition of the norite rock that comprises the SNO+ overburden.

Element	Fractional Composition (%)
Oxygen	45
Silicon	26
Aluminum	9
Iron	4

After taking air and water filled cavities into account, this overburden provides a shielding of 5890 ± 94 meters water equivalent (m.w.e.) at the center of the AV [135]. The cosmogenic muon interaction rate within SNO+ has been measured to be $3.31_{\pm 0.09(\text{sys})}^{\pm 0.01(\text{stat})} \times 10^{-10} \mu\text{cm}^{-2}\text{s}^{-1}$ [129], implying that the overburden provides a muon rate reduction of 5×10^7 .

The underground location within which SNO+ (and all associated underground infrastructure) occupies is a part of the SNOLAB underground laboratory, itself within the “6800(ft)-level” of the active Vale Limited (formerly INCO) Creighton Mine. An expansion of the original SNO facility, the SNOLAB laboratory was originally constructed at the end of the SNO experiment. At the time of this writing, SNOLAB houses a number of other experiments including the **H**elium and **L**ead **O**bservatory (HALO), DEAP-3600 (successor to DEAP-1), PICO-40L, the **S**ub-**E**lectron-**N**oise **S**kipper-**C**CD **E**xperimental **I**nstrument (SENSEI), **N**ew **E**xperiments with **S**pheres-**G**as (NEWS-G), the **C**ryogenic **U**nderground **T**est **E**xperiment (CUTE), and **R**esearching the **E**ffects of the **P**resence and **A**bsence of **I**onizing **R**adiation (REPAIR). The **S**uper **C**ryogenic **D**ark **M**atter **S**earch (SuperCDMS) and PICO-500 (successor to PICO-40L) are also currently under construction within the laboratory. The laboratory is maintained by the SNOLAB organisation, which is also responsible for providing engineering, infrastructure, and administrative support to the experiments located within. The layout of the SNOLAB facility can be seen in Figure 2.2.

The appeal of SNOLAB is not just the depth and existing SNO/SNO+ infrastructure, but the clean lab conditions within the facilities. All areas of the underground laboratory operate at a class-2000 level,^e further mitigating the possibility of contamination. In areas closer to the SNO+ infrastructure, the cleanliness is increased to class-200. Although the China Jinping Underground Laboratory (CJPL) and the experiments within are deeper underground than SNOLAB, CJPL is located within a mountain and thus admits a greater muon background from lower incident angles

^eA class-x cleanroom is one where there are fewer than x particles per cubic foot of diameter 0.5 μm or larger. In comparison, a typical outdoor space is equivalent to class 1 million.

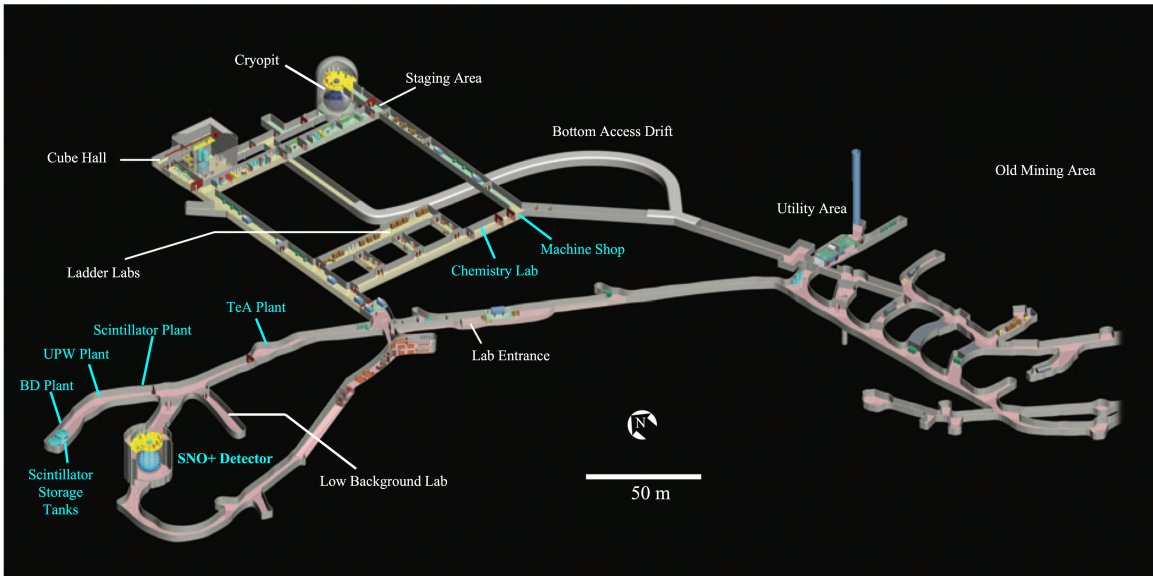


Figure 2.2: An illustrated map of the SNOLAB underground facility. The areas shaded in pink were used in the original SNO experiment, with newer lab expansions shaded in green. Labeled in cyan are major lab infrastructure components used in direct support of the SNO+ experiment. Figure adapted from [136].

compared to the flatter SNOLAB overburden. As such, SNO+ is the neutrino detector with the greatest effective overburden in the world [137]. The depth and muon rate of SNOLAB compared to other underground laboratories used for neutrino physics can be seen in Figure 2.3.

2.2.2 The Cavity

One advantage of CJPL over the SNOLAB site is the level of radioactivity in the surrounding rock, as the marble bedrock which surrounds the Chinese laboratory contains lower concentrations of ^{232}Th and ^{238}U [138]; these backgrounds are further discussed in Chapter 3. To shield from the background sources in the rock overburden




Figure 2.3: The vertical overburden depth and muon flux of underground laboratories that host neutrino experiments. The red squares are laboratories within mines and therefore have flat overburdens, while the blue triangles are laboratories in mountains that therefore do not. The blue fit (with uncertainty the shaded region) implies that mountain laboratories have a muon flux of 4.0 ± 1.9 times higher than mine laboratories. The fit and figure is from [137].

itself, the SNO+ detector is surrounded by 7000 tonnes of UPW shielding. The AV and PSUP are thus submerged in this pool of UPW within a barrel shaped cavity of 22-m diameter and 34-m height blasted out of the underground rock. Excavated between March 1990 and May 1993, the SNO+ (then SNO) cavity was the deepest cavity of its size at the time and required the removal of over 60,000 tonnes of rock.

Secured and stabilised with 800 cable bolts cemented 30 feet into the surrounding rock, the cavity was first lined with low-background concrete with a measured trace radioactivity of $0.1 \text{ Bq } ^{238}\text{U kg}^{-1}$ and $0.45 \text{ mBq } ^{232}\text{Th kg}^{-1}$ [139] in order to suppress neutrons and γ -rays from the norite. The concrete was further lined with layers of Urylon polyurethane coating to mitigate radon permeability into the cavity from the surrounding rock. Tests during SNO determined that two layers — a 0.25-inch base of Urylon HH453 Mineguard bottom coat under a 0.05-inch Urylon 201-25 top coat — formed a sufficient radon seal [140]; well over two layers were applied to minimise radon ingress.

While filled with UPW, the cavity is only accessible through a reinforced concrete deck structure situated just above the water line. The deck is primarily held up by two stainless steel trusses anchored into the cavity ceiling through four corbels with seismic dampers. Load cells provide active and live capabilities to monitor the strain on each corbel. Secondary trusses provide further stability to the deck, which is used as a storage and access point for the SNO+ calibration and data acquisition (DAQ) systems.

While the detector is running, the deck area at the top of the cavity is sealed from the rest of the detector to inhibit radon movement and potential light leaks. Repairs or inspections to areas nominally submerged such as the outer AV, PSUP, or cavity can only be performed by first lowering the UPW level and deploying an inflatable boat.^f

Of the UPW within the cavity, 1700 tonnes fill the space between the AV and the PSUP (the inner cavity) while the 5700 tonnes fills the volume between the PSUP and

^fThese operations have set the world record for deepest underground boating trip.

the cavity walls (the outer cavity). This water shielding protects the AV from backgrounds inherent in the PMT array materials and trace radioactivity from the cavity lining while further suppressing backgrounds from the surrounding norite. UPW is used as the shielding material, as the concentrations of radioactive contaminants within can be purified to several orders of magnitude lower than metal-based techniques [141].

2.2.3 Water Purification Systems

The “ultrapure” status granted to UPW is a technical standard set by the American Standards for Testing and Materials (ASTM) *D5127-13 Standard Guide of Ultra-Pure Water Used in the Electronics and Semiconductor Industries*, with select purity standards listed in Table 2.2. However, the water provided to SNOLAB is sourced from municipal water from Lively, Ontario, made potable by Vale in a surface purification plant. Although safe for human consumption, this potable water still contains sand, silt, bacteria, algae, inorganic salts, organic molecules, and gasses. The water is also super-saturated with air after falling over 2 km into the 6800-level of the mine, and therefore contains unacceptably high levels of background-producing ^{222}Rn . To achieve acceptable radiopurity levels for usage in SNO+, the water must be further treated upon reaching the underground lab through the same UPW plant built for SNO. A block flow diagram for the UPW can be seen in Figure 2.4.

The water is primarily purified using a reverse osmosis (RO) process, which uses twelve thin-film semi-permeable membranes comprised of polyamide on polysulfone to remove contamination. The RO process is more effective at removing larger and heavier particles, but is susceptible to damage or clogging by particulates or certain

Table 2.2: Select purity standards for UPW, as set by the American Standards for Testing and Materials (ATSM) D5127-13 Standard Guide of Ultra-Pure Water Used in the Electronics and Semiconductor Industries.

Parameter	Standard
Resistivity (25°C)	> 18.18 MΩ·cm
Total Organic Carbon	< 1 μg/L
Dissolved Oxygen	< 10 μg/L
Particle Count (> 0.05 μm)	< 200 particles/L
Silica	< 50 ng/L
Metals	< 1 ng/L
Ions	< 50 ng/L
Bacterial	< 10 CFU ^g /L

chemicals. Therefore, the water is pre-treated within the plant prior to RO purification.

In this pre-treatment, O₂ and N₂ bubbles are first removed through a deaerator tank before passing through both a 10 μm filter to remove particulates. The water is then passed through an additional charcoal filter to remove both organic contaminants and free chlorine, the latter of which is damaging to RO systems. Following this initial filtration, the water is then softened with Purolite C100-E cation exchange resin, exchanging divalent (e.g. Ca and Mg) ions for Na, removing iron, and stabilising colloidal particles to prevent coagulation in the RO system. A sodium ethylene diamine tetraacetate (EDTA) and sodium bisulfate (NaHSO₄) solution is injected to remove cations, and to react with O and Cl to form heavier substances. Finally, the water is passed through a 3 μm filter before being sent to the RO system.

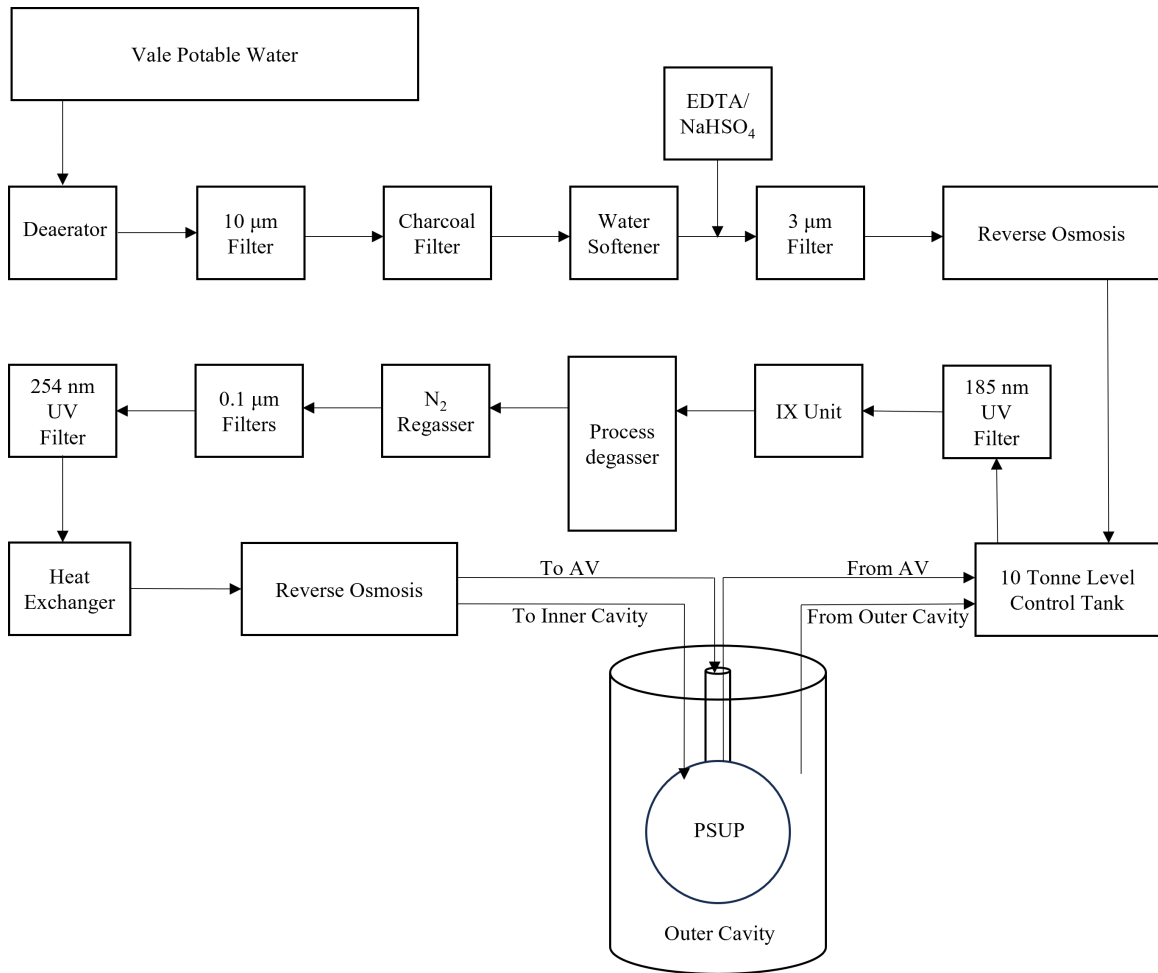


Figure 2.4: A block diagram of the UPW plant used for the SNO+ cavity, and to fill the AV during the SNO+ water phase.

The RO process provides a $> 95\%$ reduction in inorganic salt levels, and a $> 99\%$ reduction in organic molecules and particles with a molecular weight of ≥ 200 [129]. This performance is continuously verified using live online conductivity measurements. Following the RO process, the water is then subjected to a 185 nm UV filter to kill bacteria and ionise organic impurities, which are then purged in an ion-exchange (IX)

unit with Purolite nuclear grade NWR-37 mixed bed resins. The IX unit is also the main agent for removing metallic impurities.

To mitigate ^{222}Rn contamination, a process de-gasser is used to reduce O_2 and Rn levels by a factor of 1000 and 50, respectively [129]. However, this de-gassed water causes a partial vacuum on the PMT High Voltage (HV) as a result of diffusive negative pressure, thereby inducing electrical breakdowns when submersed within the cavity. Therefore, the water is regassed to atmospheric equilibrium^h with grade N5.0 (99.999% purity) N_2 gas before being subjected to 254 nm UV and $0.1\ \mu\text{m}$ filters. Finally, the UPW can be polished in an additional RO unitⁱ and cooled to the 10-12°C range before it is added to the cavity; the cooler water suppresses bacterial growth and combats the natural rock temperature of 40°C.

Altogether, this purification system is able to produce UPW at 130 L/min with impurities at or below $\mathcal{O}(10^{-15})\text{g/g}$ U or Th equivalent^j; these contaminants are further discussed in Chapter 3. Within the cavity, the UPW is added to the inner layer within the PSUP, reflecting the higher cleanliness priority of the region closer to the AV. As the outer layer is more naturally contaminated from the large content of submersed material, a plastic barrier within the PSUP largely seals the two layers from each other. The water systems are also capable of recirculating the cavity by pumping water from the outer layer and back into the UPW plant, where the IX unit is the primary tool used to polish the recirculated water.

The UPW produced within this plant was deployed within the AV and used as the SNO+ target medium during the water phase. During this deployment operation, the AV and cavity levels were filled in parallel to preserve the pressure and structural

^hThis is roughly $1.25 \pm 5\%$ atm at the depth of SNOLAB.

ⁱTo date, this polishing RO unit has been bypassed.

^jGrams U or Th for each gram of UPW.

integrity of the vessel. As the vessel was filled, inflatable boats were deployed in the cavity to repair any pinhole leaks that had formed during SNO operations or the SNO–SNO+ transition period.

After the AV is filled, the volume within could be recirculated by removing water through six sample lines extending into the AV. These sample lines can also be used to assay the AV volume.

2.2.4 Acrylic Vessel

The SNO+ AV is the heart of the experiment, and the detector structure within which the target medium of each phase is deployed. Inherited from SNO, the AV is spherical in shape and 12 m in diameter. A 7-m tall acrylic cylinder (“neck”) with a 1.5-m inner diameter (ID) is attached to the top of the AV through a joint called the “neck boss,” and connects the inner volume of the AV with the access deck above the cavity UPW. The neck at the deck thus checks the only access point to the inner volume, and is capped with a stainless steel structure known as the “universal interface” (UI). All materials within the AV from the target medium to internal calibration sources are deployed through connections on the UI, including a set of stainless steel “bubbler” lines into which N₂ gas is periodically sent in order to measure the AV liquid level. As the only AV access point, the UI and surrounding area is encapsulated within a tent that forms the Deck Clean Room (DCR), the class-200 inner sanctum of the experiment.^k

^kAs SNOLAB is a class-2000 cleanroom, the DCR has fewer airborne particulates than the rest of the lab by a factor of 10.

The AV is suspended from the deck using 10 1.85-cm diameter Tensylon™ “hold-up” ropes. These Tensylon ropes replaced the Vectran ropes used in the SNO experiment. As the Tensylon ropes use less material, this further reduces sources of radioactive contaminants in the inner cavity. Furthermore, the Tensylon material contains a lower concentration of ^{40}K , a low-energy background of concern in the scintillator phase [131]. As the liquid scintillator has a density of 0.86 g/cm^3 at the nominal 12°C temperature, the AV is expected to exert a $1.25 \times 10^6 \text{ N}$ buoyant force during the scintillator and tellurium phases of the experiment. To counteract this buoyancy, 5 additional “hold down” rope-nets were wrapped around the top of the AV and bolted more than 1.8 m into the cavity floor. Each consisting of two pairs of 3.69-cm diameter Tensylon ropes, this rope-net design crosses around the neck and was determined through Finite Element Analysis (FEA) to apply acceptable levels of stress on the AV [142]. A diagram of the AV and both rope systems, as well as the maximum strain applied by the hold-down rope-net on the AV, can be seen in Figure 2.5.

To observe creep in the tensile strength of the ropes and AV buoyancy differences caused by changing detector media, load cells at each of the hold-up and hold-down rope anchor points also allow for active live on-line measurements of the various rope tensions. The tensions and AV position can then be manipulated by manually adjusting the hold-up ropes at their anchor points on the deck.

The AV is composed of 122 thermoformed ultraviolet transmitting (UVT) panels made by Reynolds Polymer Technology Inc., which forms a sphere when assembled. This shape is advantageous due to the large volume-to-surface ratio and optimal stress distribution, thereby minimising the mass of acrylic needed, which in turn reduces

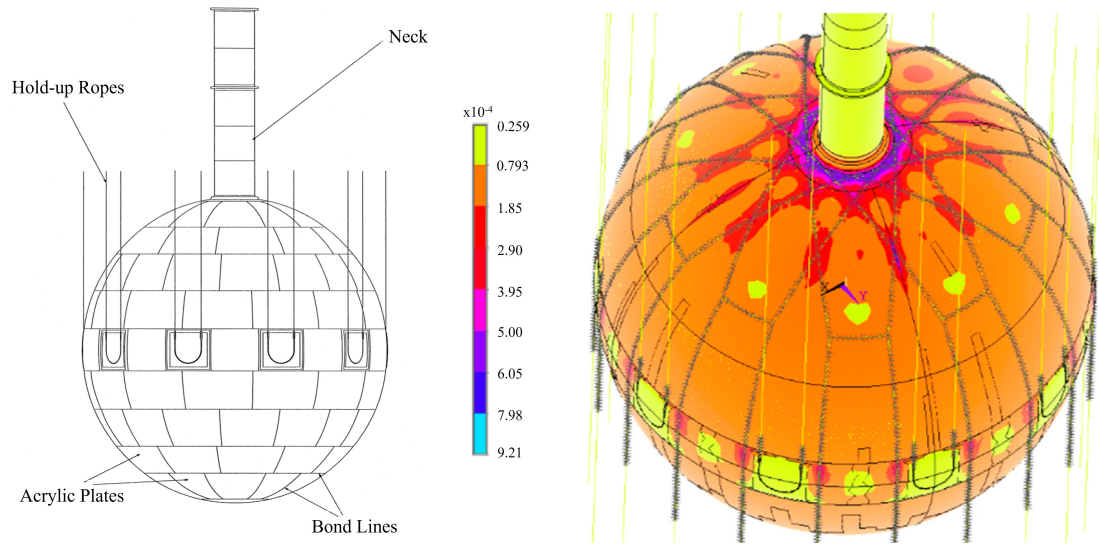


Figure 2.5: (Left) A diagram of the AV, showing the positions of the acrylic panels including those which contain grooves for the hold-up ropes to loop around. (Right) The position of the hold-down ropes, and an FEA simulation of the maximum strain they cause on the acrylic in units of strain. The hold-down ropes can also be seen in the photograph of the detector shown in Figure 2.1. Further details on the hold-down rope system including the strain measurements are published in [142].

potential acrylic-sourced contamination. The AV panels themselves are constructed entirely of Poly(methyl methacrylate) (PMMA, acrylic) with a chemical formula of $(C_5O_2H_8)_n$; as a simple hydrocarbon, the acrylic can be manufactured with very low intrinsic radioactivity. 112 of the panels are 5.6 cm thick, while 10 panels around the equator are 11.4 cm thick as they contain grooves for the hold-up ropes to loop around. To construct the AV, each panel was individually shipped underground before they were bonded together using a proprietary Reynolds Polymer Technology Inc. adhesive in a 3 mm thick layer, with a measured bond strength exceeding 27.5 MPa.

2.2.5 Cover Gas Systems

To protect the AV volume and cavity UPW from gaseous contamination, two cover gas systems — one each for the AV and cavity — fill the headspace above the liquid volumes of the AV and cavity with N₂ gas, therefore providing a physical barrier between the liquids and the Rn-rich (123 Bq/m³) laboratory air. The nitrogen used in both systems is rated to have a maximum oxygen concentration of 2 ppm and water concentration of 3 ppm. The cavity cover gas system was inherited from SNO, and is a flow-through system where the gaseous headspace is flushed at a rate of 5 L/min. This nitrogen is fed from boil-off liquid nitrogen (LN₂) and contains $< 10^{-5}$ pCi Rn per L of nitrogen gas.

In the scintillator and tellurium phases, the purity of the cover gas is even more essential, as ingress of water or oxygen could lead to a reduction of scintillator stability. Furthermore, daughters from ²²²Rn decays present a background in some low-energy regions. Thus, a fully sealed system for the AV cover gas was developed and implemented. As a sealed system, changes in the external mine air pressure could cause damage to the AV.¹ To account for this, three cover gas bags were coupled to the UI and provide up to 360 L of expansion volume to compensate for these pressure changes. A two-way pressure safety device (PSD) and set of U-traps provide additional protection to large swings in pressure. The Rn content of the AV cover gas system is actively monitored online through a dedicated radon monitor with a sensitivity of $\mathcal{O}(1)$ mBq/m³. This AV cover gas system, which has been measured to have a Rn reduction factor of 1.9×10^{-4} [131], was the last major detector system to be implemented in October 2018 — over 17 months after the start of the water phase.

¹The AV is rated to have a permissible pressure difference of ± 0.28 psi between the inside and outside of the vessel.

Therefore, the water phase data is split into an initial dataset, and lower background dataset with the fully commissioned cover gas system.

2.2.6 Photomultiplier Tubes

Events within the AV are observed with the same Hamamatsu R1408 PMTs that were originally installed for the SNO experiment. 9362 inward-looking PMTs are ever-watchful of events within the AV; 49 have a dynode tap and a second signal cable capable of reading out low-gain signals, thereby increasing the dynamic range of the detector at higher light intensities. These PMTs are supplemented by 4 PMTs installed on the UI to reject events that occur at the top of the neck, and 91 outward-looking (OWL) PMTs assist in vetoing events occurring external to the detector.

These 20.4-cm diameter PMTs were originally selected for their detection efficiencies, noise rate, transit-time spread, and radioactive trace contamination. A schematic of this PMT design is shown in Figure 2.6. Prior to installation, the raw materials of each PMT component was assayed for radioactivity [143]. Although the dominant contributor of radioactive contamination was the PMT glass, the Th/U impurities levels were measured to be only 40 ppb. Such low contamination was achieved through a glass-making process designed specifically for SNO by Schott Glaswerke, which used a glass furnace refitted with a low-radioactivity liner.

The presence of the $55\ \mu\text{T}$ magnetic field from the Earth causes electron drift in the PMTs, thereby reducing the PMT efficiency by as much as 18%. To counteract this field, 14 horizontal magnetic field compensation coils were installed, reducing the maximum residual field in the PMT region to $19\ \mu\text{T}$, thereby decreasing the magnetically-induced reduction in PMT efficiency to $\sim 2.5\%$.




Figure 2.6: (Left) A schematic of the Hamamatsu R1408 PMT, including major components. (Right) A schematic of the PMT-concentrator assembly mounted onto the PSUP hex cell. Both figures adapted from [129].

As the PMTs were deployed underwater, a waterproof base with two water barriers was designed and attached to each PMT. An outer waterproof plastic housing was also filled with waterproof silicone dielectric gel (Ge RTV6196) to provide additional water protection. Despite these protections, nearly 800 of the 9829 PMTs failed during SNO, with a majority due to electrical shorts in the base. 600 PMTs were repaired and re-implemented in SNO+,^m and the experiment started the water phase with 9362 inward-facing PMTs. Thus far in the lifetime of the SNO+ experiment, 1–2 PMTs have been lost per week due to component failure while underwater — in line with the failure rate during SNO. These failures occur uniformly throughout the

^mThe remaining 200 were not repaired due to time constraints.

detector and are accounted for in both detector simulations and data analysis, and have not meaningfully impacted the position or energy resolution of the detector.

Together, the PMTs within SNO+ provide a photocoverage around the sphere of 31%. To improve this efficiency, each PMT was fitted with 18 reflective petals arranged in a 27-cm diameter Winston Cone called a “concentrator”. The petals were made of dielectric-coated aluminum sheets, treated with a lower active coating of specular aluminum covered with magnesium flouride, an upper coating of titanium and praseodymium oxides, and mounted on a acrylonitrile butadiene styrene (ABS) plastic holder. The concentrators not only increase the effective area, but limit the angular acceptance of the PMTs to a three-dimensional critical angle of 56.4° such that the view is limited to a radius of 7 m from the center of the PSUP,ⁿ thereby increasing light collection efficiency by $\sim 75\%$ [144]. A schematic of how the PMT and concentrators are coupled is shown in Figure 2.6. After taking into account the $82 \pm 3\%$ reflectivity of the petals, the concentrators originally increased the effective photocoverage of the AV by the SNO+ PMT array to 54%. However, degradation in the petal coatings since their installation has decreased the reflectivity of the concentrators over time.^o

2.2.7 PMT Support Structure

The PMTs are mounted onto the PSUP, a stainless steel geodesic support structure that encapsulates the AV. The PSUP separates the cavity into an inner and outer volume. The cables that connect the PMTs extend from the PSUP, through the

ⁿAs the AV has a radius of 6 m, this is advantageous as it restricts events far in the external volume.

^o5 of the installed PMTs have their concentrators removed, allowing for concentrator effectiveness to be tested.

outer volume, and up into the deck where the SNO+ data acquisition and electronics systems are housed. As each PMT has an independent connection, the thousands of cables that extend into the outer cavity makes it infeasible to control the potential contamination level within the outer cavity at the same level as the inner cavity. As such, the PSUP was designed to be approximately watertight to mitigate potential water-borne contaminants in the outer cavity from leaking into the inner cavity. The PSUP also supports a number of other detector infrastructure features, including some of the cavity water recirculation piping, external calibration and camera systems, and other monitoring tools.

The PSUP is a three-frequency icosahedron composed of 92 connections between 270 stainless steel struts. 751 triangular ABS panels — each with 7–21 PMT cells capable of mounting a PMT — are then mounted between the struts. The panels are tessellated, and oriented such that every inward facing PMT is capable of observing the entire AV volume. The structure nearly fully encloses the AV; the topmost node is a hollow toroidal ring to allow the neck to pass through, and 2–4 PMT cells in 20 of the bottom panels are used to feed through the AV hold-down ropes. A further 60 PMT cells are utilised in the top hemisphere of the PSUP to accommodate the hold-up ropes and external “guide tubes”, used to deploy calibration sources externally to the AV into the inner cavity.

The PSUP is suspended on 15 stainless steel wire rope cables from the deck; the cables are monitored actively with load cells and can be manipulated to change the position of the entire array by ± 10 cm. The PSUP is also anchored onto the cavity floor to counteract buoyant forces while submerged within the cavity. As the wire

ropes used cannot be manufactured without lubrication, the cables are all coated in plastic housing to prevent contamination from leaching into the cavity UPW.

2.3 The Data Acquisition System

The DAQ is a set of electronics systems that forms the road leading from a PMT detection to an analysable data file. While many of the core SNO+ detector hardware systems such as the AV were inherited from SNO with only minor changes or refurbishments, the SNO DAQ required several upgrades for use in SNO+. As the use of liquid scintillator would increase the light yield within SNO+, the number of PMT hits for each event was expected to correspondingly increase. New, overhauled, or retrofitted electronics systems were developed for SNO+ to account for this increased signal load. Indeed, while the SNO experiment had events with $\mathcal{O}(10)$ PMT signals within each 400 ns timing window at an event rate of $\mathcal{O}(10)$ Hz and an energy threshold of 3.5 MeV, the fully commissioned SNO+ experiment now regularly measures events in excess of thousands of PMT signals at a nominal data-taking rate of ~ 2200 Hz and an energy threshold below 1 MeV. Despite the major hardware upgrades required and implemented, SNO+ nevertheless retained the electronics readout methodology of SNO [146]. A schematic of the SNO+ DAQ system that outlines systems heavily upgraded or newly installed between SNO and SNO+ is shown in Figure 2.7.

2.3.1 The Signal Flow

The PMT electronics are housed within by 19 isolated electronics crates installed as sets of one or two onto 11 racks located on the SNO+ deck. Within every crate are 16 PMT Interface Cards (PMTICs), each associated with a front-end card (FEC). Each

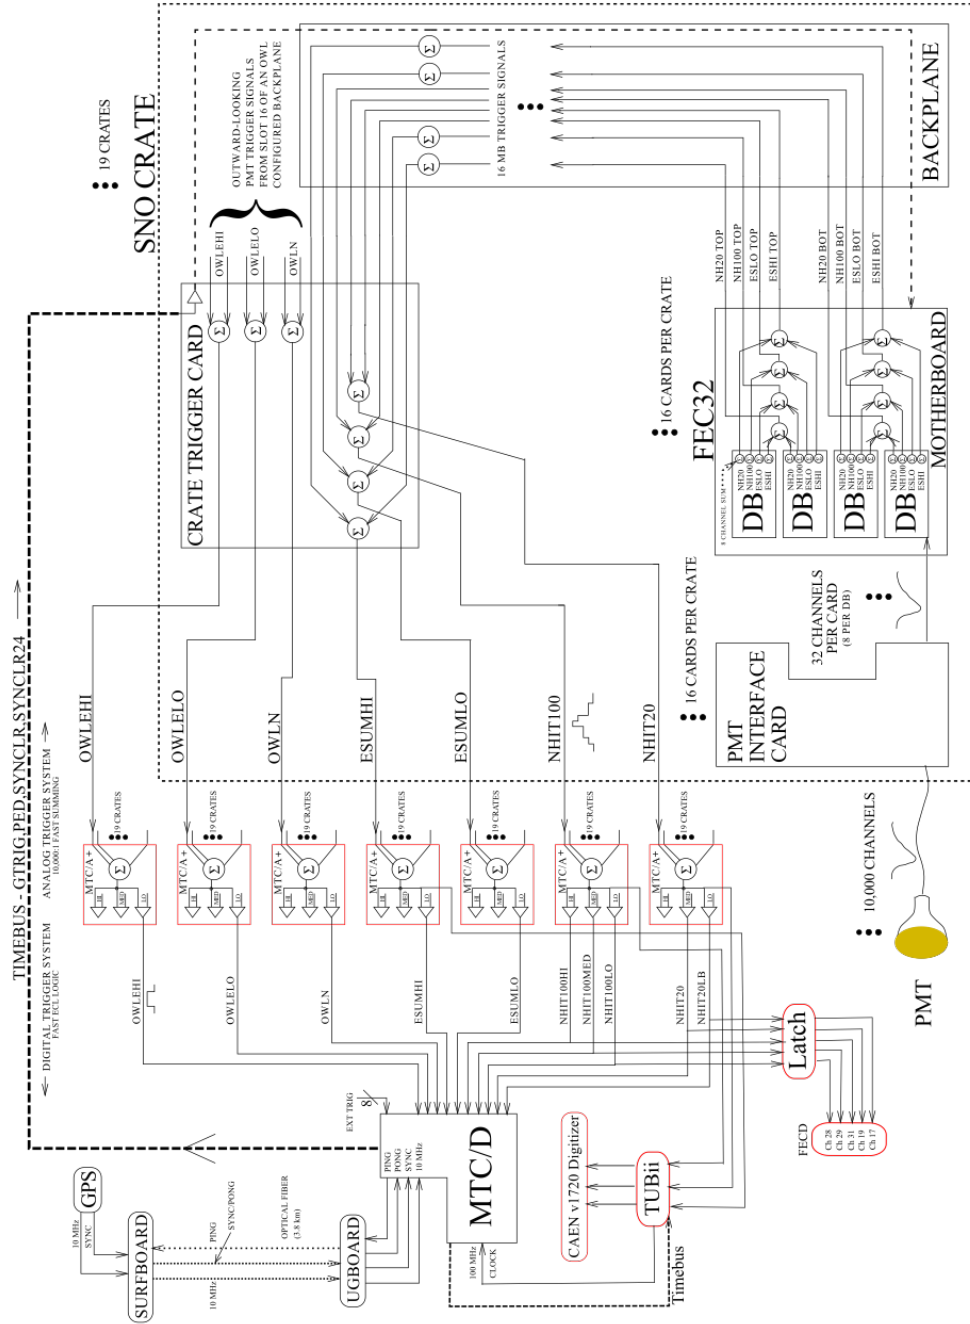


Figure 2.7: The schematic of the DAQ system. Outlined in red are systems that were heavily upgraded or newly installed between SNO and SNO+. Figure adapted from [145].

PMT is connected to a PMTIC through a 32-m $75\ \Omega$ RG59-like waterproof coaxial cable. Signals detected by PMTs are carried through this cable to the PMTIC, where the information is then processed on the associated FEC. Mounted upon each of the 16 FECs are four daughterboards (DBs) capable of reading out eight PMTs; every FEC–PMTIC pair is capable of servicing 32 PMTs for a total of 512 across the crate. The PMTICs and FECs are connected through the crate backplane, which runs through every FEC–PMTIC coupling. Connected directly to the backplane of each crate is a dedicated high voltage (HV) source, which nominally supplies ~ 60 mA to each crate and ~ 2000 V to each PMT through the PMTICs.

The electrical information received by the PMT output is sent as analogue signals along the cable through to the PMTIC, which controls the signal flow from each PMT pipeline through individual isolating, filtering, and trimming networks. The PMTICs also include diodes and fuses which prevent electrical shorts (“breakdowns”) typically caused by component failure from damaging any other part of the detector. The PMTICs also host relays that provide HV disconnect capabilities to sets of eight PMTs, important for identifying and repairing malfunctioning electronics in the case of a breakdown. Also on the PMTICs are feed-through resistors that provide HV tuning capabilities for each PMT, allowing for an adjustable PMT gain; each of the nearly 10,000 PMTs were individually characterised to determine the HV required for a gain of 1×10^7 [147].

Following initial filtering by the PMTIC, the signal then passes directly onto the associated FEC and is sent to an assigned DB for signal processing. Upon the DB are three chips: a discriminator (SNOD), integrator (SNOI), and complementary metal-oxide-semiconductor (CMOS) chip. The signals are first sent through the SNOD

and compared against adjustable voltage settings (“channel thresholds”) to eliminate electronic noise. If the signal exceeds the channel thresholds, it is saved as a PMT “hit” and passes a copy of the original pulse shape (ESUM) to both the SNOI and CMOS chips.

The SNOI calculates a high gain and low gain version of the ESUM pulse (ESUMH and ESUML, respectively), and produces charge values in three varieties: high-gain and long-integration (QHL), high-gain and short-integration (QHS), and low-gain and variable-integration (QHX).^P Simultaneously, the CMOS chip prepares timing information through a time-to-amplitude (TAC) voltage ramp and creates two fixed-current square pulses with a width of 100 ns (N100) and 20 ns (N20). The CMOS chips are also capable of disabling these pulses for individual channels, important for debugging and hunting for PMT failures on the detector. The OWL PMTs are processed independently through the same methodology, but with their square, high, and low gain pulses saved as OWLN, OWLEH, and OWLEL, respectively. The CMOS chips temporarily store the processed charge and timing information from each hit (as well as metadata such as the FEC and channel number) in analogue memory cells, and awaits permission to release the information in the form of a Global Trigger (GT) pulse from the trigger system. The data is temporarily stored for a time period of GTValid (nominally 400 ns); if GTValid expires without a GT signal, the data is permanently erased.

2.3.2 The Trigger System

While the processed hit information from each PMT is saved on their individual DBs, the channels act as independent detectors. However, the entire PMT array must be

^PIn general, a short integration time is used with QHX.

synchronised to discriminate electronics noise from potential particle interactions. This coordination is performed through the trigger system, which only saves the data if enough PMTs are coincidentally hit within a particular timing window.

The first stage of coordinating between the different PMTs is performed on the Crate Trigger Card (CTC) of each crate, which sums the N20, N100, ESUMH, and ESUML pulses from each of the individual hits for the crate. The four crate-level summed pulses from each of the 19 crates are then passed onto the central nervous system of the detector: the timing rack.

Within the timing rack, the summed crate signals are passed to the seven analogue master trigger cards (MTC/A+). Four of MTC/A+ cards are each responsible for summing the 19-crate total for one type of pulse: N20, N100, ESUMH, and ESUML. The other three are responsible for summing the OWL pulses OWLN, OWLEH, and OWLEL. Each MTC/A+ has a low, medium, and high gain path that can be used; for example, the MTC/A+ is capable of producing a summed N100L, N100M, and N100H signal. Of the 21 potential signal paths, only a handful are typically used based on the data taking conditions. Each of the raw MTC/A+ pulses are then sent to the digital master trigger card (MTC/D). Additional custom signals are added for a total of 26 channels, including the PULSEGT which fires an electronics sampling signal at 5 Hz.

Within the MTC/D, the types of signals that are considered are set by an adjustable trigger mask. For each raw signal that is masked in (i.e. allowed to proceed), the MTC/D applies a tunable threshold. If any signal passes, the MTC/D sends out the GT pulse, which prompts the entire detector to begin recording PMT data. The type of pulse that triggered the GT is saved alongside all other trigger pulses that

cross the MTC/D threshold within a 20 ns time window. The 100L trigger has the lowest trigger threshold of ~ 7 hits and is the primary trigger for the detector.

The time-keeping of every event is synchronised using an underground timing board (UGBoard), which contains a 10 MHz oscillator. The 10 MHz count is synchronised via a 4 km fibre to Universal Coordinated Time (UTC) using a commercial GPS device connected to a surface timing board (SURFBoard) every hour to associate each event with a date and time. Meanwhile, the Trigger Utility Board Mk.II (TUBii) is connected directly to the MTC/D and contains a 100 MHz clock, and provides each event with a 50 MHz (20 ns) timing resolution. TUBii can also be programmed to produce GTs through logical combinations of other raw trigger signals, providing additional flexibility on data taking capabilities.

The MTC/A+ cards were upgrades of the MTC/A cards used in SNO, and allow for simultaneous use of up to 10,000 PMTs with single PMT resolution. The other major trigger system hardware upgrade was the replacement of the SNO TUB with TUBii, which provided a plethora of additional features beyond advanced trigger logic such as signal shaping, an interface for optical calibration systems, additional logic converters, and advanced monitoring tools.

2.3.3 The Data Flow

When the GT is released by the MTC/D, a pulse is sent through the backplane of every crate to every channel on every DB, instructing for the data readout to initiate. The MTC/D also creates a lockout of 420 ns, preventing the detector from triggering again and thereby giving the data-readout process an operational time limit of 400 ns, with a 20 ns “dead time” reset period. For each GT, the MTC/D assigns an iterating

unique number known as the GTID. The GTID of every event is logged independently within each channel and is verified against the MTC/D every 2^{16} GTIDs. PMT hits within a “trigger window” of ${}_{-220}^{+180}$ ns of the original trigger pulse that are currently saved on the DB are sent to a buffer on the FEC via a $2\ \mu\text{s}$ analogue-to-digital (ADC) converter. All hits within a trigger window make up a potential particle interaction within SNO+ known as an “event”; the number of PMT hits within each event are called “nhits”.

The XL3 card⁹ outfitted on each crate then outputs the data over TCP/IP protocol via an Ethernet to a central DAQ computer. As the ADC converter has a $2\ \mu\text{s}$ processing time and GTIDs may be firing as quickly as the lockout period of 420 ns, the data is then sent to a “builder” client, which uses GTID and trigger information to consolidate data into the correct events before writing the full event information to disk in the form of a custom `.zdab` file format. Also saved are full N20, N100, and ESUMH trigger signals, which are digitised by a CAEN v1720 waveform digitiser connected to TUBii and used for monitoring the health of the trigger system. The trigger signals for N20M, N20L, N100H, N100M, and N100L are also sent through the Latch Board to a spare FEC (“FECD”) which can also be read out.

The XL3 card replaces the XL1 and XL2 cards used in SNO, which were limited to single crate communications through VME-like protocols. The XL3 hosts an on-board Xilinx ML403 integrated circuit and PowerPC processor which allows for independent, asynchronous, parallel readout of the entire crate and pushes data on a TCP/IP protocol, ultimately increasing the total detector data-readout bandwidth from 250 kB in SNO to 266 MB/s in SNO+. In the water phase, this allowed for higher trigger rates — and thus a lower energy threshold — than was possible in SNO. In

⁹Being a translator card, XL is presumed to be an abbreviation for transLator.

the scintillator phase, this accommodated the far more plentiful data afforded by the higher light-yield liquid scintillator. Furthermore, the CAEN digitiser replaces the SNO analogue measurement board (AMB) and provides digital waveforms for advanced data cleaning and monitoring tools, while the Latch Board — new in SNO+ — provides a capability to determine trigger efficiencies.

2.3.4 Detector Operations

The data is continuously saved into the same `.zdab` file until a maximum file size is reached. The `.zdab` files are numerically sorted by detector “runs”, which represent periods of time when the detector was operating under the same conditions. A data-taking run nominally iterates every one hour and is the basic package of data used for analysis. However, the run number is also frequently manually increased to reflect different or changing detector conditions.

The detector is actively monitored by a live operator at all times while HV is active in any part of the detector. This allows for the recording of any qualitative information pertaining to a run that could affect the quality of the data and a human verification that the desired suite of trigger thresholds or other detector settings are applied for the current detection mode, be it physics data-taking, detector calibrations, or electronics maintenance. Most importantly, the operator facilitates rapid response and repair times for any potential detector issues, from PMT failures to changes in hold-down rope tensions. Upon any changes to the detector, these operators record relevant information and manually iterate the run number, thus facilitating the basic data organisational structure for analysis.

2.4 Detector Calibrations

Despite the identification and stewardship of potential physics events afforded by the DAQ systems, the data gleaned are no more than sets of simple timing and charge information. A series of calibrations aimed at understanding the PMT and electronics responses must be performed to decipher these basic parameters into physical descriptions of the detected events. There are three main calibration techniques used:

- optical sources which emit photons of known wavelengths within the detector to measure absorption and scattering of the detection media, as well as the charge, timing, and angular response of the PMTs,
- radioactive sources, which emit particles of known energies within the detector to connect PMT pulses to incident particle energies, and
- electronics calibrations, in which the various channels and baselines used in the DAQ process are measured and accounted for.

2.4.1 Source Deployment Hardware

One optical source and all radioactive sources consist from objects that can be lowered into the detector through a rope-based source manipulator system (SMS). These objects (“deployed sources”) can be manoeuvred directly into the AV through the UI (internal deployments) or externally at certain locations (external deployments) through the use of one of six guide tubes that extend between the AV and PSUP. An overview image of the calibration deployment systems can be seen in Figure 2.8.




Figure 2.8: An overview of the calibration source deployment systems. Not to scale. The inset photograph is of the deployed laserball calibration source discussed in Section 2.4.2, taken by one of the permanently deployed underwater cameras.

While deployed, sources are connected to the deck using a multi-purpose 30-m long “umbilical” cable which encloses the gas capillaries, optical fibres, signal wires, or any other connections required to operate a particular source. Multiple umbilicals were produced for the various different sources in accordance of their specific needs. The umbilical is made up of an inner low density polyethelyne (LDPE) tubing that protects the gas lines or optical fibres (depending on the source type), helically wrapped with four 24-AWG hook-up wires and a 30-AWG coaxial cable for signal transmission. The outer 0.5-inch diameter tubing was made of silicone tubing in the water phase and the space between the two tubes were filled with silicone gel. For the scintillator phase, new umbilicals were made with Tygothane[®] Precision Polyurethane tubing. These new umbilicals have improved material compatibility with the scintillator while also emanating less Rn - a background with much more stringent requirements in the scintillator phase.

Deployed sources are hooked into the umbilical,^r and lowered from the deck into either the AV or one of the guide tubes with the use of an Umbilical Retrieval Mechanism (URM). The URM possesses an umbilical storage system as well as a winding drum for a central Tensylon rope which holds the weight of the deployed source. A custom software called MANMON manipulated the vertical position of a source at mm-level precision by driving the winding drum. Since the URM cannot control rotation of the source while deployed, an Umbilical Flasher Object (UFO) containing eight LEDs can be fixed to the source interface. This allows for a determination of the source orientation such that asymmetries in the source can be accounted for. The aluminum URMs inherited from SNO were used during the SNO+ water phase. In

^rFor the scintillator phase, a stainless steel source interface was constructed for this connection.

the scintillator phase, a new URM was built; the new URM is fully sealed to prevent the ingress of radon, which can go on to contaminate the AV. A cross-sectional illustration of both a URM and umbilical can be seen in Figure 2.9.




Figure 2.9: (Left) A cross-sectional illustration of a URM, showing (1) the umbilical storage system, (2) the drive system, (3) the rope mechanism, and (4) an acrylic viewing port for inspections.

(Right) A cross-sectional illustration of the umbilical, showing (1) the outer silicone or Tygothane tube, (2) silicone gel padding (3) the inner LDPE tube, and (4) the inner umbilical which feeds gas capillaries, or fibre optics lines wrapped in (5) their own protective coating. The LDPE tube is helically wrapped with (6) a coaxial cable and (7) four hook-up wires.

Both figures are adapted from [131].

The URM interfaces with either the AV or the guide tubes through a stainless steel source tube attached to the opening. During the water phase, the sources were manually attached to the umbilical and rope mechanism within the source tube, which was then fully lowered and mounted onto a flange on either the UI (for internal deployments) or the guide tube (for external deployments). Due to the more stringent purity requirements needed during the scintillator phase, this manual operation may inject unacceptable levels of contamination. Different methodologies of mounting,

swapping, and storing the sources without requiring exposure to ambient air — and therefore radon — are still under investigation, including an interface to an enclosed source box.

Gate valves mounted to the flanges on both the guide tubes and UI separates the deck from the inner cavity or AV until the URM system is properly mounted. On the UI, there are three gate valves of different sizes to accommodate various source sizes; a different source tube and URM would need to be used to access each size. A schematic on the UI highlighting features for calibration deployments is shown in Figure 2.10. After the URM and source tube is firmly sealed to the flange, the gate valve can be manually opened. The source can then be lowered through usage of MANMON to manipulate the length of rope deployed.

During internal deployments, two of four “side-ropes” can be manually hooked using butyl gloves installed on the side of the UI onto pulleys that are built into the side of the source interface. One end of each side-rope is permanently deployed down the neck and attached to the side of the AV. The other end is reeled onto a winding drum on the UI within two sets of opposing motor boxes. While the source is deployed within the AV, these side-rope tensions can also be manipulated by MANMON to move the source off of the central vertical axis. The two sets of opposing side-ropes thus allow for a full range of movement along 2 orthogonal planes. This side-rope technique was also used by SNO. However, similar to the hold-up ropes, the SNO vectron side-ropes were replaced by tensylon ropes which contains fewer ^{40}K contaminants.




Figure 2.10: A schematic of the UI, showing (1) the side rope motor boxes, (2) an extension that connects the motor boxes to the AV cover gas system, (3–5) three flanges accessed by gate valves, (6) two of the four sets of glove ports, and (7) view-ports (with removable covers) for visual inspections. Figure adapted from [131]

2.4.2 Optical Sources

There are two standard types of optical calibration sources used within SNO+. The first is a light-diffusing sphere (“laserball”), while the other is the Embedded LED/Laser Light Injection Entity (ELLIE), an array of optical fibres mounted on the PSUP.

Laserball The laserball is a deployed source that was the primary optical calibration technique used in the water phase. The light is injected into the detector using a 10.9-cm diameter quartz flask filled with silicone gel, within which 50 μm glass beads are suspended homogeneously. The quartz flask receives light through 20 fibre optic cables within the laserball umbilical from a nitrogen dye laser in the DCR and diffuses the light through the quartz sphere. This light diffusion is quasi-isotropic, with the only asymmetry from the bead distribution and equipment required to mount the laserball to the umbilical. The nitrogen dye laser is typically pulsed at a rate of 10–40 Hz and has a central wavelength of 337 nm, but can also be used to drive a dye laser, providing 5 additional wavelengths between 365–500 nm. Throughout the water phase, the laserball was internally deployed at multiple locations and was used to determine the internal and external water absorption coefficients and absorption lengths, the effective attenuation coefficients for the AV, and the angular response of the PMTs (and concentrators). The analysis of the laserball calibration data used to determine these parameters was published in [148]. A new laserball with an improved neck design to reduce self-shadowing has been developed for the scintillator phase and has yet to be deployed.

ELLIE Unlike the laserball, the ELLIE system is permanently installed on the PSUP and is not a deployed source. ELLIE injects light into the detector using optical fibres that connect to the PSUP from the DCR. The ELLIE system contains three modules to measure different parameters:

1. the timing module (TELLIE) to calibrate PMT hit-time and gain,
2. the attenuation module (AMELLIE) for monitoring the attenuation lengths of the detector media, and
3. the scattering module (SMELLIE), which measures the optical scattering cross sections of the detector media as a function of scattering angle and photon wavelength.

The subsystems differ in their fibres and light injection techniques, as can be seen in Figure 2.11. TELLIE uses 96 optical fibres coupled to 500 nm LEDs that point towards the center of the detector with a wide beam profile (22° half angle). A timing calibration is performed by comparing the light injection time with PMT hit times to a precision of 0.1 ns. AMELLIE uses eight optical fibres, of which half are coupled to 505 nm LEDs and the other half coupled to 403 nm LEDs. Two fibres are located at each injection point, pointed at 0° and 20° relative to the detector center, providing the capability of monitoring relative changes in attenuation over time. Finally, SMELLIE uses 15 fibres in 5 injection sites which are coupled to a suite of lasers ranging from 375–700 nm, and narrowly spreads light (3° half angle) at angles ranging from 0° – 20° relative to the detector center. This can therefore check for changes in the wavelength- and angular-dependent Rayleigh scattering of the detector medium.

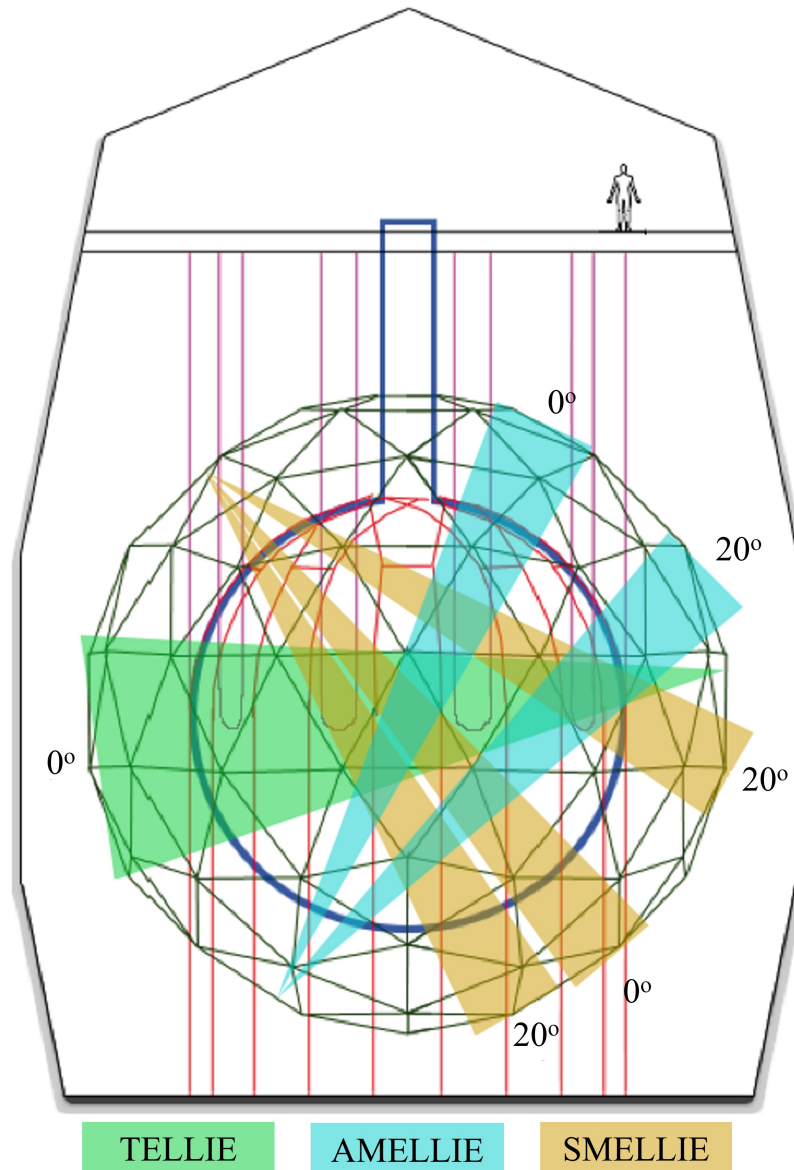


Figure 2.11: A diagram of the detector, showing one of the many light injection possibilities for each of the ELLIE subsystems. SMELLIE and AMELLIE can inject light at multiple angles, while TELLIE is always aimed at the center of the AV. Figure adapted from [131].

The various ELLIE light sources and support hardware are housed on the deck, where they can be operated without requiring access to the detector. This provides a distinct advantage over deployed sources, as their use does not risk contamination of the detector media.

2.4.3 Radioactive sources


In addition to the optical sources, three radioactive sources were deployed during the water phase: a ^{16}N , ^{241}Am - ^9Be (AmBe), and ^8Li “Cherenkov” source.

^{16}N Source The ^{16}N source inherited from SNO was the main energy calibration used in the SNO+ water phase, and deployed both internally and externally. ^{16}N undergoes β decay in the form of,



where $E_\gamma = 6.13\text{ MeV}$ with a branching ratio (BR) of 66.2%. As this ^{16}N decay occurs within the source, the e^- released is captured by a plastic scintillator that lines the source encapsulation and produces light that is detected by a dedicated PMT inside the source enclosure. Simultaneously, the γ is emitted from the source and Compton scatters in the detector medium (or cavity water if deployed externally) before producing light in the detector. This coincidence provides a powerful tool for quantifying the response of the PMTs when measuring a known energy. *Ex situ* measurements have determined the detection efficiency of this coincidence tag is over 95% [149].

To produce the ^{16}N isotope, a deuterium-tritium (DT) generator that accelerates DT particles at a DT target is first used to produce neutrons with the reaction,



A CO_2 gas is then passed through this DT generator, where the ^{16}N is produced through neutron interactions with ^{16}O ,



The resulting ^{16}N isotope ($t_{1/2} = 7.13\text{s}$) is then carried through a gas capillary within the umbilical and into the deployed source.

Although it was the main energy calibration in the water phase, it is unlikely that the ^{16}N source will be internally deployed during the scintillator phase. As the energy of the decay is well above the $0\nu\beta\beta$ ROI, the risk of contaminating the scintillator volume (inherent in any internal deployment) is likely too high considering the utility of the calibration. Instead, other radioactive sources will shoulder the weight of internal deployments. Nevertheless, an external deployment of the ^{16}N is still desirable, as it is capable of providing a direct comparison between water and scintillator phases.

AmBe Source The AmBe source consists of a mixture of ^{241}Am and ^9Be powders within a stainless steel capsule surrounded by a polyoxymethylene (Delrin) encapsulation. The ^{241}Am isotope α decays while emitting a $59.5\text{ keV } \gamma$ [BR = 84.6%]. The α particles are then absorbed by the ^9Be , with $\mathcal{O}(10^{-4})$ captures resulting in the (α, n)

reaction,



$\sim 60\%$ of the time, the ${}^{12}\text{C}$ daughter is produced in an excited state, which immediately de-excites while releasing a prompt 4.4 MeV γ . Meanwhile, the neutron is thermalised and typically captured on hydrogen in the detector, producing a delayed 2.2 MeV γ . This coincidence signal provides two additional energy calibration points near the $0\nu\beta\beta$ ROI.

The AmBe source is also used to measure the neutron capture cross-section and neutron capture time, which were determined to be $336.3^{+1.2}_{-1.5}$ mb and $202.35^{+0.87}_{-0.76}$ μs , respectively. The efficiency of detecting neutron capture on a proton was also found to be 50% with a 1% variation within the AV, and above 30% in the external water region between the PSUP and AV [150].

The AmBe source used during the water phase was inherited from SNO, but borrowed by other experiments in the intervening period prior to the SNO+ water phase. Since the history of the source usage — and therefore explicit estimate of cleanliness — could not be reliably tracked, an additional Delrin encapsulation was designed, machined, and installed for the water phase deployments. The final encapsulation was subjected to a pressure test at 60 PSI within an aluminium vessel, as well as an air leak test where the source was submerged in UPW and visually monitored for bubble formation. The air leak test was performed prior to each internal or external deployment, and the interior of this external encapsulation was inspected for moisture after each use for further verification.

Although the 59.5 keV γ released with the α decay from ${}^{241}\text{Am}$ falls below the ~ 260 keV kinematic energy threshold required to produce Cherenkov light in the

water phase, these γ s will present a background to the AmBe coincidence in the lower threshold scintillator phase. Therefore, a new encapsulation capable of absorbing these γ s without compromising the rest of the source output will need to be constructed, tested, and installed prior to AmBe deployment in the scintillator phase. Simulations of early designs have projected that utilising two 2.6 mm lead layers will block 99.996% of these γ s but only 5.9% of the neutrons.

Cherenkov Source The deployed Cherenkov source is based on a similar device used in SNO [151], and relies on the β decays of ${}^8\text{Li}$ within a 6 cm thick ultraviolet absorbing (UVA) acrylic shell with a black opaque inner lining. The ${}^8\text{Li}$ source is created by lining the same DT generator used in ${}^{16}\text{N}$ production for the ${}^{16}\text{N}$ source with a 99.9% enriched ${}^{11}\text{B}$ target, thereby inducing the (n, α) reaction,



The ${}^8\text{Li}$ then β decays as,



which is followed by the double α decay,



As ${}^8\text{Li}$ has a half-life of 0.8387 s, the charged cationic lithium ions are bound with anionic Na ions in aerated NaCl, produced by heating NaCl salt at 610°C. The isotope is then transported through the Cherenkov source umbilical to the deployed source

chamber at speeds limited by the laminar flow through the conduit using helium gas.^s The β particles released in the ${}^8\text{Li}$ decay creates Cherenkov light in the acrylic sphere, while the delayed α particles scintillate the helium gas and are either detected in a dedicated PMT within the source or stopped by the opaque inner lining. This source provides a tagged energy deposit with an end-point energy of 12.96 MeV, with only Cherenkov light produced. This allows for calibration of the PMTs independent of the properties of the scintillator.

The Cherenkov source was tested in the water phase. Due to the unavoidable presence of oxygen within the DT generator, ${}^{16}\text{N}$ events were also produced and released into the Cherenkov source, creating an ${}^{16}\text{N}$ $\beta - \gamma$ coincidence event. As the γ is released as in the case of ${}^8\text{Li}$ with the β picked up by the Cherenkov source PMT, this presents a background when attempting to identify tagged ${}^8\text{Li}$ events. However, it was demonstrated during the water phase test deployment that the α and β events could be discriminated by comparing the QHS-time distribution of hits from the internal source PMT, as seen in Figure 2.12. If desired, further discrimination could also be performed through pulse-shape discrimination techniques on the digital waveforms by connecting the source PMT to the CAEN digitiser.

Beyond those already constructed and deployed, several additional radioactive sources have been considered to provide more energy calibration points. This would further improve the energy calibration thereby reducing systematic uncertainties in the energy reconstruction of events. Some of the radioactive sources considered are listed in Table 2.3.

^sFaster delivery could be performed at near-supersonic gas flow speeds using high pressures, but were deemed impractical due to safety and engineering challenges [151].




Figure 2.12: Plotting the QHL charge against the time distribution of hits from the Cherenkov source PMT. α events resulting from ${}^8\text{Li}$ events have higher charges and slower timing, and can be therefore discriminated against the β events from the ${}^{16}\text{N}$ background.

2.4.4 Electronics Calibrations

To ensure the uniform response of all PMT channels, there are multiple charge offsets within the DAQ system that must be determined and accounted for such that PMT pulses can be correctly translated into charge quantities. These are measured through a suite of electronics calibrations (ECALs), which are performed either monthly or after any electronics have undergone repairs or upgrades. ECALs quantify baselines by running the detector without HV on the PMTs. As explained in Section 2.3.1, the PMT pulses are first sent to the SNOD on the DB for filtering, done through

Table 2.3: List of considered sources for additional energy calibrations beyond the ^{16}N and AmBe sources already developed and deployed.

Source	Energy (MeV)
^{46}Sc	0.9, 1.1
^{48}Sc	1.0, 1.2, 1.3
^{57}Co	0.122
^{137}Cs	0.7

comparison against a threshold voltage. ECALs measure the SNOD no-PMT baseline in ADC counts, and sets the threshold to 1 ACD count above this baseline to mitigate electronics noise. Similarly, ECALs also configure the QHL, QHS, and QHX baseline values (“pedestal”) and timing windows, as well as the GTVALID timing windows, through adjustment of digital-to-analogue converters (DACs) on the FECs. ECALS also verify that FEC voltages are running at expected values, the XL3 card is correctly translating data, and that all channels can count GTs correctly, output expected pedestal charges, have working N20 and N100 signals, and otherwise respond properly to a fake PMT pulse sent by the PMTIC.

After the electronics baselines have been measured, the first stage of turning PMT pulses into meaningful information is through another set of electronics calibrations (ECAs). The first ECA process sends a zero-point “pedestal” charge (PDST) to each channel to verify that they remain corrected from the ECALS, and records the pedestal value such that it can be appropriately removed during the ADC readout process. The second ECA process then uses zero-point charges with a delayed GT to check the timing slope (TSLP) of the TAC operation on each DB CMOS chip, thereby

measuring conversion from ADC to time. Both PDST and TSLP ECA processes are run weekly to account for baseline drifts and identify faulty channels.

In addition to their functionalities described in Section 2.4.2, the laserball or TELLIE LEDs can be used to perform PMT electronics calibrations (PCAs). By shining the light isotropically and tuning the intensity such that the PMTs only detect a single photoelectron (SPE) in each event, the charge associated with a single photon can be determined for each PMT.

Altogether, the multitude of complementary calibration systems are used to turn collections of PMT charges into meaningful, analysable data for use in physics analyses.

2.5 Secondary Physics Goals

Although the prime objective of the SNO+ experiment is the search for $0\nu\beta\beta$, the detector design, DAQ, and calibration systems allow for the pursuit of numerous high interest secondary physics goals in both the water and the scintillator phases. Furthermore, the addition of $^{\text{nat}}\text{Te}$ into the detector for the tellurium phase will still allow for continued measurements of secondary scintillator physics goals.

2.5.1 Water Phase Technique

While filled with water, the detector operated as a water Cherenkov detector, continuing a highly successful legacy championed by SNO and SK. As a Cherenkov detector, the primary detection technique was to measure the Cherenkov radiation emitted by a charged particle moving faster than the speed of light in the detector medium. In

an electromagnetic analogy to the shock wave of a sonic boom, the resulting emission is released as a cone of light.

The angle of this Cherenkov cone θ_C is [84],

$$\cos\theta_C = \frac{1}{n\beta}, \quad \beta = \frac{v}{c}, \quad (2.11)$$

where n is the refractive index and v is the velocity of the particle. In the case of relativistic particles ($\beta \sim 1$) in water ($n = 1.33$), this angle is $\sim 41^\circ$. From Equation 2.11, the minimum kinetic energy required for a particle of mass m to cause this Cherenkov emission can then be determined to be,

$$E_C = m \left(\frac{1}{\sqrt{1 - n^{-2}}} - 1 \right). \quad (2.12)$$

Particles below this kinetic energy threshold ($E_C = 265$ keV for electrons) are therefore undetectable within water Cherenkov detectors.

A strength of Cherenkov detectors is the capacity to recover the directionality of the incident particle, afforded by the shape and location of the cone. Within the spherical symmetry of the AV and PSUP, the isotropy of the event can also be used in particle identification. The number of Cherenkov photons N released can be estimated through the Frank-Tamm formula,

$$\frac{d^2N}{dx d\lambda} = \frac{2\pi\alpha z^2}{\lambda^2} \sin^2\theta_C, \quad (2.13)$$

where $\alpha = 1/137$ is the fine structure constant and z is the charge of the particle. As the energy of the propagating particle is linearly proportional to the number of photons produced, it can be estimated by counting the number of detected photons. The relative timing of the detection between each PMT also allows for the reconstruction of the time and location of the originating interaction.

2.5.2 Water Phase Results

Data taken during the water phase were differentiated into two main datasets used for analysis of secondary physics goals. The 115 live-day “initial” dataset started on 4 May 2017 after the completion of the water fill, and lasted until December 2017. In this initial period, the background level fluctuated, as the various detector systems were commissioned and tuned. The evolution of this background was studied and determined to broadly fall into six discrete time bins. Following the completion of the AV cover gas systems, a 190 live-day “low-background” dataset was taken between the period of October 2018 and the start of the scintillator fill in July 2019. Backgrounds in the initial dataset were measured to be at the level of $\mathcal{O}(10^{-14})$ g/g U or Th equivalent within the AV, and reduced to $\mathcal{O}(10^{-16} - 10^{-15})$ g/g U or Th equivalent in the low background dataset.

2.5.2.1 Invisible Nucleon Decay

Numerous experiments have searched for hypothesised BSM decays of n and p to final states where no visible energy is deposited, both in mono-nucleon and dinucleon (pp , np , nn) decay modes. During the water phase, this “Invisible Nucleon Decay” (IND)

was searched for in the case of spontaneous n decay in oxygen,



The relaxation of the daughter ${}^{15}\text{O}^*$ isotope would then release a detectable γ of $E = 6.18 \text{ MeV}$ (BR = 44%). Similarly, the hypothesised p , nn , np , and pp IND modes of ${}^{15}\text{O}$ were searched for in the relaxation of ${}^{15}\text{N}^*$ (6.32 MeV [BR = 41%]), ${}^{14}\text{O}^*$ (6.09 MeV [BR=10.9%] and 7.01 MeV [BR=20.1%]), ${}^{14}\text{N}^*$ (6.45 MeV [BR=7.7%] and 7.03 MeV [BR=8.9%]), and ${}^{14}\text{C}^*$ (numerous channels between 5–9 MeV with a total [BR = 4.53%]), respectively.

A search for these invisible nucleon decays was performed during the initial dataset, for which new world-leading lifetime limits were set in the p , np , and pp decay modes [152]. An improved search using both initial and low-background datasets improved on these measurements, and also set world-leading limits in the n decay mode [153]. Current best limits on mono- and di-nucleon IND decay modes are seen in Tab. 2.4, with decays of p , n , pp , and np set by SNO+; the nn IND half-life limit continues to be held by KamLAND [154].

2.5.2.2 Reactor Antineutrinos

Within SNO+, antineutrinos can interact with protons through the IBD interaction seen in Equation 1.9. Due to the large discrepancy between the e^+ and n masses, the resulting e^+ carries away a majority of the energy. The positron is annihilated within $\sim 1 \text{ ns}$ and releases two photons of 0.511 MeV, while the neutron thermalises and is typically captured by hydrogen at around $\sim 200 \mu\text{s}$, releasing a distinctive 2.2 MeV photon. The coincidence between the prompt energy deposited by the positron and

Table 2.4: Current best limits on the mono- and di-nucleon IND decay modes following the SNO+ result published in [153]. The limits on the p , n , pp , and np decay modes are set by SNO+, while the limit on nn mode is set by KamLAND.

Decay Mode	Current Best Limit $\times 10^{28}$ y	Experiment
p	96	SNO+ [153]
n	90	SNO+ [153]
pp	11	SNO+ [153]
np	6.0	SNO+ [153]
nn	140	KamLAND [154]

the distinctive 2.2 MeV energy deposited by the delayed neutron provides a method for identifying antineutrino interactions within the detector. Additionally, the AmBe source described in Section 2.4.3 provides a particularly formidable calibration tool due to the similarities between AmBe and IBD interactions.

SNO+ is coincidentally situated 240 km, 340 km and 350 km away from three nuclear reactors with a total of 18 cores. These three reactors provide 60% of IBD interactions within the AV, and their fluxes and positions relative to SNO+ are well known, potentially allowing for a measurement of Δm_{21}^2 and θ_{12} . Antineutrinos released from these reactors have been detected by SNO+ at a significance of 3.5σ [155]. This result is the first ever measurement of antineutrinos in water. However, higher rates afforded by a higher light yield would be required for a robust measurement of the oscillation parameters.

2.5.2.3 ^8B Solar Neutrinos

Using the initial dataset, the ^8B solar neutrino spectrum was measured down to 5 MeV. This analysis capitalises on the directionality provided by Cherenkov nature of detector by investigating $\cos\theta_{\text{sun}}$ of detected particles, where θ_{sun} is the angle between the direction of the detected particles and the position of the sun. The resulting measured flux^t of $\Phi = 2.53^{+0.31(\text{stat.})+0.13(\text{syst.})} \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$ is consistent with measurements performed by SNO and SK [156].




Figure 2.13: 5–15 MeV ^8B solar neutrino spectrum measured during the water phase using (left) the initial dataset as published in [156] and (right) the low background dataset. The low background dataset features lower backgrounds primarily due to the introduction of the AV cover gas system, which reduced the level of ^{222}Rn brought into the detector medium from the air. The left figure is adapted from [156] while the right figure is adapted from [157].

The background levels seen in Figure 2.13 (left) are dominated by a Rn background that has been removed from the low-background dataset following the installation of the AV cover gas system, as seen in Figure 2.13 (right). An improved

^tThis is the rate inferred in the case of no neutrino oscillations.

flux measurement utilising the low-background dataset is underway. A preliminary energy spectrum can be seen in Figure 2.14. The lowest energy bin fits anomalously low compared to Monte-Carlo (MC) simulated expectations; this unusual downturn in the low energy peak has also been observed by other solar neutrino experiments. Similar to the measurement of Δm_{21}^2 using reactor antineutrinos, repeating this analysis with a lower energy threshold of the detector is desirable to further explore this anomaly.

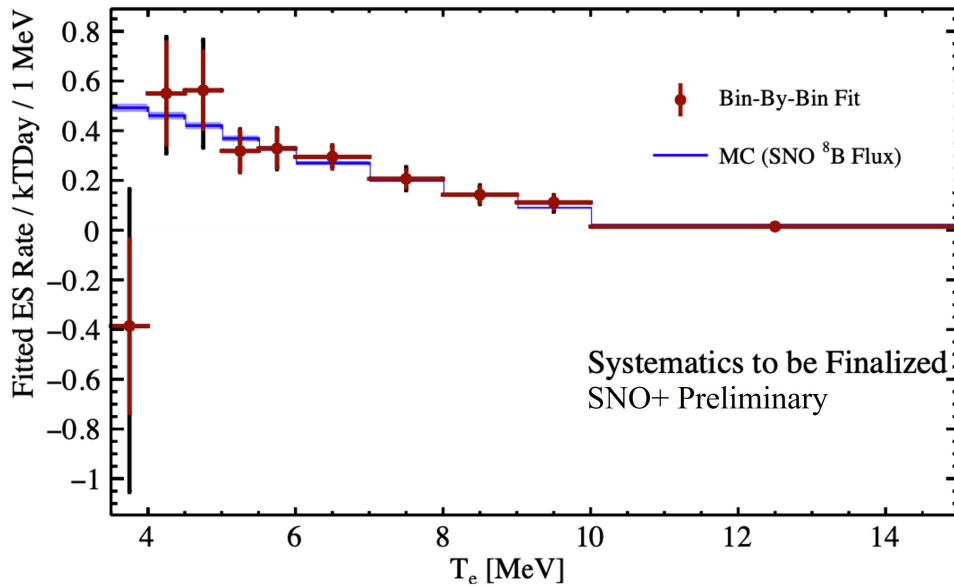


Figure 2.14: Preliminary bin-by-bin fit of the ^8B solar neutrinos as measured using the low background dataset, with variable bin widths from 3.5–15 MeV. Figure adapted from [157].

2.5.3 Scintillator Phase Technique

Replacing the water with liquid scintillator lowered the energy threshold of the detector. While necessary to perform a competitive measurement of $0\nu\beta\beta$, this also enables SNO+ access to a new lower-energy physics programme not possible while it was a water Cherenkov detector.

The scintillation process occurs when the molecules that make up the detector medium — the scintillator — are excited by ionising radiation from particle interactions that occur within the AV. The scintillator then luminesces as it de-excites, isotropically producing photons; this effect is discussed further in Section 4.1.1. Although the light output is much higher than that of water, the scintillation effect utterly dominates the scintillator Cherenkov light, thereby severely decreasing possibilities of recovering particle directionality.

For a particle depositing energy in the scintillator over a path length of dx , the empirical Birks' law provides the expected light yield,

$$\frac{dL}{dx} = S \frac{dE/dx}{1 + k_B(dE/dx)}, \quad (2.15)$$

where S is the scintillator efficiency in photons per energy and k_B is the material-dependent Birks' constant. The Birks' constant causes a non-linearity to the light yield, which is known as the quenching effect. The energy deposited (per dx) is dependent on both the property of both the incident particle and the scintillator, and is given by the Bethe-Bloch formula,

$$\frac{dE}{dx} = \frac{z^2 e^4}{8\pi\epsilon_0^2} \frac{1}{m_e} \frac{1}{v^2} \left(Z \frac{N_A}{A} \right) \ln \left(\frac{2m_e \beta^2}{E_{min}(1 - \beta^2)} \right) \quad \text{where} \quad \beta = \frac{v}{c}. \quad (2.16)$$

Here, ϵ_0 is the vacuum permittivity, N_A is the Avogadro number, m_e and e are the mass and charge of the electron, z and v are the charge and speed of the incident particle, and N and Z are the atomic and mass number of the scintillator, and E_{min} is the scintillator ionisation potential (minimum energy required to ionise one molecule).

The light yield is heavily dependent on the material chosen as the scintillator. Indeed, the choice of the scintillator is the defining experimental characteristic of a scintillation experiment, and the considerations behind the choice of a LAB-based liquid scintillator in SNO+ is elaborated further in Section 4.2.1.

2.5.4 Scintillator Phase Prospects

The SNO+ scintillator phase began on 29 April 2022 following a scintillator fill that lasted 33 months, including the 8-month partial fill period where the AV was filled with 46% LAB + 0.6 g/L PPO and 54% UPW. The data taken during this scintillator phase will be further split into two data-taking periods. The first dataset has been completed, with a scintillator composition of LAB + 2.2 g/L PPO. The second “enhanced” scintillator fill period will commence following the addition of 1,4-Bis(2-methylstyryl)benzene (Bis-MSB), a wavelength shifter to further enhance the light yield of the detector. The enhanced scintillator is discussed in Section 4.5.

2.5.4.1 Antineutrinos

As they were in water, antineutrinos are detected in scintillator with IBD reactions. With a greater light yield, the energy threshold is reduced to sub-MeV levels, allowing for a robust measurement of the reactor antineutrinos. As discussed in Section 2.5.2.2,

the source location and flux of a majority of these reactor antineutrinos are well known, thus providing the capability for a precise measurement of Δm_{21}^2 and θ_{12} .

As seen in Figure 2.15, the main background behind such a reactor antineutrino measurement are $^{13}\text{C}(\alpha, n)^{16}\text{O}$ interactions within the detector,^u



As in the case of IBD events, the n will be captured after $\sim 200 \mu\text{s}$ and provide the distinctive 2.2 MeV signal. However, in the case of alpha-n events, the n carries most of the kinetic energy;^v as it moves, the n causes scattering recoils off other particles (typically protons) before it is thermalised, providing a corresponding prompt signal. Thus, alpha-n interactions cause a coincidence event that is challenging to distinguish from that of IBD interactions. The alpha-n background is further discussed in Chapter 3.

Fortuitously, the scintillation time profile of proton recoils in alpha-n events are slower than that of the positron signal in IBD events. A pulse shape discriminator taking advantage of this difference has been developed, and was demonstrated in the partial fill phase to reject 69.46% of alpha-n events while only sacrificing 6.60% of IBD events [158]. As shown in Figure 2.16, employing this classifier will reduce the livetime required to make a world-leading measurement of Δm_{21}^2 from 3.2 to 2.6 years. Even earlier than this, a measurement of Δm_{21}^2 after 6 months or less of livetime can

^uThis interaction with a ^{13}C parent and ^{16}O daughter will hereafter be shortened to “alpha-n”.

^vRecall that the positron carries away nearly all of the kinetic energy in IBD events due to Conservation of Momentum.

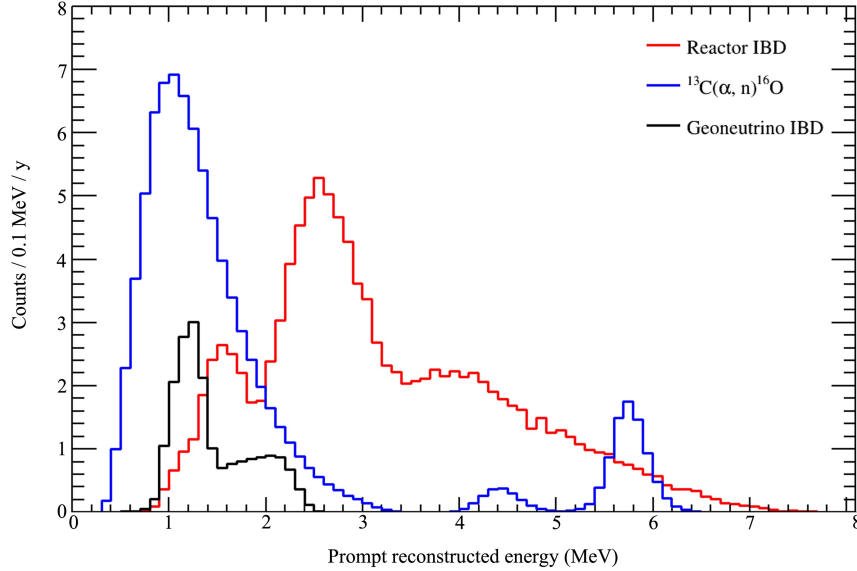




Figure 2.15: Simulated scintillator phase spectra for (red) reactor IBD events in the detector, (blue) the alpha-n background, and (black) geoneutrinos.

inform on the current tension between the solar and reactor experiments explained in Section 1.2.2.1 and seen in Figure 1.3.

As shown in Figure 2.15, electron antineutrinos from the crust and mantle of the Earth — collectively known as geoneutrinos — are also detectable. These geoneutrinos are primarily released through decays of ^{238}U and ^{232}Th ,



These decays are thought to be responsible for a large fraction (potentially over 50%) of the heat produced by the Earth [159]. A measurement of the geoneutrino flux by

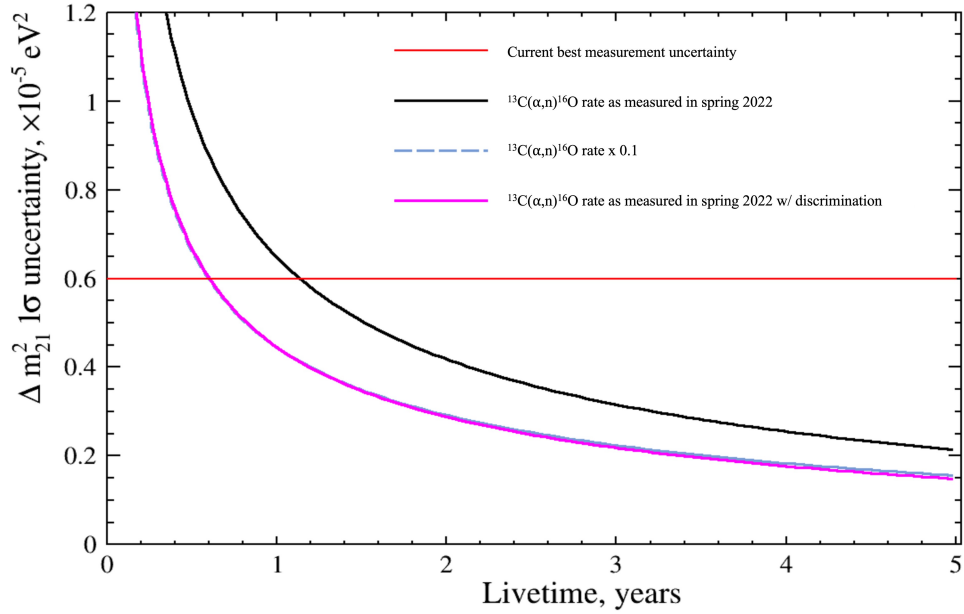


Figure 2.16: The sensitivity of SNO+ to Δm_{21}^2 with and without the classifier developed to discriminate alpha-n events. The alpha-n background rate has been reduced by an order of magnitude, reducing the livetime required to make a world-leading measurement from 3.2 to 2.6 years. Plot adapted from [158].

SNO+ would be the first in North America, complementing measurements performed in Japan by KamLAND [160] and Italy by Borexino [161]. Such a measurement would provide constraints towards a global model of the radiogenic heat flux of the Earth, thereby playing an important role in building a contemporary model of the interior of the Earth [162].

2.5.4.2 Solar Neutrinos

Using a scintillator target, the ^8B solar neutrinos can be measured down to lower energies than was accomplished in the water phase, thereby providing insight into the 3.5 MeV bin anomaly seen in Figure 2.14. Furthermore, lower energy neutrinos

from solar processes such as the CNO, pep , and ${}^7\text{Be}$ processes could be measured if the scintillator purity is sufficiently well controlled. While photons produced in solar processes may be trapped within the Sun for uncountable periods of time due to electromagnetic interactions with the Solar envelope, neutrinos would mostly shoot out of the solar core unscathed. Therefore, a measurement of these neutrinos would probe the inner workings of the Sun in a way prohibited in traditional electromagnetic techniques, and provide constraints towards Solar models.

While the higher energy ${}^8\text{B}$ solar neutrinos already measured have a P_{ee} dominated by the MSW effect, the P_{ee} of solar neutrinos below 1 MeV would be dominated by vacuum oscillations.^w Investigating the 1–5 MeV vacuum-matter transition region through measurements of solar neutrinos would not only enhance our understanding of the MSW mechanism, but provide a look into an energy region particularly sensitive to new physics due to the resonant nature of the MSW effect [130].

In the scintillator phase, the $\cos\theta_{\text{sun}}$ method used in the water phase is more challenging, as isotropic scintillation light dominates the Cherenkov signal. However, as scintillation light is released more slowly than Cherenkov light, they could be separated through an event-by-event comparison of prompt timing profiles. This reconstruction technique was successful when applied to partial fill data, as seen in Figure 2.17. This result is the first ever recovery of particle directionality in a scintillation detector on an event-by-event basis, and can be used to enhance solar neutrino analyses.

^wFurther discussion on vacuum/matter oscillations and P_{ee} (the electron neutrino survival probability) is found in Section 1.2.1.




Figure 2.17: Distribution of event trajectory relative to the position of the sun, showing a directional preference. Measured using partial fill data, this demonstrates the capability of retrieving directional information in a scintillation event. Figure adapted from [163].

2.5.4.3 Supernovae and Exotics

In the tumultuous twilight moments of a star with masses over the Chandrasekhar limit, overwhelming gravitational forces cause the core of the star to collapse in on itself until supported by neutron degeneracy pressure (in the case of a neutron star) or nothing at all (in the case of a black hole). In the resulting shockwave, 99% of the gravitational binding energy of the star is released as neutrinos [164]. While the photonic signal may take hours or days to escape from the dense plasma of the stellar remains, the neutrinos escape relatively easily, similar to the case of Solar neutrinos from the core of the Sun.

The neutrinos — released in quantities of $\mathcal{O}(10^{58})$ — would provide a wealth of information on the supernova mechanism. Furthermore, these neutrinos could be used to alert the astronomical community of the supernova event, enabling the potential to observe early electromagnetic signals from the rare occurrence. To this end, SNO+ will join a network of neutrino detectors that form the SuperNova Early Warning System (SNEWS) [165].

Beyond supernovae, the low backgrounds, large volume, and high light yield of SNO+ will also render the detector sensitive to a plethora of other exotic physics, from a measurement of the diffuse supernova neutrino background (DSNB) and axion-like particles (ALPs) to various classes of dark matter candidates [130].

Although encompassing a wide range of phenomena, all SNO+ physics goals are challenged by a common adversary: the detector backgrounds. Understanding potential backgrounds, deploying methods to suppress them, and quantifying what remains is the central objective behind the science behind the SNO+ experiment — especially in the search for $0\nu\beta\beta$.

Chapter 3

The Background Model

Even in the heart of our stronghold the Enemy has power to strike us: for his will it is that is at work.

– Gan. 5:7

Despite the utmost lengths undertaken to mitigate background events through the design of the detector, the infiltration of even the most formidable defences is expected from both internal and external sources.^a Regardless of their origins, these backgrounds pose a critical threat to the search for $0\nu\beta\beta$, as they contribute towards the background count in the ROI, thereby reducing the experimental sensitivity to $0\nu\beta\beta$ as shown in Eq. 1.67. As discussed in Sec. 2.2, suppressing these background events is the chief experimental challenge of SNO+.

^aInternal backgrounds are non-signal.(i.e. not $0\nu\beta\beta$) events that originate from within the AV, while external backgrounds are those that originate from outside the AV but propagate into it.

3.1 Expected ROI Backgrounds

An accurate accounting of all $0\nu\beta\beta$ ROI backgrounds cannot be ascertained until the final atom of ^{130}Te is deployed within the detector. Indeed, the irony behind developing a low-background experiment like SNO+ is that the only device capable of determining if the background rate is sufficiently low to achieve a competitive sensitivity is the detector itself. However, simulations^b based on existing decay measurements,^c studies of the water and partial fill phases, and radiopurity assays of SNO+ detector components allow for the construction of a model of expected backgrounds. These investigations have been conducted extensively by the SNO+ collaboration, with official predicted rates collected internally in [166]. The resultant expected number of background counts in the ROI for a one-year exposure with the initial Te-loaded scintillator are calculated in [167], with the relative share shown in Figure 3.1.

A simulation of the energy spectrum of these expected background sources can be seen in Figure 3.2. From this simulation, we can explicitly define the SNO+ ROI as,

$$-0.5\sigma < Q_{\beta\beta} < 1.5\sigma \quad (3.1)$$

which corresponds to,

$$[2.42 - 2.56] \text{ MeV}. \quad (3.2)$$

The uncertainty σ is the Poisson error on the expected nhits for the $0\nu\beta\beta$ energy peak of 2.46 MeV; this central value is lower than the ^{130}Te Q -value of $2.527.518 \pm 0.013$ keV

^bAll simulations performed in this chapter follow the same Geant4-based Monte Carlo technique described in Sec. 6.1.1.

^cAll half-life and Q -values used in this chapter were taken from NuDat3.0: <https://www.nndc.bnl.gov/nudat3/>

due to quenching. This asymmetric ROI would retain most of the $0\nu\beta\beta$ decays while minimising contributions from the $2\nu\beta\beta$ background discussed in Section 1.4.1 and 3.1.3.

Other unforeseen backgrounds may yet appear. Indeed, a significant advantage of the SNO+ $0\nu\beta\beta$ technique — wherein the liquid scintillator is deployed prior to loading with tellurium — is the capability to quantify the detector backgrounds without the $0\nu\beta\beta$ target. This “target out analysis,” discussed in Chapter 6, provides a quantification of the SNO+ $0\nu\beta\beta$ ROI backgrounds, clues towards the efficacy of the scintillator fill, and an understanding of the detector during and after the tellurium fill. However, expected backgrounds must first be discussed — only then can the target out analysis be used to root out those unexpected post factum.




Figure 3.1: The expected share of backgrounds in the $0\nu\beta\beta$ ROI during the tellurium phase. Figure adapted from [167].




Figure 3.2: Simulated energy distribution of expected backgrounds in the $0\nu\beta\beta$ ROI after 5 years of tellurium phase data-taking with 0.5% Te and a fiducial volume of 3.3 m. In red is the projected $0\nu\beta\beta$ signal in the case of $m_{\beta\beta} = 100$ meV while using the IBM-2 NME formalism ($M = 4.03$) and a phase space factor of $G = 3.688 \times 10^{-14} \text{ y}^{-1}$. Figure adapted from [167].

3.1.1 Cosmogenics

As discussed in Section 2.2.1, the deep underground location of the detector is what keeps the otherwise overwhelming rate of cosmic rays at bay. While the electromagnetic and hadronic particles of the shower are attenuated within metres of the crust, some of the highest energy muons are able to penetrate this shielding and interact with the detector. Although suppressed through this crustal shielding by a factor

of 5×10^7 to a rate of $3.31_{\pm 0.09(\text{sys})}^{\pm 0.01(\text{stat})} \times 10^{-10} \text{ cm}^{-2} \text{ s}^{-1}$ [129], the size of the detector implies that 68.9 ± 1.8 muons still interact each day (for a total detector rate $\sim 0.80 \text{ mHz}$) [168].

Unfortunately, as muons travel through matter — including the detector — they can cause the nuclear breakup of particles in a process known as spallation. This process is expected to occur at a rate of $\sim 1/\mu$ within the detector [169], and primarily occurs on ^{16}O in the cavity UPW and ^{12}C in the AV scintillator. Although numerous types of radioactive isotopes that are created from these interactions (“cosmogenically activated”) may pose risks to various secondary physics goals, none are both sufficiently energetic and long-lived to endanger the ROI.

Of greater concern is cosmogenic activation of the tellurium nuclei. Spallation products of tellurium are numerous [170], and those with half-lives of $t_{1/2} > 20$ days and $Q > 2.3 \text{ MeV}$ are summarised in Table 3.1. Unlike spallation products of UPW or the scintillator, simulations of these cosmogenic backgrounds shown in Figure 3.3 demonstrate that ^{22}Na , ^{44}Sc , ^{60}Co , ^{88}Y , ^{110}Ag , and ^{124}Sb have energy peaks which overlap with the ROI. Problematically, these backgrounds are sourced from the tellurium nuclei itself and thus scale with the amount of tellurium added.

Although the rate of cosmogenic activation while the tellurium is in the laboratory (or AV) is insignificant, these isotopes are sufficiently long lived such that irradiation while en route to the experiment can cause a lasting background. The simplest way to mitigate this is to bring the tellurium underground and wait for the cosmogenic products to decay away prior to the tellurium fill. For this purpose, the tellurium to be used in SNO+ was brought underground between 2015–2018, and has been “cooling off” for the greater part of a decade in anticipation for the tellurium fill.

Table 3.1: Half-lives and Q -values of isotopes with $t_{1/2} > 20$ days created as a result of spallation on tellurium nuclei. Only backgrounds with a $Q > 2.3$ MeV are considered. In orange are those with energy peaks that overlap with the ROI. These isotopes were identified in [170].

Isotope [Parent]	$t_{1/2}$ (days) [Parent]	Q -value (MeV)
^{22}Na	950.6	2.84
^{26}Al	2.62×10^8	4.00
^{42}K [^{42}Ar]	0.51 [1.2×10^4]	3.53
^{44}Sc [^{44}Ti]	0.17 [2.16×10^4]	3.65
^{46}Sc	83.79	2.37
^{56}Co	77.2	4.57
^{58}Co	70.9	2.31
^{60}Co [^{60}Fe]	1925.27 [5.48×10^8]	2.82
^{68}Ga [^{68}Ge]	4.7×10^{-2} [271]	2.92
^{82}Rb [^{82}Sr]	8.75×10^{-4} [25.35]	4.40
^{84}Rb	32.8	2.69
^{88}Y [^{88}Zr]	106.63[83.4]	3.62
^{102m}Rh	1366.77	2.32
^{102}Rh [^{102m}Rh]	207.3(1366.77)	2.46
^{106}Rh [^{106}Ru]	3.47×10^{-4} [371.8]	3.54
^{110m}Ag	249.83	3.01
^{110}Ag [^{110m}Ag]	2.85×10^{-4} [249.83]	2.89
^{124}Sb	60.2	2.90
^{126m}Sb [^{126}Sn]	0.01 [8.4×10^7]	3.69
^{126}Sb [^{126m}Sb]	12.35 [0.01]	3.67

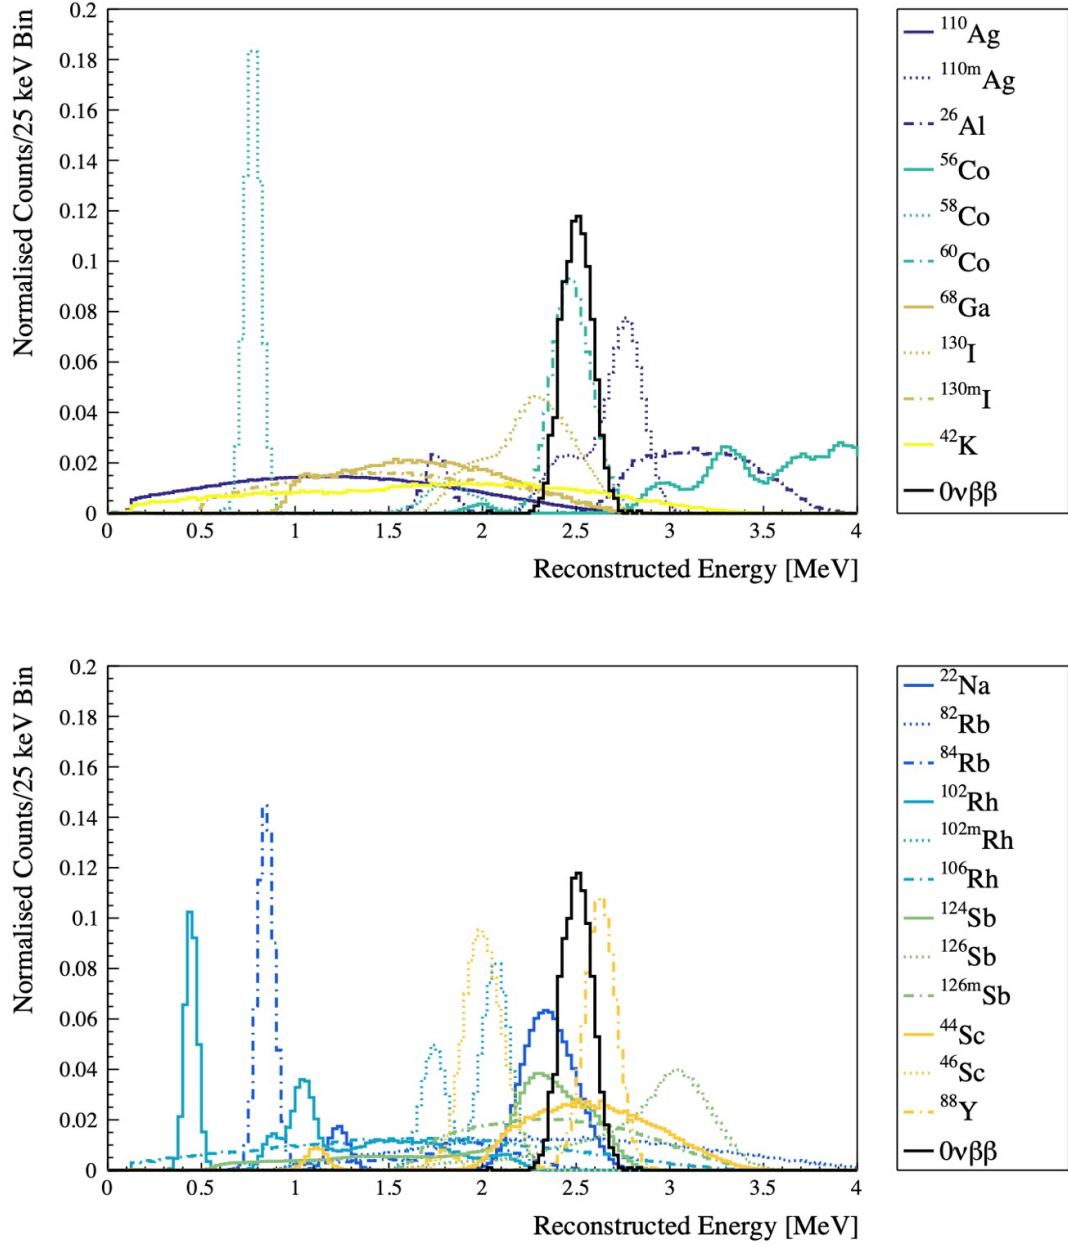


Figure 3.3: Simulated energy distributions of isotopes created through the cosmogenic activation of Te. The spectra have been normalised for comparison. A simulated $0\nu\beta\beta$ peak is shown in black. Figures adapted from [171].

For many of the isotopes, over 10 half-lives have already lapsed during their time underground, corresponding with a reduction in activity by a factor of 9.76×10^{-4} . These backgrounds can be further removed through the tellurium purification process, discussed in Section 5.4. Furthermore, studies show that analysis methods involving timing and PSD techniques are also capable of discriminating between $0\nu\beta\beta$ signal and cosmogenic background events [171].

3.1.2 Solar Neutrinos

Undeterred by the overburden that shields against cosmogenics, solar neutrinos penetrate deep into the detector and elastically scatter off the electrons within the liquid scintillator. Upon being struck by a neutrino, the recoiling particle ionises the scintillator, thereby causing detectable light. Both the ν_e and oscillated ν_μ are capable of scattering off of electrons to cause these ES events.

There are several solar neutrinos that could be detected using SNO+, with their fluxes shown in Figure 3.4. However, only the very same ${}^8\text{B}$ solar neutrino flux measured in the water phase (as discussed in Section 2.5.2.3) is expected to deposit this persistent background in the ROI. Although the *hep* solar neutrinos also exist in the ROI, their fluxes are comparatively negligible; a search using 2.47 kTy^d of SNO data resulted in a measured flux that was consistent with zero [172]. In a cruel twist of fate, the solar neutrinos which were once so desperately hunted for may now endanger SNO+.

^dkTy is kilotonne-years, a standard measure of detector data-taking time used within the neutrino physics community.




Figure 3.4: The expected solar flux of the neutrinos produced from the pp (solid) and CNO (dashed) fusion chains. Only ${}^8\text{B}$ and hep solar neutrinos overlap with the ROI, though the rates from the latter are negligible [172]. Figure adapted from [173].

This interaction rate is dependent on the cross-section between the solar neutrino and the recoiling electron [174],

$$\frac{d\sigma}{dT}(E_\nu, T_e) = \frac{2G_F^2 m_e}{\pi} \left[g_R^2 \left(1 - \frac{T_e}{E_\nu} \right)^2 - g_R g_L \frac{m_e T_e}{E_\nu^2} + g_L^2 \right], \quad (3.3)$$

where T_e is the kinetic energy of the recoil electron, σ is the interaction cross-section, G_F is the Fermi constant, and m_e is the electron mass. The flavour constants are defined as,

$$g_R = \sin^2 \theta_w \sim 0.23116 \quad (3.4)$$

$$g_L = g_R \pm 0.5, \quad (3.5)$$

where positive values correspond to electron neutrino scattering while negative values correspond to muon or tau neutrino scattering. The cross-section in Equation 3.3 does not account for radiative corrections, which applies an energy-dependant suppression of up to $< 2\%$ in the ROI [175]. For the purpose of estimating the background rate within this model, this radiative correction is not taken into account.

From the cross section, the rate can be determined by integrating

$$R = \Phi n \int S_\nu(E) [(P_{ee}(E)\sigma_e(E) + (1 - P_{ee}(E))\sigma_{\mu,\tau}(E)] dE \quad (3.6)$$

over the full ${}^8\text{B}$ energy spectrum. Here, n is the total number of electron targets within the AV, S_ν is the normalised energy spectrum of the neutrino, P_{ee} is the normalised electron neutrino survival probability, and $\sigma_{e,\mu,\tau}$ are the cross-sections for each neutrino flavour. Φ is the full-spectrum ${}^8\text{B}$ solar neutrino flux [176],

$$\Phi = 5.16_{-1.7\%}^{+2.5\%} \times 10^6 \text{ cm}^{-2}\text{s}^{-1}. \quad (3.7)$$

The rate can then be calculated through numerical integration of known S_ν and P_{ee} distributions after estimating n for a particular detector medium.

Based on the SNO+ scintillator composition of LAB + 2.2 g/L PPO (as discussed in Section 4.2.3) and the molecular composition as received from the suppliers prior to purification (summarised in Table 4.9), the SNO+ liquid scintillator is made up of $n = 3.357 \times 10^{29}$ electrons/tonne.

The full numerical integration provides an estimated rate of $R = 1061.25$ events/kTy, of which ν_e and ν_μ events contribute 75.4% and 24.6% of the rate, respectively. The fractional rate expected to occur within the ROI was determined through simulations.

This provides an expected ${}^8\text{B}$ rate of,

$$R = 31.2304 \pm 0.6075 \text{ counts/kTy} \quad (3.8)$$

within the ROI throughout the full volume of the AV.

In addition to this ES solar neutrino rate, neutrinos can undergo CC interactions on tellurium nuclei, creating ${}^{130}\text{I}$ in both ground and excited states [177]. The ground state β decays with $t_{1/2} = 0.515$ days and a Q -value of 2.94 MeV, while the metastable state decays with $t_{1/2} = 6.1 \times 10^{-3}$ days and a Q -value of 2.99 MeV. This contribution is included as part of the cosmogenics count, to differentiate it from solar neutrino ES events.

There is no feasible technique to prevent these solar neutrinos events from entering the ROI. However, if the event directionality can be recovered in the scintillator, a veto on events coming from the direction of the Sun could effectively aid in the rejection of solar neutrino ES signals in the ROI. Recovering directionality in scintillation detectors is traditionally challenging, as the isotropic scintillation light vastly outshines the directional Cherenkov component. As discussed in Section 2.5.4.2, it has been demonstrated from partial fill data that directionality can be extracted from scintillator phase events by taking advantage of the faster Cherenkov timing profile.

The PPO concentration during the partial fill phase was 0.6 g/L, far short of the 2.2 g/L deployed in the scintillator and tellurium phases. As discussed in Section 4.1.1, higher PPO concentrations will result in faster scintillator timing. Thus, the separation of scintillation and Cherenkov light will be more challenging than what has already been demonstrated. In lieu of this directionality-based signal extraction, the solar neutrino ES interaction flux will simply provide a sensitivity floor with which

to contend; SNO+ sensitivity estimates conservatively assume no suppression of the solar neutrino signal in the ROI.

3.1.3 The Double Beta Background

Although the deployment of tellurium is fundamentally required for the $0\nu\beta\beta$ search, doing so unfortunately adds a background: the $\Delta L = 0$ standard model $2\nu\beta\beta$ process. As discussed in Section 1.4.3, the $2\nu\beta\beta$ decay channel is produced by the isotope itself, rendering it an irreducible background.

In principle, the $0\nu\beta\beta$ signal should stand out as a peak at the Q -value of the decay. However, the overwhelmingly higher decay rate of the 2ν mode and limited energy resolution of the detector is expected to cause the upper tail in the $2\nu\beta\beta$ energy spectrum to leak into the ROI. In the case of ^{130}Te , the $2\nu\beta\beta$ half-life has been measured by CUORE to be [178],

$$\mathcal{T}_{1/2}^{2\nu} = 7.71_{-0.006}^{+0.008}(\text{stat.})_{-0.015}^{+0.012}(\text{syst.}) \times 10^{20} \text{ years}, \quad (3.9)$$

while the world-leading limit on the $0\nu\beta\beta$ half-life has been measured (also by CUORE) to be [99],

$$\mathcal{T}_{1/2}^{0\nu} > 2.2 \times 10^{25} \text{ years}. \quad (3.10)$$

Considering the initial 0.5% loading of $^{\text{nat}}\text{Te}$ in the tellurium phase, the $2\nu\beta\beta$ decay rate is expected to be 6.6×10^6 events/kTy. This irreducible background will only fall to the lower side of the ROI, and is the primary reason behind the asymmetry in the selected ROI.

3.1.4 Uranium

Trace amounts of ^{238}U are intrinsic in all Earthly materials, and those that make up the detector are no exception. With a long half-life of 4.5×10^9 years, ^{238}U has a rich decay chain containing daughters that have decay energies within the ROI, as shown in Figure 3.5. This is accentuated in the underground laboratory, as the surrounding norite rock — and thus the dust that fill the cavities within — have elevated levels of this radioactive isotope relative to the crustal average. Although significant efforts are undertaken to suppress ^{238}U through detector design (discussed in Chapter 2) and purification (discussed in Chapters 4 and 5), no mitigation strategy has a 100% reduction efficiency.

Through the ^{238}U decay chain, the decays of ^{210}Tl , ^{214}Bi , and ^{234m}Pa have sufficient energies to endanger the ROI. The α decay of ^{214}Bi to ^{210}Tl has a [BR = 0.021%], greatly suppressing that background. Meanwhile, the β decay of ^{234m}Pa has a Q -value of 2.12 MeV, which may leak into the ROI given a sufficient rate. This leaves the decay of ^{214}Bi as the isotope that poses the greatest threat. Thankfully, the decay is accompanied by a decay of ^{214}Po , with a $t_{1/2} = 164 \mu\text{s}$. Therefore, the ^{214}Bi β decays can be tagged by searching for events with the decay energy consistent with ^{214}Po event within a time frame of ~ 1 ms. This “ $^{214}\text{BiPo}$ ” coincidence provides a powerful tool for subjugating this otherwise insurmountable ROI background.

As the half-life of ^{238}U is over 1.8×10^4 times the half-life of the next longest lived isotope in the decay chain (^{234}U), the decay rate of ^{238}U constrains the decay rates of its multitudinous progenies in what is known as “secular equilibrium”. Thus, the rate of tagged ^{214}Bi can be used to assume a naïve g/g concentration of ^{238}U within

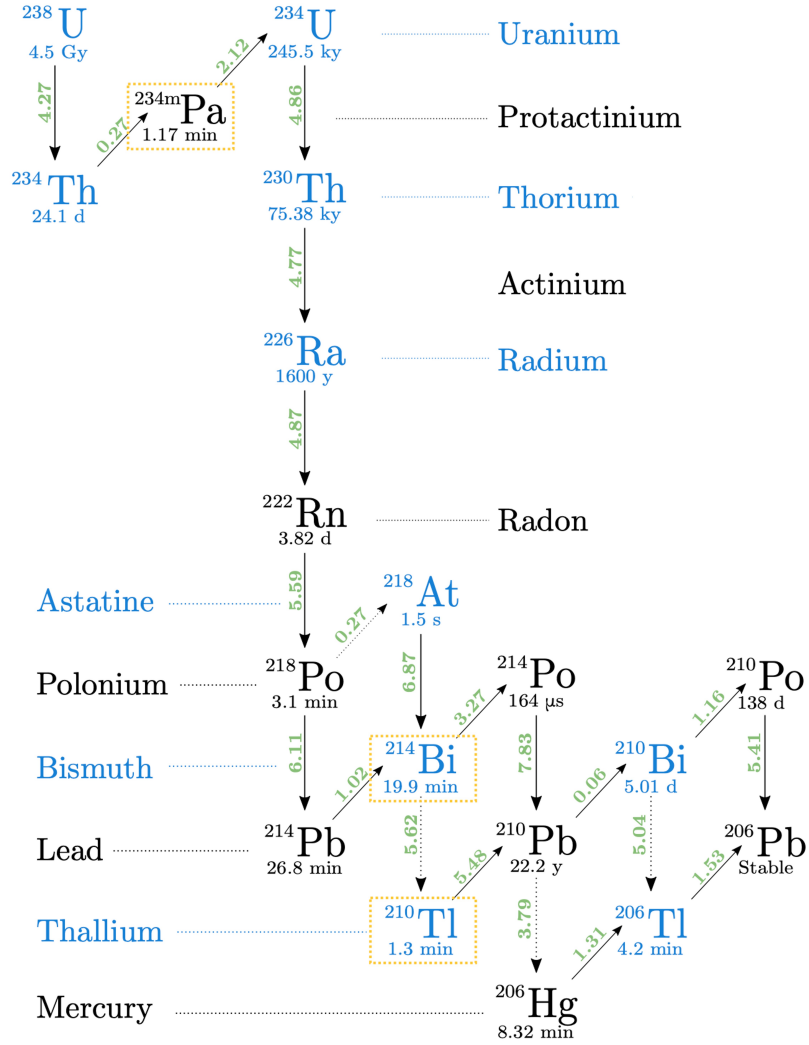


Figure 3.5: The radioactive decay chain of ^{238}U , along with the half-lives and Q -values (in green) of the associated isotopes. Vertical and diagonal arrows represent α and β decays, respectively. Orange boxes indicate isotopes with decay energies that may contribute background counts in the ROI.

the detector medium. However, secular equilibrium can be broken upon introduction of other long-lived isotopes within the ^{238}U chain such as ^{226}Ra .

3.1.4.1 Radon Ingress

One particular isotope in the ^{238}U chain that can break the secular equilibrium between ^{238}U and $^{214}\text{BiPo}$ is ^{222}Rn . Abundant in the laboratory air, the ambient decay rate of ^{222}Rn has been measured as 123Bq/m^3 within SNOLAB. [131]. Therefore, any contact with the laboratory air, including amounts dissolved into the detector media, will impart the detector with trace amounts of radon.^e This consequently breaks secular equilibrium and causes higher decays rates in isotopes further down the chain. Indeed, much of the detector design described in Chapter 2 — from aspects of the water purification systems to the scintillator phase URM — are specifically built to minimise air ingress for this reason. In particular, the installation of the AV cover gas system discussed in Section 2.2.5 so greatly reduced the ^{222}Rn rate in the detector that it ushered in a new low-background dataset during the water phase, as discussed in Section 2.5.2.

Although extreme lengths have been undertaken to protect the AV from radon ingress, any accidental exposures to air need not be an existential threat to the $0\nu\beta\beta$ search so long as the air leak is identified and eliminated. With a half-life of 3.82 days, the radon will decay back to rates sustained by secular equilibrium with ^{226}Ra within weeks. The typical method of determining the current radon concentration within

^eThe ^{222}Rn isotope is of such a great concern in low-background experiments that it is referred to associatively as “radon”. As is customary, both previous and subsequent uses of “radon” should be assumed to mean ^{222}Rn unless otherwise explicitly stated.

the detector is by tracking $^{214}\text{BiPo}$; the rates of these coincidence events should exponentially decrease with the radon decay half-life until it returns to the steady rate supported by secular equilibrium.

3.1.5 Thorium

Similar to ^{238}U , another long-lived isotope intrinsic to detector materials and abundant in the surrounding rock is ^{232}Th . Like its ^{238}U counterpart, this isotope also has a long decay chain, as seen in Figure 3.6. Within this chain are also three isotopes that imperil the ROI: ^{208}Tl , ^{212}Bi , and ^{228}Ac .

As with ^{234m}Pa from the ^{234}U chain, the decay of ^{228}Ac also falls short of the ROI with a Q -value of only 2.13 MeV; only the high energy tail of the decay may overlap with the signal given sufficiently high rates. And like the ^{214}Bi from the uranium chain, the decay of ^{212}Bi is also followed by a short-lived polonium isotope, in this case to ^{212}Po . However, with a half-life of only 300 ns, many of the ^{212}Po decays will occur within the same 400 ns event timing window as the prompt ^{212}Bi , making the follower isotopes too short-lived to be managed in the same “ $^{212}\text{BiPo}$ ” coincidence-based manner as $^{214}\text{BiPo}$. Therefore, “in-window” classification techniques that use charge timings are required to identify occurrences of $\beta - \alpha$ coincidences within single events.

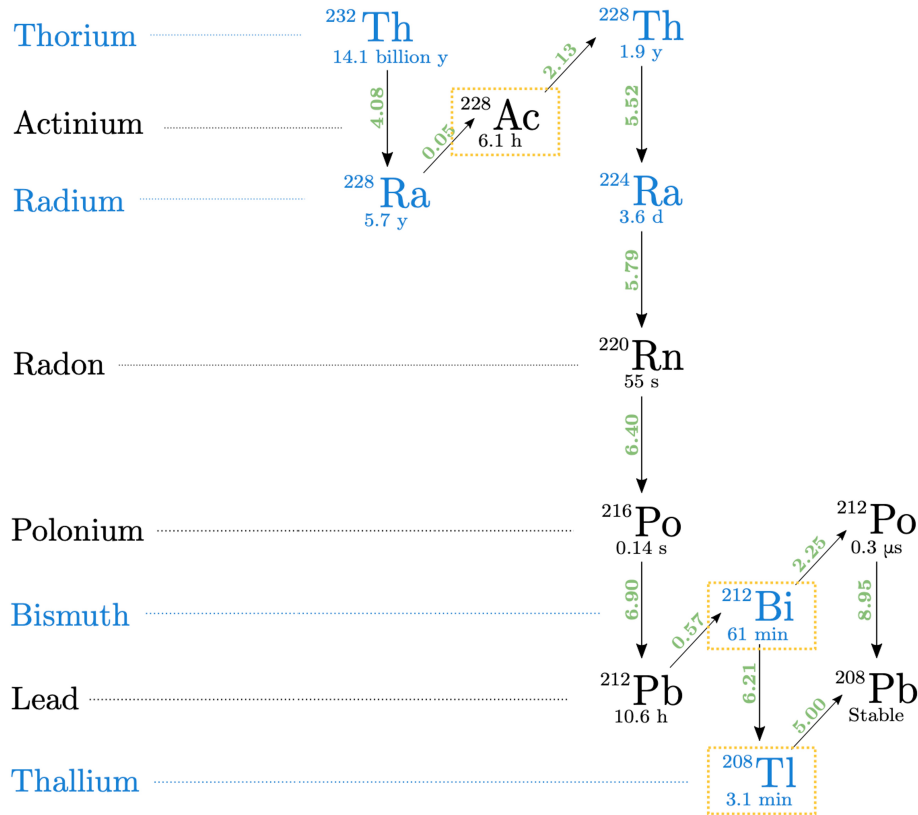


Figure 3.6: The radioactive decay chain of ^{232}Th , along with the half-lives and Q -values (in green) of the associated isotopes. Vertical and diagonal arrows represent α and β decays, respectively. Orange boxes indicate isotopes with decay energies that may contribute background counts in the ROI.

Also different from the uranium chain is the α decay of the bismuths. While the ^{214}Bi - ^{210}Tl (from the ^{238}U chain) decay has a minute branching ratio, the ^{212}Bi - ^{208}Tl (^{232}Th chain) decay has a [BR = 35.94%]. The ^{208}Tl itself β decays with γ s and a total Q -value of 5.00 MeV, thereby affecting the ROI. Thankfully, much of the resulting energy spectrum occupies an area above the ROI that is otherwise expected to be relatively background free. Along with $^{212}\text{BiPo}$ tagging, this consequently allows for ^{208}Tl to provide a method of estimating the ^{232}Th content within the detector.

3.1.6 Alpha-n Reactions

During the construction of SNO, large quantities of radon from the air had embedded onto the acrylic, eventually decaying to ^{210}Pb with a half-life of 22 years. The β s released from the resulting decays down the chain do not have sufficient energy to interfere with the ROI. However, the various α decays down the chain — primarily from ^{210}Po — can sustain a relatively high rate of alpha-n interactions on the ^{13}C in the scintillator and acrylic as described by Equation 2.17.

The embedded ^{210}Pb can also leach from the AV into the detector media. Bench-top tests using acrylic immersed in UPW and scintillator have determined a leach rate of $(1.05^{+0.18}_{-0.19} \times 10^{-3})/\text{day}$ in UPW at 18°C, corresponding to a rate of $\sim 0.40 \text{ Bq}/\text{m}^3$ after one year of exposure. This rate was measured to be an order of magnitude lower in scintillator. In both cases, there may be a possibility to remove these leached isotopes through recirculation and re-purification of the detector media as discussed in Section 4.3.2.

Regardless of the source, the α released by the ^{210}Po can be observed in the data, providing an estimate of the alpha-n rate. Although the ^{210}Po α carries 5.304 MeV

when released, quenching effects by the scintillator (discussed in Section 4.2.4.2) reduces the observed energy to 0.4 MeV [179]. A measurement performed during the partial fill phase provided an estimated ^{210}Po rate of 0.24 Bq/m³. Due to improvements in the scintillator deployment techniques later in the scintillator fill (discussed in Section 4.3.2.2), preliminary measurements have determined the rate in the scintillator phase to be 0.08 Bq/m³. Taking into account the $\alpha \rightarrow n$ yield of $(6.27 \pm 0.44) \times 10^{-8}$ [158], the predicted alpha-n rate in the scintillator phase is 110.2 events/kTy.

As can be seen in Figure 2.15, some leakage of alpha-n events into the ROI is expected. Furthermore, the 2.2 MeV γ that accompanies each neutron is precariously close to the ROI; high rates may again see their energy tails entangled within the signal region.

3.1.7 External Backgrounds

All materials used in the construction of the detector contain trace amounts of intrinsic bulk contamination, as well as surface contamination embedded during manufacturing, shipping, or installation. Thankfully, nearly all of these detector materials are installed externally to the detector medium, and the resulting externally produced α particles, β particles, and all but the most energetic external γ particles are attenuated by the external UPW by the time they reach the detector volume.

The only surviving γ s are expected to be from decays of ^{214}Bi (from ^{238}U contamination) and ^{208}Tl (from ^{232}Th contamination). Of particular danger to the ROI are the 2.6 MeV γ s that are released from ^{208}Tl decays from sources near the detector. As the hold-down ropes and inner cavity UPW press against the AV, it is expected

that contamination in those detector items will contribute the greatest share of these external backgrounds. The AV itself also has an intrinsic (“bulk”) contamination, with limits measured by SNO to be < 1.1 ppt in both ^{238}U and ^{232}Th [129].

3.1.7.1 Detector Fiducialisation

Even the most energetic of these γ s are sufficiently attenuated as they propagate deeper into the detector volume, as can be seen by Figure 3.7. Therefore, these external backgrounds can be strongly mitigated by rejecting events that originate from outer regions of the AV. This technique is called fiducialisation, and the remaining sub-volume of the detector that is kept for analysis is the “fiducial volume (FV)”. Nominally, a 3.3-m FV^f (corresponding to the inner 16.6% of the detector or 131 tonnes of liquid scintillator) is considered for the $0\nu\beta\beta$ ROI as a means to suppress these external backgrounds without excessively sacrificing detector mass.

3.2 ROI Background Model

Having surveyed the expected backgrounds mustered against the SNO+ ROI, the rate from each contribution has been extensively studied by the SNO+ collaboration [166, 167]. From these rates, a model estimating the number of background counts in the ROI can be constructed to predict the $0\nu\beta\beta$ sensitivity of the experiment.

^fExplicitly, an x -m FV means that only events that occur within a spherical volume of radius x centered on on the AV will be considered. This will hereafter be the definitional usage of fiducial volume (FV) unless otherwise specified.




Figure 3.7: The simulated number of ROI counts as a function of radius. The assumptions for the $0\nu\beta\beta$ signal are the same as they were in Figure 3.1 and Figure 3.2. The attenuation of external events as they penetrate deeper into the AV is well described by a simple single Compton scatter model.

A summary of the various backgrounds that endanger the ROI is listed in Table 3.2. Within this model, there is a total expected background count rate of,

$$b = 9.48 \text{ counts/year} \quad (3.11)$$

within the SNO+ ROI with a 3.3-m FV. This accounts for background rejection inefficiencies (such as in the tagging of ^{212}Bi) that may leave residual counts in the ROI.

Table 3.2: The predicted backgrounds in the SNO+ $0\nu\beta\beta$ ROI, and the number of counts they are expected to deposit within a fiducial volume of 3.3 m after one year exposure.

Background Source	SNO+ ROI Counts (FV = 3.3 m)	
	1 Year Exposure	5 Year Exposure
Cosmogenics	0.1081	0.5405
Solar Neutrinos	4.61152	23.0576
$2\nu\beta\beta$	1.2105	6.0525
^{238}U Chain	0.39726	1.9863
^{232}Th Chain	1.9204	9.6018
(α, n)	0.0224	0.1119
Externals	1.2108	6.0542
Total	9.4810	47.4049

From this background rate, the SNO+ $0\nu\beta\beta$ sensitivity can be determined with a counting analysis using Equation 1.67,

$$\mathcal{S}_{1/2}^{0\nu} = \ln(2)\epsilon \frac{1}{n_\sigma} \frac{xaN_A}{A} \sqrt{\frac{MT}{b\Delta_E}}. \quad (3.12)$$

Recalling that b is the number of background counts in events/year/keV/kg, the number of background counts in the ROI per year can instead be described as,

$$b_{\text{ROI}} = b\Delta_E M, \quad (3.13)$$

where b_{ROI} is the number of ROI counts. The efficiency of the detector is defined as,

$$\epsilon = f_{\text{FV}} f_{\text{ROI}}, \quad (3.14)$$

where $f_{\text{FV}} = \left(\frac{R_{\text{FV}}}{R_{\text{AV}}}\right)^3$ is the fraction of the detector used for the FV, and f_{ROI} is the fraction of $0\nu\beta\beta$ signal that falls within the ROI. Therefore, the sensitivity of the experiment given live time T and number of ROI background counts can be written as,

$$\mathcal{S}_{1/2}^{0\nu} = \ln(2) \left(\frac{R_{\text{FV}}}{R_{\text{AV}}}\right)^3 f_{\text{ROI}} \frac{1}{n_{\sigma}} \frac{x a N_A}{A} M \sqrt{\frac{T}{b_{\text{ROI}}}}. \quad (3.15)$$

The parameters that are constant within SNO+ are summarised in Table 3.3. Combined, they can be extracted and calculated to be,

$$\mathcal{C} = \ln(2) \frac{1}{R_{\text{AV}}^3} f_{\text{ROI}} \frac{1}{n_{\sigma}} \frac{x a N_A}{A} M = 7.8847 \times 10^{24} \text{ m}^{-3}. \quad (3.16)$$

Table 3.3: The parameters applicable for the SNO+ experiment in the calculation of the experimental $0\nu\beta\beta$ sensitivity.

Parameter	Description	Value
r_{AV}	AV Radius	6 m
f_{ROI}	Fraction of $0\nu\beta\beta$ signal in ROI ^g	0.6457
n_σ	Z-value for a 90% confidence level	1.645
x	Fraction of mass that is ^{nat} Te	0.005
a	Fractional abundance of ¹³⁰ Te	0.3408
A	Atomic number of Te	129.9 g/mol
N_A	Avogadro Constant	$6.022 \times 10^{23} \text{ mol}^{-1}$
M	Deployed mass	$792.4 \times 10^6 \text{ g}$

The SNO+ $0\nu\beta\beta$ sensitivity can thus be simplified to,

$$\mathcal{S}_{1/2}^{0\nu} > (7.8847 \times 10^{24}) R_{FV}^3 \sqrt{\frac{T}{b_{ROI}}}. \quad (3.17)$$

For the background model with an expected ROI count rate of $b_{ROI} = 9.48/\text{year}$ within $R_{FV} = 3.3 \text{ m}$, the projected SNO+ sensitivity is,

$$\mathcal{S}_{1/2}^{0\nu} > \begin{cases} 9.2029 \times 10^{25} \text{ years} & \text{after 1 year live time (90\% C.L.)}, \\ 2.0577 \times 10^{26} \text{ years} & \text{after 5 year live time (90\% C.L.)}. \end{cases} \quad (3.18)$$

In this model, the current world-leading sensitivity limit of $\mathcal{S}_{1/2}^{0\nu} > 2.2 \times 10^{25} \text{ years}$ [99] would be surpassed with 21 days of detector live-time.

^gCalculated through comparison to simulations.

3.3 The Phantom Menace

Even minuscule amounts of dust entering the AV could contain sufficient ^{238}U and ^{232}Th to a background rate that debilitates the $0\nu\beta\beta$ search. This is the motivating factor behind the stringent cleanliness conditions in all activities related to the detector. Moreover, it stresses the vital need to reach extreme purity levels in the detector media, as there may already be unforeseen sources of backgrounds that were deployed alongside the liquid scintillator.

The projected rates used within this background model are derived from measurements in the water and partial fill phases. As such, it does not give consideration to the most severe experimental challenge of all: the capability of maintaining extreme purity of the detector medium through to the end of the tellurium fill. Thus, as compelling as the predicted sensitivity in Equation 3.18 may look, the model merely concludes that the detector is likely capable of providing a competitive $0\nu\beta\beta$ measurement. The promise of the SNO+ experiment thus rests on the quality of the liquid scintillator through deployment during the scintillator fill and beyond.

Chapter 4

The Deployment of Liquid Scintillator

May it be a light to you in dark places, when all other lights go out.

– Gal. 2:8

The core identity of every particle physics experiment is the detection technology used, for which in SNO+ is scintillation counting. Scintillation counting techniques have been used since the dawn of particle physics, as scintillating ZnS crystals were used by E. Rutherford, H. Geiger, and E. Marsden to perform the famous “Rutherford gold foil” scattering experiment which discovered the atomic nucleus [180]. The modern scintillation counter in which a scintillator is coupled to a photodetector (e.g. PMTs) was invented for the Manhattan Project by S. Curran [181]. This rich legacy of scintillation counting is carried forth within the SNO+ experiment.

It is the SNO+ scintillator — specifically designed for the experiment — that gives the experiment its greatest advantages and enables a rich physics programme. It may also be the greatest peril in SNO+, as only diminutive amounts of contamination can undermine the entire scientific programme. It is in the deployment and understanding of this scintillator upon which all hopes of finding $0\nu\beta\beta$ rest.

4.1 Enlightening Scintillators

All scintillation counters utilise the same fundamental principle: the detection of light emitted from scintillators following excitation by ionising radiation. Broadly defined, a scintillator is a material that can quickly fluoresce this absorbed energy as light over $\mathcal{O}(\leq 10)$ ns, differentiating it from traditional phosphorescent materials which are left in a metastable state following excitation and emit light much later than the initiating event [182]. Charged particles that pass through scintillators can excite some of the atoms that comprise it, which then emit isotropic “scintillation light” through direct radiative emission or by passing vibrational energy onto other emitting components of the material [183].

4.1.1 Scintillation Mechanisms

Scintillation capabilities have been discovered in both inorganic and organic (carbon-based) forms. The former are typically crystal lattices doped with impurities. The scintillation process utilises the electronic band structure of the crystal, wherein electrons can be excited from the valence to the conduction band. The resulting electron-hole pair then propagates through the lattice structure until captured by an impurity center and de-excited; this relaxation produces the scintillation light. The most popular inorganic scintillator is undoubtedly the Tl-doped NaI crystal. Common in industry and astroparticle physics alike, NaI(Tl) crystals have been used most famously in the **Dark Matter NaI** (DAMA/NaI) experiment and its successor, the **Dark Matter Large Sodium Iodide Bulk for Rare processes** (DAMA/LIBRA) experiment^a [184].

^aThe myriad of opinions on the merits of this experiment will not be discussed here.




Figure 4.1: (Left) The standard structural formula for benzene, (middle) an illustration of the various σ bonds from sp^2 orbitals which keep the structure together, and (right) the bond length and angles of the resulting planar hexagon. Figure adapted from [185].

In contrast, organic scintillators are made of aromatic molecules^b and do not utilise electronic bands to enable their scintillation capabilities. Instead, it is the molecular structure of the benzene (C_6H_6) hydrocarbons — the heart of all organic scintillators — which give the materials their scintillating mechanism. The carbon atoms that form the scintillator molecule each individually have an electron configuration $1s^22s2p^3$, such that there is one s orbital and three p orbital valence electrons available in the binding. The s electron and two of the p electrons (sp^2) form axially symmetric σ covalent bonds at 120° , which bind the structure together into the distinctive planar cyclic shape of benzene; each carbon is attached to two other carbons and a hydrogen at an angle of 120° degrees, as seen in Figure 4.1.

The final p orbital valence electron is oriented perpendicular the plane. However, given the cyclic shape of the molecule, they overlap with the other p orbitals to form a delocalised π system above and below the plane of the molecule, as seen in Figure

^bAromatic molecules are those with cyclic structures.




Figure 4.2: An illustration of (left) various π bonds from the p orbital, which overlap with neighbouring p orbitals to form (right) the delocalised π system. Excitations of p electrons in the π system give rise to scintillation light. Figure adapted from [185].

4.2. The p electrons are thus able to move about this π structure, acting as free π electrons [185]. When the molecule is affected by a charged particle, it is these π electrons that are excited to provide the resultant scintillation mechanism.

From the free electron model for π electrons as seen in the Jablonski diagram in Figure 4.3, the π electron can be energised from the ground state (S_0) to excited singlet ($S_1, S_2, \dots S_x$) or triplet ($T_1, T_2, \dots T_x$) states. Each singlet or triplet states also contain their own vibrational sub-states (e.g. $S_{x0}, S_{x1}, \dots S_{xy}$). Excited $S_{(x>1)(y>0)}$ states immediately de-excite to the S_{10} state through heat or vibrational excitation without photon emission; the S_{10} state promptly ($\mathcal{O}(1)$ ns) relaxes to the ground state by emitting a fluorescent photon. Excited triplet states also immediately degrade to the T_{10} state, which is notably less energetic than the S_{10} state. As such, molecules in the S_{10} state can undergo intersystem crossing and relax to the T_{10} state as well.




Figure 4.3: A Jablonski diagram of free π electrons after excitation. Singlet and triplet states immediately vibrationally relax to S_{10} and T_{10} states without emission. S_{10} then emits fluorescent light. T_{10} emits delayed phosphorescent light unless it is excited to a S state through interaction with another T_{10} molecule. Inset boxes are electron spin configurations for excited singlet and triplet states. Figure adapted from [185].

These T_{10} states are metastable, as decays of T_{10} to the ground state S_0 requires a spin-flip; eventually ($\mathcal{O}(100) \mu\text{s}$), this will occur giving rise to phosphorescent light.

In the time before they decay to the ground state, two molecules with T_{10} may interact and produce a transient dimer;^c one molecule is excited to the S state such that the other can relax to the ground state. The re-excited molecule then relaxes while producing fluorescent light after a delay ($\mathcal{O}(10) \text{ ns}$).

Heavier ionising particles can fully excite available local S states, thereby causing a greater ratio of T to S excitations. This results in an overall slower timing profile

^cDimers are individual molecules held together through intermolecular forces.

and is the basis behind the indispensable capability to differentiate between heavier α (3.727 GeV rest mass) and lighter β (511 keV rest mass) interactions within the detector.

As there are numerous non-radiative sources of energy loss throughout the various processes, the emission spectrum of the scintillator will occur at longer wavelengths than the absorption spectrum; the change in the absorption and emission peaks ($\Delta\lambda$) is known as the “Stokes’ shift”. Furthermore, as scintillation light is affected by the interplay between singlet and triplet states, the quantum yield^d of organic scintillators is significantly less than in their inorganic counterparts. Nevertheless, they are the favoured medium in the modern era of colossal neutrino detectors due to their relatively easy large-scale fabrication.

4.1.2 Liquid Scintillators

While they are widely used in solid (plastic) forms, the primary advantage of organic scintillators is the capability for some to exist stably as a liquid under a wide range of conditions. This allows for the deployment of large bulk volumes in flexible vessel sizes, as well as the potential to dissolve radioactive isotopes^e for *in situ* counting.

There are multiple widely-used types of liquid scintillators, each mostly comprised of their defining “solvent”: the benzene-based aromatic molecule which provides the fundamental scintillating functionality of the material. As with all scintillators, the

^dIn this context, the quantum yield is defined as the fraction of deposited energy that is re-emitted as light.

^e ^{130}Te , for example.

solvent is fundamentally characterised by its range of excitation energies (corresponding with the absorption spectrum) and the wavelengths of radiated photons (corresponding to an emission spectrum). However, liquid scintillators can have insufficiently large Stokes' shifts, resulting in a self-absorption that suppresses the amount of emitted light as the absorption and emission spectra overlap.

To counteract this effect, the solvent is doped with a fluorophore solute (“fluor”), to which it can form a dipole-dipole bond when excited. As the π electrons in the solvent relax to the ground state, the resultant energy released may not be in the form of a photon. Instead, the energy can transfer to the fluor, exciting it in a non-radiative^f process called Förster resonance energy transfer (FRET) [183]. The fluor then relaxes, thereby emitting the light at a different wavelength; choosing the appropriate fluor is important to minimise the overlap between the solvent absorption and solute emission spectra, thereby maximising the light yield. These solvent-solute mixtures form a deployable “scintillator cocktail,” and can be supplemented by additional wavelength shifters to match the emission spectra with the wavelength range that confers maximal PMT quantum efficiency.

A variety of scintillator cocktails have been successfully deployed in multiple large-scale neutrino experiments such as Borexino, KamLAND, and Double Chooz, which used 1,2,4-trimethylbenzene (pseudocumene, PC) [186], PC diluted in dodecane [187], and phenyl-*o*-xylylene (PXE) [188] as their respective solvents and 2,5-diphenyloxazole (PPO) as the fluor. However, existing widely used liquid scintillators such as PC and PXE are uncompromisingly destructive to the SNO+ infrastructure, as they are incompatible with the acrylic that makes up the AV. As the AV is the

^fMeaning that there is no light emission.

central component of the detector, a new liquid scintillator needed to be developed for SNO+.

4.2 The SNO+ Scintillator

There were certain mixtures and dilutions of existing solvents that have been demonstrated to be acceptably compatible with acrylic, such as the mixture of 80% dodecane + 20% PC + 1.52 g/L PPO used in KamLAND. However, such scintillators — particularly those based on PC — were also infamous in their strong odours, volatility, toxicity, hazardous handling, and potential for environmental damage [189]. Coupled with the increasing complexity brought about by creating and purifying multicomponent solvents, there was a high desire from both SNO+ and the global community at large to develop a new scintillator that is easy to handle, relatively cheap, and compatible with acrylic.

4.2.1 Linear Alkylbenzene

In the search for a new solvent, six were selected following a survey of liquid aromatic compounds as potential candidates to be compared against PC, and summarised in Table 4.1. Although the chief factor in the initial evaluation was the compatibility with acrylic, it was also desirable for a density close to that of the external water (1 g/cm^3) to reduce buoyancy strain on the AV as it is immersed in the cavity UPW. A higher flash point was also coveted to improve safety and handling.

To test for acrylic compatibility, 80 acrylic cubes ($6.35 \text{ mm} \times 6.35 \text{ mm} \times 6.35 \text{ mm}$) were submerged in 175 mL of each liquid for two weeks. Incompatibilities would cause the acrylic to dissolve or be otherwise damaged, the latter of which would result in

Table 4.1: A list of solvents identified as potential candidates for a new SNO+ liquid scintillator, along with their densities and flash points.

Solvent		Density (g/cm ³)	Flash point (°C)
Linear Alkylbenzene	(LAB)	0.856	143
Dodecylbenzene	(DCB)	0.870	140
Phenylcyclohexane	(PCH)	0.950	98
Di-isopropylnaphthalene	(DIN)	0.960	140
Phenyl-o-oxylyetane	(PXE)	0.985	167
1-Methylnaphthalene	(1-MN)	1.020	82
Pseudocumene	(PC)	0.889	44

suspended acrylic particulates. The solvents were then measured every two days using a Photon Technology International (PTI) QuantaMasterTM fluorescence spectrometer. As the acrylic residue would likely be larger than the wavelengths emitted by the spectrometer, attenuation due to Mie scattering would be indicative of acrylic degradation. Indeed, PXE, DCB and PCH were deemed insufficiently compatible from this test.^g

Beyond the acrylic compatibility, the light yield^h of the scintillator — being the main advantage conferred by its use — was the next main factor for consideration. To investigate this, the scintillators were doped with 2 g/L PPO and excited with a 10 mCi ¹²⁵I source. The emission spectra were then measured using the PTI spectrometer. The measured emission intensities were integrated between 350–425 nm (the quantum efficiency of the PMTs peaks near 386 nm [147]) and compared to that

^gSimilar to PC, PXE is compatible when diluted down to $\sim 20\%$ with dodecane. However, like PC, there were no obvious advantages to justify the complexity.

^hThe light yield is the number of emitted photons for each absorbed unit of energy.

of PC, with results shown in Table 4.2. As can be seen, 1-MN was measured to have a poor light yield; although this is known to be improved with intense handling and purification [190], 1-MN was excluded from contention as it afforded no obvious advantages when compared to DIN or LAB even after applying extreme handling considerations.

Table 4.2: The measured light yield of SNO+ liquid scintillator solvent candidates, after doping with 2 g/L PPO and excited with a 10 mCi ^{125}I source. The reported measurement is the integrated emission intensity between 350–425 nm

Solvent		Light yield (PC = 1)	Acrylic Compatibility
Linear alkylbenzene	(LAB)	0.963	Acceptable
Dodecylbenzene	(DCB)	0.870	Poor
Phenylcyclohexane	(PCH)	0.304	Poor
Di-isopropylnaphthalene	(DIN)	1.17	Acceptable
Phenyl-o-oxylyetane	(PXE)	0.898	Poor
1-Methylnaphthalene	(1-MN)	0.00	Acceptable
Pseudocumene	(PC)	1	Poor

Of the remaining acrylic-compatible candidates — LAB and DIN — the latter afforded a superior light yield and a better density match with UPW. However, DIN was found to have attenuation lengthsⁱ of 3.3–4.4 m at 430 nm [191], notably less than the scale of the 12-m diameter AV. Although it would be an excellent solvent in smaller experiments, this would render DIN ineffective within SNO+. In contrast,

ⁱThe attenuation length, also known as the extinction length, is the distance the light travels within the scintillator before the original intensity is reduced by a factor of $1/e$. This is further discussed in Section 4.2.4.3.

LAB was found to have attenuation lengths of over 20 m, as discussed in Section 4.2.4.3.

LAB thus stood alone as the ultimate compromise between light yield, simplicity, and attenuation lengths. As for acrylic compatibility, LAB has since been found to have a negligible chemical effect on acrylic after years of exposure when tested in accordance with the ASTM *D543 Standard Practises for Evaluating the Resistance of Plastics to Chemical Reagents* [192,193]. Furthermore, LAB is safe to handle; the material has been certified for environmental release, and the Material Safety Data Sheet (MSDS) provided by CEPSA Química Bécancour (formerly Petresa Canada) notes that the only risk to human health is as a choking hazard.

4.2.2 2,5-diphenyloxazole

For an efficient transfer from the solvent to solute, the emission spectrum of the solvent and absorption spectrum of the solute must overlap. To achieve high light yields, high quantum efficiencies and large Stokes' shifts are also required. To challenge the popularity of PPO, alternative potential fluors were investigated, as summarised in Table 4.3.

Of the considered fluors, p-terphenyl, p-quinquephenyl, and p-sexiphenyl were deemed unsuitable due to their low solubility in LAB. Both remaining challengers — BPO and NPO — were determined to have slightly superior light output characteristics to PPO. However, the absorption peak of PPO matched more favourably with the emission spectrum of pure LAB as shown in Figure 4.4, which was measured using a fluorescence spectrometer with an excitation wavelength of 250 nm. Furthermore, low costs, wide availability, and a long legacy of successful use in the particle physics

Table 4.3: A list of fluors identified as potential candidates for a new SNO+ liquid scintillator. Measurements performed with fluors dissolved in cyclohexane. Data retrieved for PPO and BPO from [194], and others from AAT Bioquest: <https://www.aatbio.com/>

Solvent	Absorption Peak (nm)	Emmission Peak (nm)	Stokes' Shift (nm)	Quantum Efficiency (%)
2,5-Diphenyloxazole (PPO)	303	355	52	84.2 ± 4.2
2-(4-Biphenyl)-5-phenyloxazole (BPO)	320	384	64	91 ± 5
2-(1-Naphthyl)-5-phenyloxazole (NPO)	330	399	59	94
P-terphenyl	276	338	62	93
P-quinquephenyl	295	370	75	100
P-sexiphenyl	292	373	81	94

community had resulted in extensive existing literature on the chemical. Ultimately, these factors allowed for PPO to retain its long-standing status as the pre-eminent fluor, and chosen as the solute for the SNO+ scintillator. However, NPO and BPO (if the toxicity of the latter is ignored) may still be fluors of interest towards the design of a scintillator that maximises light yield above all else.




Figure 4.4: The emission spectrum of LAB after dilution to 2.5 mL LAB in 1 L cyclohexane to minimise self-absorption. The emission spectrum of the diluted LAB was measured using a fluorescence spectrometer with an excitation wavelength of 250 nm.

4.2.3 The SNO+ Scintillator Cocktail

The final SNO+ scintillator cocktail was determined to be linear alkylbenzene (LAB) as the solvent doped with 2 g/L of 2,5-diphenyloxazole^j (PPO) as the fluor. The chemical formulae of these molecules are shown in Figure 4.5. This 2 g/L target concentration of PPO was selected, as the light yield afforded by the fluor exhibits steep diminishing returns with higher concentrations in LAB as shown in Figure 4.6 [196], which is further offset by a self-absorption that increases with the PPO concentration.

^jBecause of a typographical error in [195], this is being erroneously propagated as diphenyloxazole. Sorry about that.




Figure 4.5: Structural formulae for the two primary components of the SNO+ scintillator cocktail. (a) LAB is a phenyl group attached to a carbon chain varying between 9–14 (> 95% 9–12) carbons in length. (b) PPO is the fluorophore of the scintillator cocktail, with O–C bonds illustrated in red, and N–C bonds illustrated in blue.

Since the introduction of LAB to the community following this search [197], it has since been adopted as the premier liquid scintillator across the field of particle astrophysics. LAB-based liquid scintillators have been successfully deployed in a number of neutrino and dark matter experiments such as Daya Bay, RENO, and COSINE-100, and will be used in future neutrino experiments including JUNO and SABRE. Following the publication of the SNO+ liquid scintillator in [195], an independent study that re-investigated various solvents and fluors has confirmed LAB as the preferred scintillator in large-scale detectors, both in terms of safety and performance [198].

4.2.4 Cocktail Performance

With the issue of compatibility between the scintillator and the SNO+ hardware put to rest, the attention can then be turned towards the performance of the material. Such investigations were necessary prior to deployment in the AV to improve the




Figure 4.6: The light yield of LAB + PPO depending on the PPO concentration, as measured by [196]. The measured light yields (red dots) are normalised to the extrapolated maximum (assuming no LAB self-absorption, as is the case in lower concentrations) and fit by a double exponential. The difference between the measurements and extrapolated maximum (green boxes) is well described by an exponential fit, suggesting an exponential increase in LAB \rightarrow PPO light transfer as the PPO concentration increases.

understanding of detector through simulations^k during and after the scintillator fill period. These preliminary measurements thus set the groundwork behind the understanding of the SNO+ detector and ultimately the physics that it can investigate.

The most important quantity — and the core motivation behind the deployment of scintillator at all — is the light yield. It is this which will deliver unto the experiment the physics which it seeks, through the gifting of both a lower energy threshold and higher energy resolution. As explained in Section 2.5.3, the light yield is described

^kThe simulations are further discussed in Section 6.1.1

by Birks' Law shown in Equation 2.15 and restated,

$$\frac{dL}{dx} = S \frac{dE/dx}{1 + k_B(dE/dx)}, \quad (4.1)$$

recalling also that dE/dx is described by the Bethe-Bloch equation shown in Equation 2.16.

4.2.4.1 Scintillation efficiency

The amount of light that the scintillator is intrinsically capable of producing is based on two factors: the FRET light transfer efficiency from LAB to PPO, and the quantum yield of the PPO. The former is discerned by comparing the scintillator cocktail decay time τ_{0x} to that of pure LAB τ_{LAB} [148],

$$f_{xy} = 1 - \frac{\tau_{0x}}{\tau_{\text{LAB}}}, \quad (4.2)$$

both of which were measured using a time-based laser-excited fluorimeter. For the 2 g/L scintillator cocktail, τ_s and τ_{LAB} were measured as 5.1 ± 0.1 ns and 22.7 ± 0.1 ns, respectively, for a non-radiative transfer efficiency of $78.2 \pm 1.5\%$. As summarised in Table 4.4, this transfer efficiency increased with the PPO concentration; this is expected, as the presence of more PPO allows for a greater interaction probability between the solvent and solute molecules. Since greater transfer efficiencies are associated with shorter decay times (as can be seen in Equation 4.2), higher PPO concentrations increase the challenge of separating the prompt Cherenkov from slower scintillator light components, as discussed in Section 2.5.4.2 and Section 3.1.2.

Table 4.4: The measured FRET transfer efficiencies between LAB and PPO at various concentrations

PPO Concentration (g/L)	FRET efficiency
0.25	48.7 ± 5.0
0.5	59.3 ± 3.2
1	67.7 ± 2.3
2	78.2 ± 1.5
4	86.0 ± 0.8

Meanwhile, the quantum yield of PPO has been extensively studied by the community, with measured values ranging from 0.71–1.00; a relatively recent measurement of 0.842 ± 0.042 [194] is commonly used. Multiplying the quantum yield with the FRET efficiency provides the LAB–PPO re-emission probability¹ of $68.5 \pm 5.3\%$. As the LAB quantum efficiency is low ($20 \pm 2\%$ [194]), and much of the emitted light is re-absorbed, the contribution of LAB–LAB re-emission to the light output is negligible.

4.2.4.2 Absolute light yield

The light yield was initially measured through the same ^{125}I irradiation technique used to determine the light yields of the candidate solvents of Table 4.2, in which the emission spectrum as measured by the PTI spectrometer was integrated between 350–425 nm and compared to a sample of the well-studied Borexino (PC + 1.5 g/L PPO) liquid scintillator. This yielded an absolute SNO+ scintillator cocktail light yield of 11900 ± 1100 photons/MeV.

¹The probability that a photon absorbed by LAB will be re-emitted by PPO

This measurement was supplemented with another performed through use of a $0.1 \mu\text{Ci}$ ^{90}Sr source embedded into an acrylic cube. ^{90}Sr β decays to ^{90}Y with Q -value of 0.546 MeV β^- , the latter of which is observed with a trigger PMT. The ^{90}Y subsequently decays to ^{90}Zr with a Q -value of 2.28 MeV , which penetrates the acrylic and irradiates the scintillator and is observed with three PMTs observing the sample itself. This methodology yielded a light yield measurement of 10830 ± 570 photons/MeV. This measurement was improved to 11920 ± 630 photons/MeV, after the scintillator cocktail was sparged with N_2 gas to remove oxygen, in excellent agreement with the ^{125}I measurement.

The improved light yield following removal of oxygen was an expected outcome, as oxygen contamination was previously observed to contribute towards quenching in LAB [199]. Quenching reduces the effective light yield of the scintillator and thus decreases the performance of the entire detector, and is driven through FRET by absorbing contaminants such as oxygen which non-radiatively de-excite.

Even in lieu of such contaminants, quenching will nevertheless occur through interactions between excited molecules as they propagate. This “ionisation” quenching effect accounts for non-radiative energy losses and is included within Birks’ law, seen as dE/dx within the denominator of Equation 2.15, and implying that heavier ionisation particles are more severely quenched. The effect is parameterised by Birks’ constant k_B , and critical towards modelling light production in the detector. As such, the Birks’ constant has been measured extensively for α , β , and protons traversing through LAB, as summarised in Table 4.5.

Table 4.5: Measured Birks' constants for α , β , and p propagating through LAB.

Particle type	Birks' Constant ($\mu\text{m}/\text{MeV}$)	
α	$[66 \pm 16, 76 \pm 3]$	[179, 200]
β	~ 74	[201]
p	$[70.8 \pm 4.5, 98 \pm 3]$	[179, 202, 203]

4.2.4.3 Attenuation

Ionisation quenching is unfortunately not the only source of light yield reduction that must be accounted for. In the absence of impurities, the LAB molecules become a limiting factor, as they cause a scattering effect on the scintillation light.^m This in turn causes diffuse interreflection of the photons propagating through the medium, thereby attenuating the luminescence that is detected. This reduction in light is parameterised by an attenuation length which corresponds to the distance required for the light to be reduced by a factor of $1/e$.

As the wavelengths of light are larger than the size of the molecules,ⁿ the attenuation effect is dominated by Rayleigh scattering. As such, the attenuation length can be deduced by first determining the Rayleigh ratio of LAB. This ratio can be determined by comparing the intensity of light scattered at a particular angle in LAB against that of other solvents with known Rayleigh ratios,

$$R_{\text{LAB}} = R_s \frac{I_{\text{LAB}}}{I_s} \frac{n_{\text{LAB}}}{n_s} A. \quad (4.3)$$

^mPPO molecules would have a negligible scattering contribution in comparison effect, as they are substantially less abundant ($\sim 0.2\%$).

ⁿThe emission peak of PPO is 355 nm, while the σ bond length between carbon in benzene is 140 pm.

Here, R is the Rayleigh ratio, I is measured intensity, and n is the refractive index for LAB and the solvent s . A is a correction factor implemented to account for discrepancies in refractive indices between the solvent and quartz sample cuvette used for the measurement.

As summarised in Table 4.6, the intensity of LAB was compared to that of UPW, Acetone, Cyclohexane, and Toluene when illuminated by using a 546-nm monochromatised Xe lamp. Between the four solvents, the average LAB Rayleigh ratio was $(16.60 \pm 3.14) \times 10^{-6} \text{ cm}^{-1}$, corresponding to a mean attenuation length of $71.90 \pm 13.6 \text{ m}$ at 546 nm. Further details on this measurement can be found in [204].

The strong wavelength dependence of Rayleigh scattering ($R \propto \lambda^4$) should be noted, as shorter wavelengths will result in suppressed attenuation lengths. The Rayleigh scattering has been externally measured at $22.9 \pm 0.3(\text{stat.}) \pm 1.7(\text{sys.}) \text{ m}$ at 405 nm [205]; these attenuation lengths are still compatible with the scale of the detector, as the diameter of the AV (and maximum travel distance for photons in LAB) is 12 m.

This scattering is strongly dependant on the purity of the scintillator — even minor impurities will shift the dominant scattering mechanism from the Rayleigh to Mie paradigm, further reducing the attenuation length. Along with the effects of contaminants on quenching, this implies that ensuring extreme purity of the scintillator is paramount to allow for maximum scintillation light output. Furthermore, any contaminants would also carry the risk of introducing radiogenic backgrounds discussed in Chapter 3. Therefore, the primary experimental challenge is to ensure that the scintillator quality is maintained while being deployed into the detector. The fate of the entire experiment rests upon success of the scintillator fill.

Table 4.6: Using the Rayleigh ratio of known solvents, the Rayleigh ratio of LAB was determined by comparing the intensity of scattered light between LAB and each solvent. The correction factor accounts for discrepancies in refractive indices between the solvent and quartz cuvette used for the measurements.

Solvent	Rayleigh ratio (cm^{-1})	Correction Factor	Corresponding Rayleigh ratio of LAB (cm^{-1})
Ultrapure Water	1.05	21	$(16.63 \pm 2.99) \times 10^{-6}$
Acetone	4.47	5	$(16.79 \pm 3.14) \times 10^{-6}$
Cyclohexane	4.68	4	$(16.46 \pm 3.00) \times 10^{-6}$
Toluene	20.5	1	$(16.61 \pm 3.45) \times 10^{-6}$

4.3 The Scintillator Fill

The replacement of 904 tonnes of UPW with the novel liquid scintillator underground was an undertaking as daunting as it was arduous. That even a gram of contamination from the surrounding environment^o would render the entire project futile was an existential danger not lost on those who faced this challenge.

Yet *audentes fortuna iuvat*, for not only is LAB widely available due to its industrial uses,^p but a major production plant for the chemical — the only in Canada — is located only hours away. The SNO+ LAB is thus sourced locally, purchased from CEPSA Química Bécancour (formerly Petesa Canada) as high quality PETRELAB[®] 500-Q LAB from their LAB manufacturing plant in Bécancour, Québec, and shipped to SNOLAB in sets of 20 tonne tanker trucks.

^oThe ambient mine dust contains $1.11 \pm 0.13 \mu\text{g/g}$ U and $5.56 \pm 0.52 \text{g/g}$ Th [136] — a billion times the maximum concentration required for a competitive $0\nu\beta\beta$ sensitivity [130].

^pLAB is the most important raw material used in the manufacture of detergents, having outclassed all alternatives as it is both relatively cheap and biodegradable [206].




Figure 4.7: Preliminary density measurements in the scintillator cocktail, taken in one degree intervals between 10–26°C. Samples were cooled in a refrigerator and measured as it warmed back to ambient temperature. Instrumental uncertainties on the densiometer are represented by the shaded regions.

4.3.1 LAB Deployment Strategy

The relatively short distances and lack of international borders between the CEPESA LAB plant and the experiment also provided flexibility when transporting the SNO+ LAB to SNOLAB, a premium in the logistics of any large operation. Based on preliminary measurements of LAB density, as seen in Figure 4.7, a projected 780 tonnes of LAB would be needed to fill the AV, since a temperature of 12°C is desired to suppress creep in the hold-down ropes and bacteria growth in the external water. It would be infeasible to store the entire volume of LAB required underground prior to

deployment. As such, shipments from CEPSA had to be coordinated with the rest of the operation for the duration of the scintillator fill. A summary of this operation can be seen in Figure 4.8.

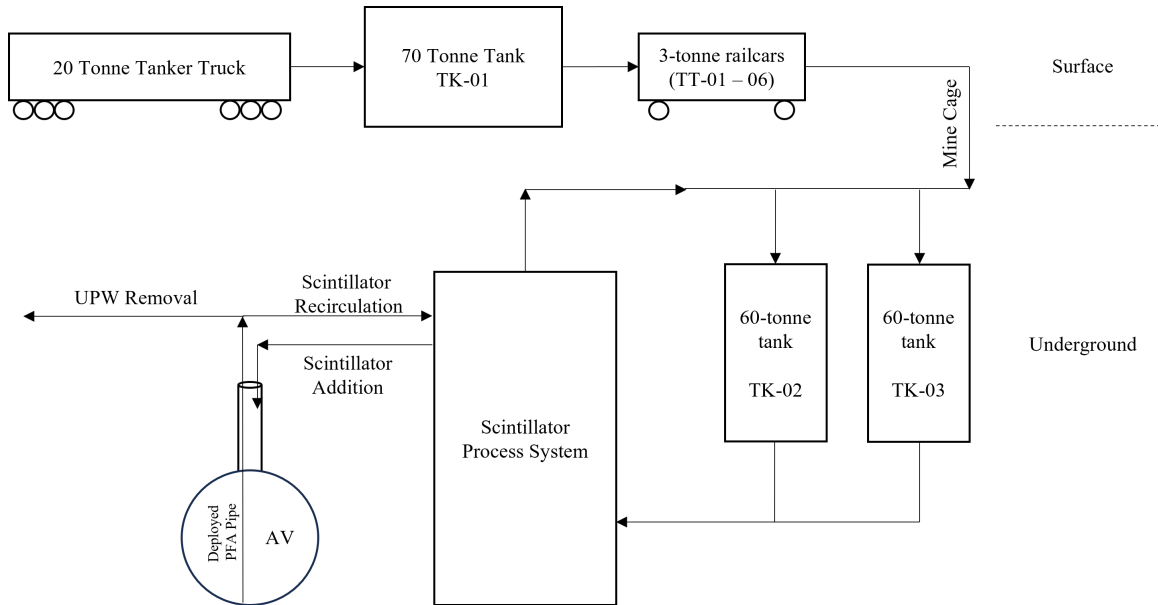


Figure 4.8: A simplified block flow diagram summarising the LAB deployment strategy.

A tanker transfer facility was set up at SNOLAB, where a swinging overhead crane could be rapidly deployed over incoming tanker trucks. The trucks were then connected to a hose via a Hansen quick-release coupling, and the volume pumped into TK-01, a 70-tonne vessel in the SNOLAB surface building.⁹ TK-01 was then unloaded into 6 3-tonne railcar tankers (TT-01 – TT-06), and continuously trammed back and forth from the underground laboratory to offload their scintillating cargo.

⁹Unless otherwise specified, vessel volumes are stated hereafter as their rated water carrying capacity. Thus, a 70-tonne vessel is only capable of holding ~ 60 tonnes of LAB after accounting for density differences.

Within the laboratory, the railcars were processed in the “dirty-side car-wash,” the second room in a four-stage compartmentalised processing area, where all equipment entering the laboratory is hand-cleaned to SNOLAB standards and inspected prior to admittance into the facility. An underground transfer facility was installed at this location in a manner similar to the surface transfer facility, with a swinging overhead crane to allow for the railcars to be connected to a hose. The LAB was then offloaded from the railcars to TK-02 and TK-03, two identical 60-tonne tanks. The railcars were then be raced back to surface, desperate to stay ahead of the next impending tanker truck.

The LAB within the 60-tonne tanks was then purified and sent into the AV. To protect the AV from both radon in the air and external pressure from the cavity UPW, the LAB was loaded into the top of the AV while the UPW from the bottom and sent to waste.^f This allowed for the detector to be constantly full of fluid while the filling was under way. Throughout the duration of the fill, the hold-up ropes had to be repeatedly adjusted to account for the changing buoyancy of the AV.

The removal of UPW was planned to occur through the same lines that were used to recirculate the AV volume during the water phase. However, the lines were determined to have a large leak rate; even early in the fill, accidental removal of scintillator constrained the net addition of scintillator to 20% nominal flow rate. To counteract this, a 2.5-cm diameter perfluoroalkoxy alkane (PFA) line with a 1.6-mm wall thickness was deployed through a calibration gate-valve on the UI, which was then used for UPW removal.

^fThis mechanism is possible because LAB is both non-polar and less dense than UPW.

The first drop of scintillator was added to the AV on 25 October 2018. However, the water phase would continue until the scintillator volume reached the AV neck boss, which it did on 19 July 2019. The addition of LAB from the 60-tonne tanks to the AV was facilitated by the SNO+ scintillator process system. Within this system was among the most critical of all operations: the purification of LAB.

4.3.2 The Scintillator Purification Plant

In the journey from the CEPSA tanker truck to the 60-tonne tanks, all possible precautions were undertaken to minimise potential contamination. Every vessel — including the railcars — were pre-cleaned through citric acid passivation, and the airspace within kept constantly pressurised with grade N5.0 (99.999% purity)^s N₂ gas. This cover gas minimises air contamination which could carry oxygen, radon, and dust containing trace radioactivity from entering the vessel volume. TK-01, TK-02, and TK-03 also each contained a polypropylene liner, and were cleaned using both a 1% Alconox[®] detergent solution and UPW to reach cleanliness levels in accordance with MIL-STD-1257C^t level 50.

The kilometers of tightly packed lines and vessels that make up the process and purification systems that guide the LAB from the 60-tonne tanks to the AV were designed specifically for maximal cleanliness and vacuum-tight requirements. Each of the hundreds of valves, fittings, and couplings of all types were individually leak-checked to have a leak tightness of $\mathcal{O}(10^{-11} - 10^{-7})$ mbar·L/s. Every metallic surface that would come into potential contact with the scintillator was constructed with

^sThis was the only purity grade used for of N₂ gas, and thus will hereafter not be specified.

^tMIL-STD-1257C is the United States Military Standard on Product Cleanliness Levels and Contamination Control Programme.

electropolished 316L stainless steel, rendering it chemically inert to organic solvents such as LAB.

Nevertheless, traversing even the purest lines into the cleanest vessels will not remove the contamination potentially already within the LAB itself. As discussed in Chapter 3, as even the most minor environmental exposures experienced by LAB from the moment of manufacturing — and from the manufacturing process itself — may have contaminated the fluid with ^{238}U , ^{232}Th , and ^{40}K from ambient dust, as well as Rn, Ar, Kr, and O_2 from dissolved ambient gasses.

From the multitudinous channels of potential contamination ingress, it is incontrovertible that a purification system was necessary. The process most effective at removing such varied contaminants to the excruciatingly low levels required is distillation [207]. This distillation process thus forms the heart of the scintillator purification plant, upon which the quality of the scintillator — and the scientific prospects of the SNO+ experiment — relies upon.

4.3.2.1 Primary Distillation Systems

Although it was distillation that spearheaded the cleansing of the LAB, there were two additional techniques built into the plant to rid the scintillator of impurities: solvent-solvent “water extraction,” and steam/nitrogen gas “stripping”. Supplemented by systems of filters of varying sizes down to 50 nm, these three purification techniques were highly effective at removing a large range of potential contaminants, as summarised in Table 4.7. Furthermore, the capability to remove scintillator from the AV and recirculate it within the plant for re-purification was also built in.

Table 4.7: Contaminants removed from the three main purification mechanisms in the scintillator purification plant.

Contaminant Type	Distillation	Water Extraction	Gas stripping
Heavy metals	Bi, K, Pb, Po, Ra, Th, U	K, Pb, Ra, Th, U	—
Dissolved gasses	—	—	Ar, Kr, O ₂ , Rn
Oxidised Organics	Carboxyl groups, 1,4-benzoquinone	—	—
Volatile Liquids	—	—	Residual water

A simplified block flow diagram of the scintillator plant can be seen in Figure 4.9. The scintillator was drawn from the 60 tonne tanks, and could pass through any of the three purification columns: C-100 (distillation), C-200 (water extractions), and C-300 (stripping). The scintillator could bypass any of the columns at any time, and be sent back into the 60 tonne tanks, forming the “primary” distillation loop.

The scintillator could also be sent into the V-01: the AV head tank, and the cleanest part of the scintillator process systems. From there, the LAB could either be sent back into the purification columns; this method of continuously cleaning just the fluid within the process systems without involving the 60-tonne tanks is known as “closed-loop recirculation”. It is also from V-01 that the scintillator could be sent into the AV itself.

Distillation The scintillator plant keystone is C-100, the 32” (0.81-m) diameter and 13’7” (4.14-m) tall column that performs the distillation process. An ancient and well-understood technique of separating fluids through heating, distillation remains the

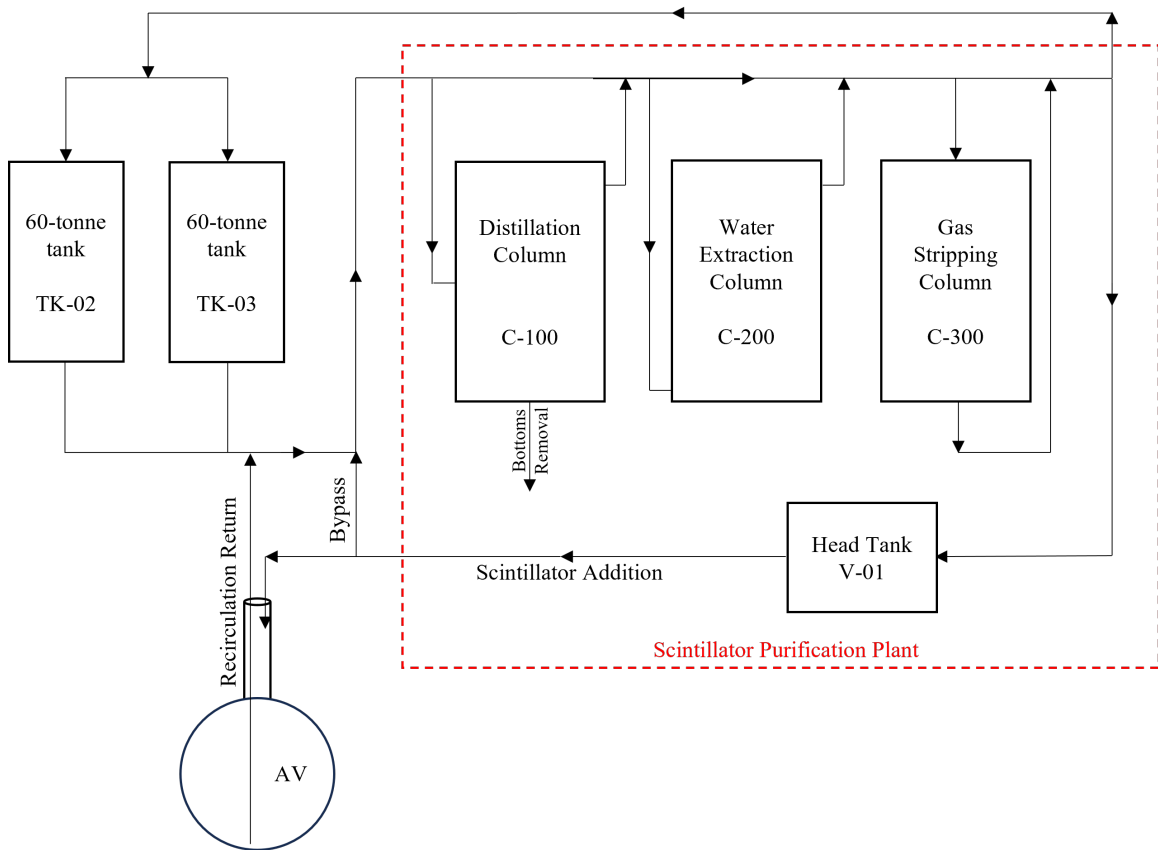


Figure 4.9: A simplified block flow diagram summarising the primary purification systems.

most effective method of contaminant removal [207]. This six-compartment fractional column allows for highly effective multi-stage distillation and allows for continuous (as opposed to batch) distillation. The compartments themselves were separated by dual-flow trays, allowing for superior efficiency per stage when compared to traditional downcomer trays as they allow for phase contact in the full stage volume.

The LAB is introduced at the midpoint of the column, flowing down until it is vapourised. As such, dense contaminants flow downwards while pure LAB is sent upwards. The vapour at the top of the column is then re-condensed, and can be sent

back (“refluxed”) into an upper stage in the column for re-vapourisation leading to increasing levels of contamination removal. The increasingly contaminated liquid at the bottom of the column (“the bottoms”) are removed in small batches. The boiling point of LAB at atmosphere is 278–314°C depending on the chain length, close to the autoignition temperature of 323°C. As such, the column was kept under vacuum, thereby reducing the boiling point to $\sim 200^\circ\text{C}$.

Operation of this column required constant manipulation of the heater power levels, bottoms removal rates, distillation feed rate, and reflux rate to maintain a stable output of LAB with acceptable quality. After a commissioning period, optimal operational conditions were discovered that allowed for the column to be operated stably at 220°C and 44 Torr, resulting in a constant distillate (i.e. distilled scintillator) flow rate of 750 kg/hr.

As the start-up temperature management was a complex manual operation, it could take upwards of 8 hours to get the column to a condition where distillate flow could be achieved. Similarly, the plant required direct attention while shutting down to ensure that the instrumentation would be left in a safe state. As such, the plant was nominally operated 24-hours a day and 6 days a week, and drove the schedule (and logistics train) of the entire scintillator fill.

The distillation process was expected to be highly effective at removing lower volatility impurities [207], including all heavy metals of concern (Bi, K, Pb, Po, Ra, U, and Th) as discussed in Section 3.1.4 and Section 3.1.5. Oxidised organics such as carboxyl groups and 1,4-benzoquinone that are not radioactive but can reduce the optical clarity of the scintillator were also effectively removed through this process.

Water Extraction The C-200 column within the plant provided a polishing technique through which water and LAB — immiscible with each other — are brought into close phase contact with one another. UPW from the water purification systems (described in Section 2.2.3) would be introduced near the top of the 30" (0.762-m) diameter and 18'4" (5.6-m) tall column, while LAB could be fed near the bottom. Due to their density differences, the water would flow to the bottom of the vessel, where it would be removed and sent to waste. Meanwhile, the LAB would flow to the top and be extracted. As they interacted, a motorised impeller would mix the fluids to maximise molecular contact; the column is also designed to operate at $\sim 80^{\circ}\text{C}$ to increase efficiency.

Due to the polar structure of H_2O , the water would be highly effective at removing ionic heavy metals such as U, Th, Ra, K, and Pb, as well as any other fine dust particulates (which are typically charged). However, the use of the water extraction column was deemed not immediately critical due to the extreme effectiveness of the distillation column and complexities of running both C-100 and C-200 in parallel. As such, the C-200 was never operated during the scintillator fill, but can be used in the future during recirculation of the AV medium.

Gas stripping The final major purification system on the scintillator plant occurred within the 24" (0.61-m) diameter 22'8" (6.1-m) tall C-300 column. Relatively simple compared to the other columns, this stripping column was filled with Koch-Glitsch FlexiPac high density SS316L packing, a high surface area porous material. The scintillator was fed into the top of the column, which drains to the bottom through the packing at a maximum rate of 4000 kg/hr. Meanwhile, gas is sent into the bottom of the vessel, and vacuum pumped to the top of the vessel.

As this gas comes into contact with the scintillator, it carries away contaminants in a gaseous analogy to the water extraction counter-current equilibrium process. As a gas, the stripping is most effective at removing impurities with large vapour pressures, such as O₂, Rn, Ar, and Kr. As C-300 is also operated at 100°C and 150 torr, water is also removed from the scintillator.

Nitrogen gas was typically used as the stripping agent, as it can be easily sent into the scintillator plant from either cylinders or underground dewars. However, it was also possible to use superheated steam from the UPW plant as well. As radon (being non-polar) has a lower Henry coefficient in steam, which itself has an exceedingly low radon content due to its degassing in the water purification systems, this allows for a more effective radon removal. Using steam would also reduce the logistical challenges around nitrogen, which needed to be shipped into the laboratory. However, between the operational challenges of the steam system and the commissioning of an underground liquid nitrogen facility at SNOLAB in early 2022, the stripping within C-300 was performed nearly entirely with nitrogen gas.

4.3.2.2 Secondary Distillation Systems

The primary scintillator plant was observed to be highly effective at purifying the scintillator, to the extent that the effectiveness presented a complication. In particular, the column reduction factor for PPO was determined to be well over 100 before refluxing, even despite having a low relative volatility $\alpha \sim 0.11$ between LAB and PPO. As such, distilling PPO-loaded LAB would cause the PPO to be removed from solution and sent to the bottoms. This meant that the PPO had to be purified and

added independently to the LAB. To accommodate the purification of PPO, a secondary distillation system was developed within the scintillator plant, as summarised in Figure 4.10.




Figure 4.10: A simplified block flow diagram summarising the secondary distillation systems.

Approximately 1800 kg of PPO was required to achieve the desired concentration of 2.0 g/L; as some losses were expected, 1904 kg of PPO powder was initially procured in 119 16 kg barrels that were shipped into the underground laboratory. The PPO was procured from PerkinElmer Inc., and manufactured as “neutrino-grade,” wherein a special synthesis procedure that took care in mitigating radioactive contaminants was applied. The PPO powder would then be loaded and processed in batches of 36–48 kg.

Master Solution Preparation In order to load and purify the PPO, ~ 400 L of distilled LAB was transferred from V-01 into V-501, a 1000 kg preparation vessel. PPO powder was then added manually through a hopper at a rate of 100 g/min into V-501 to achieve a high concentration “master solution,” which was then mixed with an impeller while sparged with nitrogen gas. The original master solution concentration

target was 120 g/L (48 kg batches), but later optimised to 80 g/L (36 kg batches) to better match SNOLAB access constraints during the COVID-19 pandemic; the reduction in concentration allowed for full dissolution between the end of one access period and the start of the next. In later batches, a heater was also added to V-501 to promote dissolution speed.

This powder addition process was non-trivial, as the PPO powder exhibited adhesive properties that frequently caused the powder to solidify into solid clogs within the 3" line that separated the hopper from the top of V-501. The removal of these clogs not only introduced many contamination channels, but was physically labourous in an area of the scintillator plant that reached nearly 40°C. As such, the powder loading underwent five major iterations:

- Series 1** The original design of powder loading had operators adding up to 4 kg into a hopper, which was then sealed and sparged with nitrogen for several minutes. A valve separating the hopper from V-501 was then opened, and the PPO dropped in. This method caused severe clogging in line between the hopper and V-501.
- Series 2** The valve between the hopper and V-501 was left open, and the powder loaded directly into the vessel; light amounts of clogging occurred frequently. However, this left the V-501 volume briefly exposed to air, leading to ingress of radon and oxygen.
- Series 3** To counteract the ingress of air in Series 2 batch types, vacuum was pulled on V-501, after which the head space was re-filled by sparging the entire volume with nitrogen gas. This operation was repeated three times after each batch prior to mixing to counteract any gaseous contamination that may have entered the head space of the vessel and begun leaching into the master solution volume.

Series 4 The timing of every operation was optimised for reduced laboratory access conditions as a result of the COVID-19 pandemic.

Series 5 A set of 1"-diameter PFA pipes were constructed, which would then be deployed into the line between the hopper and V-501. PPO was then added into this interior pipe, which was removed and swapped out when clogged.

Following full dissolution, the master solution was then subjected to several (typically two) solvent-solvent water extractions, where ~ 250 L UPW was added to V-501 from the water purification systems discussed in Section 2.2.3. The impeller then mixed the volume for at least 3 hours to ensure full molecular contact, then turned off to allow for solvent separation. Although a majority of fluid separated within several minutes, an emulsified interface layer remained, which would take up to several days to separate out. Nominally, a settling period of 3 hours was allotted before the water was removed and sent to waste along with the interface layer. Through concentration measurements of V-501, 2.64% of the PPO was calculated to have been lost in this manner.

Secondary Distillation The master solution preparation process for each batch was optimised to take two days, after which the master solution was sent through 200 nm and 5000 nm hydrophobic filters to a 1000 kg holding tank, capable of storing up to two batches. From there, the master solution was fed into the E-110 kettle heated to 230°C and 20 Torr. At this temperature and pressure, the LAB was instantly vapourised while the PPO was slowly distilled. Operating this “flash distillation” process took similar degrees of manual intervention as the C-100 column, and the process was optimised to produce a distillate flow rate of 15 L/hr.

Following distillation, the master solution was cooled and recombined in a 280 kg recovery tank, from which it could be re-added to the primary distillation loop and mixed in-line with pure LAB after distillation in C-100 but before stripping in C-300.

4.3.2.3 Recirculation

Due to the challenging powder loading process and the slow flow rate achievable through secondary distillation, the purification of pure LAB outpaced the addition of PPO — when the last drop of AV UPW was replaced with LAB, the concentration of the PPO within was only 0.531 g/L. Furthermore, UPW had leached into the scintillator volume in the 907 days of the “AV bulk fill,” a name ascribed to the period between the first drop of scintillator and the final removal of UPW from the AV on 9 April 2021.^u Indeed, the humidity of the AV scintillator was measured using Karl-Fischer (KF) titration to exceed 55 ppm at the end of the bulk fill.

Since dissolved water may endanger the chemical stability of tellurium-loaded scintillator as discussed in Section 5.2, the AV volume was recirculated through the scintillator process system, where it was stripped in C-300 for dehumidification. Meanwhile, secondary distillation operations continued, and the AV became increasingly doped with PPO as recirculation proceeded over the period of a year. Unconstrained by the limited output flow rate of the primary distillation column, recirculation of the AV could proceed at rates exceeding 10 tonnes/hr. By the time the PPO top-up was completed on 29 April 2022, the humidity within the AV had dropped below 5 ppm,

^uThe bulk fill is occasionally erroneously noted to have ended on 27 March 2021, when a lengthy de-watering operation was projected to start. This de-watering process aimed to remove the final 5 tonnes of UPW beneath the level of the deployed PFA pipe through a bubbler line, only possible with a low-flow peristaltic pump. However, the PFA pipe had stretched such that the de-watering operation was not needed, and the final amounts of UPW was replaced with LAB through the regular process on 9 April 2021.

the detection limit of the KF technique. A summary of this PPO addition and water removal through the duration of the recirculation period can be seen in Figure 4.11.



Figure 4.11: A summary of the recirculation and PPO top-up period showing (top) the addition of PPO within the detector, and (bottom) the removal of dissolved water from the AV scintillator.

As the AV is saturated with up to 2% water by mass from both the water phase and inner cavity, water is continuously diffused into the scintillator bulk. As such, recirculation of the deployed scintillator through the gas stripping column will need to be occasionally repeated throughout the lifespan of the experiment in order to minimise water concentration in the scintillator.

The scintillator process system was capable of recirculating the detector in either regular or reverse directions.^v In general, adding scintillator to the top of the AV

^vThe regular direction is the mode akin to the bulk fill, when scintillator is added to the top of the AV while removed from the bottom.

promoted mixing of PPO within the detector [208], while adding to the bottom suppressed mixing. The latter was occasionally used when testing various changes to the process systems, as much of the volume could be recovered if necessary by reverting to recirculating in the regular direction.

4.3.3 Scintillator Quality Assurance

To undertake this process engineering of the scintillator plant, the expertise of Koch Modular Process Systems, LLC (KMPS) was called upon, for they had previously spearheaded the design of the famously successful Borexino scintillator purification systems [209]. However, unlike SNO+, Borexino was located at the Laboratori Nazionali del Gran Sasso (LNGS) and could be accessed by vehicle, greatly simplifying the design considerations and construction of the plant.

As the scintillator plant was built to be underground to enable the capacity to recirculate the detector, the design of the scintillator plant was placed under severe engineering constraints. All components of the purification plant would need to be transported through the mine shaft and thus limited in size. The scintillator plant — effectively a highly specialised oil refinery — would have to safely fit into a mere corridor within the SNOLAB facility itself. There were further constraints on electrical power, cooling, and reagent availability that would be trivialities in the case of similar process systems built in typical surface conditions. A photograph of the completed purification plant showing $\sim 15\%$ of the process systems is seen in Figure 4.12.

To monitor the hardware and performance of the scintillator plant, hundreds of instruments fed active on-line monitoring information and were paired with automated




Figure 4.12: A photograph of the completed scintillator purification plant showing $\sim 15\%$ of the process systems.

alarm systems. Furthermore, many unmonitored failures would themselves preclude stable operating conditions on the plant. Nevertheless, there were uncountable possibilities for inconspicuous and unmonitored hardware failures that could imperil the purify of the scintillator, and thus the potential of the entire experiment.

To confront this, an expansive quality assurance (QA) campaign was initiated to ensure that the purity of the scintillator maintained only the highest possible standards at every stage of the filling process. This both aided in the identification and repair of otherwise unseen damage on the scintillator process systems, and served as the last line of defence for the experiment.

4.3.3.1 QA Testing

Of the potential issues on the scintillator plant, the most common was introduction of an air leak, which could occur in multifarious ways ranging from ageing weld points to mechanical wearing on valves. Such air could carry with it suspended dust particles containing U and Th, as well as oxygen and radon. Although leaks in the system after the purification processes would lead to the introduction of these backgrounds into the AV, leaks prior to distillation were no less serious. The introduction of oxygen causes LAB to undergo oxidation to form 1,4-benzoquinone (para-benzoquinone, *p*-benzoquinone) [210], which is severely problematic as the latter causes strong distinctive absorption peaks at 368 nm, 389 nm, and wavelengths above 500 nm, resulting in a “yellowing” effect that diminishes both the light yield and the attenuation length of the scintillator.

This reaction is slow at room temperature — samples exposed to oxygen and measured with ultraviolet-visible wavelength absorption spectroscopy (“UV-Vis”) do not show significant yellowing after several years — but the oxidation reaction is vehemently accelerated by the heat of the distillation column. Although the C-100 *p*-benzoquinone column reduction factor was measured to be 44 before refluxing, the distillation capabilities were greatly outpaced by the production rate of the contaminant in the column when exposed to minimal amounts of air. Even small amounts of this *p*-benzoquinone that ended up in the AV would cause a darkening of the detector.

Defence Against the Dark Arts Yet for the danger it brought, the *p*-benzoquinone could be wielded to uphold the quality of the scintillator. The distinctive, quick-forming peaks allowed for immediate identification of oxygen in the process fluids,

which itself was indicative of the ingress of air and the radioactive contaminants carried with it. A UV-Vis absorption measurement of the scintillator can be therefore used as a proxy for the radiopurity of the scintillator and by extension, the performance of the purification plant. This measurement could be performed rapidly, with turnaround times of under 10 minutes from the time of sampling to result.

Thus, the defence of the AV primarily rallied around UV-Vis 1 cm absorption measurements of the scintillator using a Thermo Scientific™ Orion™ AquaMate 8000 UV-Vis spectrophotometer. Each sample was compared to UV-Vis absorption measurements of a “control” sample — the best possible LAB produced using small-scale benchtop distillation in the surface laboratory facilities — as well as a “filtered” sample, which was LAB taken from TK-01 and sent through a simple open-air paper filter. This “filtered” LAB noted the quality of the solvent as received from CEPESA and served as the benchmark of the scintillator process systems. Although the peaks associated with *p*-benzoquinone were the primary observation targets, all spectral anomalies were investigated. In lieu of such anomalies, the baseline of the spectra in the 360-400 nm region above the absorption of PPO — relative to that of the 600-800 nm region where no absorption is expected — was tracked to verify the effectiveness of the distillation, as superior purity would result in advanced optical clarity and thus a lower absorbance. Figure 4.13 is a representative UV-Vis measurement of purified, distilled LAB of moderate quality undergoing closed loop recirculation during scintillator plant commissioning, compared to the another sample taken only 20 minutes later. Through this, it was discovered that a valve had been improperly closed by mere degrees. An absorption measurement showing excellent quality LAB approved for filling the AV can be seen in Figure 4.14.




Figure 4.13: 1 cm absorbance spectra showing (left) LAB purified within the scintillator plant during scintillator plant commissioning to a moderate quality when compared to control and filtered LAB and undergoing closed loop recirculation, and (right) the same amount of LAB within the scintillator plant only 20 minutes later after a small leak exposed the fluid to oxygen. The 368 and 389 nm peaks are distinctive of *p*-benzoquinone.

These UV-Vis measurements were supplemented by turbidity^w measurements performed using a Hach[®] TL2310 LED Turbidimeter, as well as density measurements performed using an Anton Paar GmbH DMA 35 Portable Density and Concentration Meter. The turbidity and density measurements allowed respective elucidation of suspended and dissolved contamination within the sample, that may not be observable using a 1 cm absorption measurement. The QA constraints on these two measurements are seen in Table 4.8. These three analyses were nominally performed on every sample during the fill, and together are hereafter referred to as the “QA analysis”. During the PPO top-up and recirculation period, regular KF titrations were also included as part of the standard QA analyses to ensure that the LAB was being appropriately dehumidified.

^wTurbidity is also referred to as nephelometry, which is more precisely the technique used by the turbidimeter.




Figure 4.14: Absorbance spectra for scintillator of excellent quality, approved for use in the AV. (Left) The scintillator has superior optical clarity compared to control LAB at all wavelengths above where PPO absorbs. (Right) The PPO absorbance spectrum, recovered from subtracting the control LAB absorbance spectrum from the total absorbance spectrum.

Table 4.8: The acceptable range of QA measurements for the turbidity and density (corrected to 20°C) of pure LAB.

Constraint	Acceptable range
Density (Pure LAB)	[0.8540 – 0.8560] g/cm ³
Density (Variable PPO up to 3 g/L)	[0.8540 – 0.8800] g/cm ³
Turbidity	[0.50 – 0.65] NTU
Water Concentration	< 5 ppm

Testing Regime Armed with these capabilities, the scintillator samples were tested for quality at every stage of the process, from the delivery by CEPESA to the last moments of recirculation. For every measurement, the scintillator was collected into an ultrasonically cleaned glass jar with a PTFE lined lid; the glassware and lid were confirmed to be compatible by measuring the absorption spectrum several times per

week over a month. No change in the absorption spectrum was observed. Prior to taking the sample, the port from which the scintillator was retrieved was cleaned using methanol, and drained for 10 L to eliminate residue that had collected externally to the port. The glassware was then rinsed with the fluid three times before the sample was taken. While transporting or storing each sample, the lid was secured with several layers of teflon tape to mitigate air seepage into the jar. The external surface of the jar was then cleaned with methanol before it was placed in a new pre-cleaned sliding channel storage bag. Samples that included high concentrations of PPO (i.e. master solution) were placed in amber jars or amber bags to minimise degradation from light exposure.

With each tanker shipment, CEPESA provided a Certificate of Analysis (CoA) to demonstrate that their measurements of the manufactured LAB met previously established standards. Measurements from a representative CoA are shown in Table 4.9. Nevertheless, QA analyses were performed with every tanker shipment prior to the transfer to TK-01. At this stage, samples sourced for “filtered” LAB and for bench-top distillation into “control” LAB were also taken. Similarly, QA analyses were also performed regularly on samples drawn from the line between the railcars and the 60-tonne tanks.

Whenever the scintillator process systems were running stably and distillate flow into the AV was possible, the primary distillation loop was primed with LAB and run in closed-loop recirculation mode. Three consecutive QA analyses were then taken to ensure that the distillation process was monotonically improving the LAB before the in-line mixing of PPO master solution or scintillator flow to V-01 was allowed to proceed.

Table 4.9: LAB parameters as measured by CEPESA. These parameters were measured for each of the 45 tanker shipments received and presented in a CoA to confirm that the material fell within desired specifications.

Parameter	Measurement	Specification
Density (20°C)	0.8595	0.85–0.87
Water (ppm)	13	< 100
Bromine (mg/100g)	1	< 10
Saybolt colour	30	> 29
Tetralins/indanes (%)	0.6	< 1
Unsoluble matter (ppm)	0	0
Acid wash test transmittance (%)	83	> 35
2 phenyl LAB (%)	28.2	27–31
Molecular Weight (g/mol)	234.8	233–237
9 C chain (%)	0.9	< 3
10 C chain (%)	18.0	< 25
11 C chain (%)	45.5	–
12 C chain (%)	30.4	> 25
13 C chain (%)	5.2	< 15
14 C chain (%)	0.0	< 2
> 14 C chain (%)	0.0	< 0.3
Paraffin content (%)	0.04	–

As the detector was filled from V-01, QA analyses were performed at the last accessible sampling port, immediately before the scintillator entered V-01. If any anomalies were detected at any time, the scintillator plant reverted to closed-loop

recirculation and the V-01 volume returned to the primary distillation loop or sent to waste until the anomalies were understood, corrected, and the scintillator quality was verified; this process could take minutes (such as the example noted earlier with the improperly closed valve) to months (in the case of major infrastructure failure). Each time the primary distillation paused, a full QA approval process involving the three consecutive analyses were required for the AV fill to resume.

PPO Testing The QA regime during the preparation and distillation of PPO master solution was more challenging yet no less vigorous. The high concentration in the PPO master solution meant that turbidity measurements were untenable as undissolved, suspended solid PPO was inevitable short of several weeks of heated mixing. Furthermore, the PPO absorption spectrum overlapped with the lower *p*-benzoquinone peaks, making identification more challenging at higher PPO concentrations. To ensure that there were no spectral anomalies hidden beneath the PPO absorption spectrum, the UV-Vis measurement of the control LAB was subtracted from the UV-Vis measurement of the master solution. The resulting “control-subtracted” spectrum could then be surveyed for anomalies. A typical UV-Vis absorption spectrum used to approve the scintillator for deployment into the AV can be seen in Figure 4.14, which includes the control-subtracted PPO spectrum. The approved scintillator has superior optical clarity when compared to control LAB at all wavelengths above where PPO absorbs.

The control-subtracted spectra also provided an opportunity to measure the PPO concentration. 20 LAB + PPO calibration samples were developed, with concentrations ranging from 0–5 g/L. It was determined that the PPO concentration between 0.5–6 g/L could be recovered with a third degree polynomial fit on the absorption at

350 nm. Similarly, lower concentrations (0–0.5 g/L) could also be recovered with a third degree polynomial fit at 375 nm, albeit with a different slope. The fit parameters for recovering PPO concentration can be found in Table. 4.10.

Table 4.10: The relationship between the concentration of PPO in the 0–0.5 g/L to 0.5–6 g/L ranges and 1 cm absorbance at 346 nm and 350 nm, respectively, follows a third degree polynomial. These third degree polynomial fit parameters for $C = aA_x^3 + bA_x^2 + cA_x + d$ where C is the concentration and A_x is the absorbance at wavelength x are shown.

Parameter	Concentration Range (g/L)	
	0–0.05	0.5–6 ($\times 10^{-3}$)
a	0.192207196	1.612200877
b	-0.28597763	9.402655972
c	2.307639142	102.4706
d	0.0637178	0.38035072

Both the capability to recover the control-subtracted spectrum and the PPO concentration was built into the software used for UV-Vis measurement analysis and taken with every sample as part of the standard QA analysis. As removing a sample from the AV for direct measurement was not possible while UPW was still deployed, the recovery of PPO concentration from the hourly measurements allowed for an integrated estimate of PPO within the system.

Such PPO measurement capabilities were also critical in ensuring the quality of the master solution through the preparation and distillation processes. Recovering the PPO concentration of master solution was more complicated, as the fits in Table. 4.10 do not extrapolate to high concentrations. Therefore, the concentrations could only be

recovered by first performing a 50:1 dilution using a volumetric flask. Nevertheless, to assess purity and proper dissolution, QA analyses (which all included a measurement of PPO concentration after dilution) were performed at every stage of the master solution preparation:

- after V-501 was refilled with distilled LAB and prior to loading of PPO,
- after the loading and mixing of PPO, before each water extraction,
- after each water extraction had completed; the output was also examined with hydrocarbon test strips to verify that the water and interface layer had been properly removed,
- after the master solution was moved to the holding tank prior to distillation, and
- after the master solution had been moved to the PPO recovery tank prior to in-line mixing.

Master solution batches that had failed to meet the QA standard could either be sent to waste or siphoned back to V-501. In the latter case, the master solution would be fully reprocessed as a new batch (excluding addition of more powder).

Light Yield Testing As seen in Figure 4.6, the addition of PPO should bring with it an increase in light yield. To assess this, the Scintillator Counter Of Uranium and Thorium (SCOUT) was set up underground near the scintillator plants. SCOUT is a 9.86” diameter and 5.75” tall acrylic cylinder placed within copper and lead shielding. Built into the shielding are four ADIT 3” PMTs, which observe the volume of the acrylic cylinder from below. A source compartment above the acrylic cylinder allowed

for deployment of radioactive sources to activate the scintillator; typically used for SCOUT was a 3.6 MBq ^{60}Co source. A CAD drawing showing the four PMTs under the acrylic cylinder can be seen in Figure 4.15.




Figure 4.15: A CAD drawing of SCOUT, showing the acrylic cylinder above four PMT support units enclosed within scopper and lead shielding.

As the master solution was mixed in-line into the closed loop, samples were drawn from the plant and placed within the acrylic chamber before being measured in ten-minute exposures. The integrated charge measured by the PMTs was then compared

to those of calibration standards created using undistilled P-500-Q LAB to achieve a relative light yield measurement. In all samples measured, the relative light yield of the samples were superior to that of the calibration standards. Due to the complexity and labour-intensity of the SCOUT light yield measurement, it was only performed once for each batch of PPO. Later in the fill, the confidence was so high in the light yield output of the deployed scintillators that the measurements were scarcer yet.

Yet despite these successes in SCOUT, the device was originally intended to be a $^{214}\text{BiPo}$ counter to estimate the amount of U and Th within the scintillator. This functionality was never used, as there was a superior instrument that was making the same measurement - the detector itself.

4.3.3.2 Radon Tracking

As had been mentioned in Chapter 3, the irony of an experiment such as SNO+ is that the detector is the only instrument sufficiently sensitive to ensure its own success; preliminary measurements taken using SCOUT limits the sensitivity to 10^{-10} g/g U and Th equivalent, even after 24 hours of data-taking, a factor of $\mathcal{O}(10^4)$ above target contamination levels. However, measurements with sensitivities exceeding 10^{-17} g/g U and Th could be made using the detector within the same time frame. Given the rigour of the QA programme and the fact that the first 100 tonnes of scintillator (corresponding to the combined volumes of the 60-tonne tanks) could be removed from the AV with relative ease, the SCOUT U/Th measurement was therefore bypassed in favour of a direct *in situ* measurement using the detector.

To perform this measurement, $^{214}\text{BiPo}$ coincidences occurring within the detector were tagged on a run-by-run basis, using the coincidence-based methodology discussed

in Section 3.1.4. The coincidence tagging parameters (“cuts”) used are found in Table.4.11. From the rate of decay (discernible from the exponential coincidence time between the Bi and Po), it was immediately apparent from the tagged $^{214}\text{BiPo}$ coincidences that the $^{214}\text{BiPo}$ count rates were dominated by ^{222}Rn . This analysis constrained the sustained contamination level of ^{238}U to under 9.759×10^{-17} g/g U during the partial fill period after sufficient time was afforded for the radon to decay away. This agreed with other background analyses that constrained this rate in the partial fill to $(4.7 \pm 1.4) \times 10^{-17}$ g/g U [211].

Table 4.11: The cuts used to tag $^{214}\text{BiPo}$ coincidences during the scintillator fill. The skyShine cut is a classifier comparing the number of side to bottom hits, while fitValid and the DC bitmask ensure that correct data cleaning parameters described in Section 6.1.2 were used.

Parameter	Cut
^{214}Bi nhits	[400, 1700]
^{214}Po nhits	[160, 400]
Δr	[0,1] m
Δt	[3690, 1×10^6] ns
skyShine	> 1
fitValid	==1
DC bitmask	2x210000000242
FV	5500 m
PSUP-AV z-correction	15 cm

This analysis was also able to determine that high quantities of ^{222}Rn was entering the detector whenever it was being filled with scintillator. The amount of radon was

further increased when in-line mixing of PPO master solution had been carried out, as can be seen in Figure 4.16. This implied that there was indeed an ingress of air in both the primary and secondary scintillator plant systems. Interestingly, the fact that newly added PPO had a higher radon content meant that the addition of PPO within the detector could effectively be tracked by reconstructing the decay position within the detector. Films showing the movement of the PPO within the detector were produced, with each frame corresponding to an hour of tagged ^{214}Po locations within the detector; one such frame is shown in Figure 4.17. This provided information on internal flow and mixing mechanics within the detector, and allowed for an understanding of fill conditions that promoted a quickest path to a homogeneous detector.

Tagged ^{214}Po from multiple runs were also frequently collated into single larger datasets in order to find sources of sustained radon ingress. One such example can be seen in Figure 4.18, which collated the tagged ^{214}Po from 1030 runs in April 2020, early in the partial fill phase. An obvious hot spot at the top of the detector was identified, with an x-y position offset that correlated with the deployed PFA pipe. It was postulated that the line had lost prime near the top of the AV, and had backfilled with air which was then diffusing through the PFA wall and into the AV volume. The PFA pipe was sealed off, and the radon began to decay away; as can be seen in Figure 4.19, a collection of 950 runs across July 2020 — after roughly 10 half-lives of Rn had passed — show that the hot spot was no more.^x

^xAs the partial fill phase persisted until October 2020, no AV filling had occurred in the time between the data shown in Figure 4.18 and Figure 4.19.




Figure 4.16: The tracking of radon within the AV depending on various fill modes using a representative time frame of 23 October 2020 - 12 December 2020, while the Series 3 PPO loading procedure was used. The top bar shows the status of the various PPO batches in this time period. The radon within the detector increases at a higher rate while PPO is added.

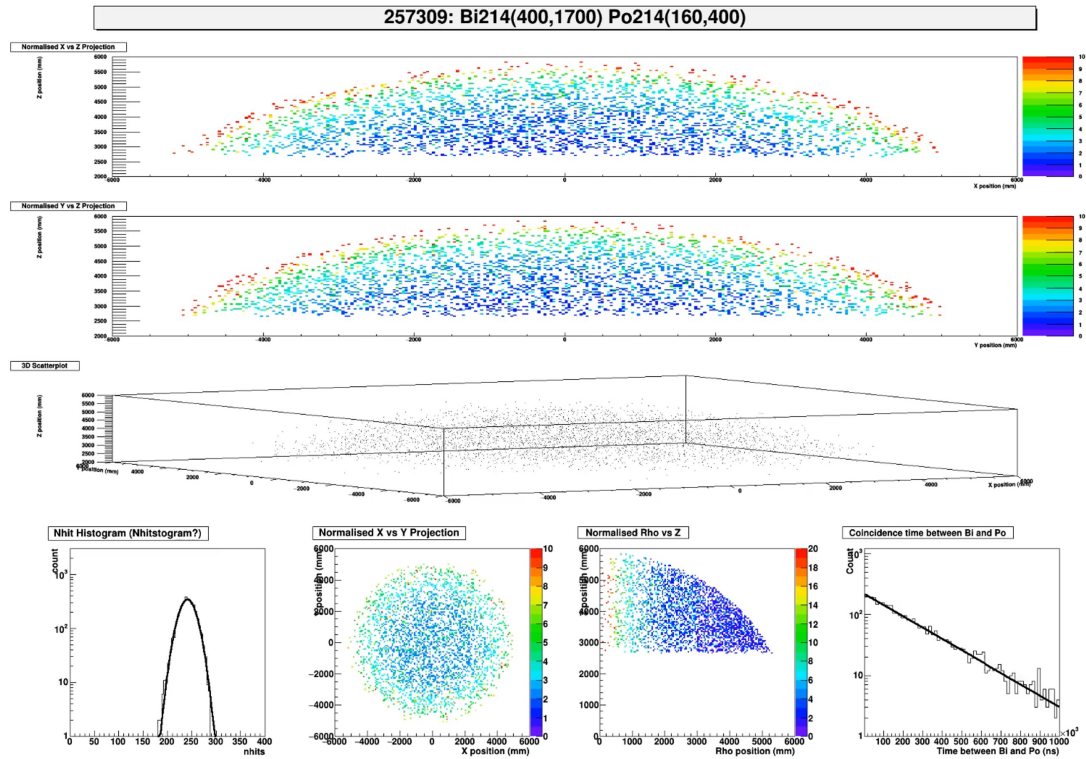


Figure 4.17: The ^{214}Po tagged using the parameters in Table 4.10 for a representative run during the scintillator fill. Similar plots were created for every run during the scintillator fill, and constructed into videos to elucidate movement of radon within the detector.

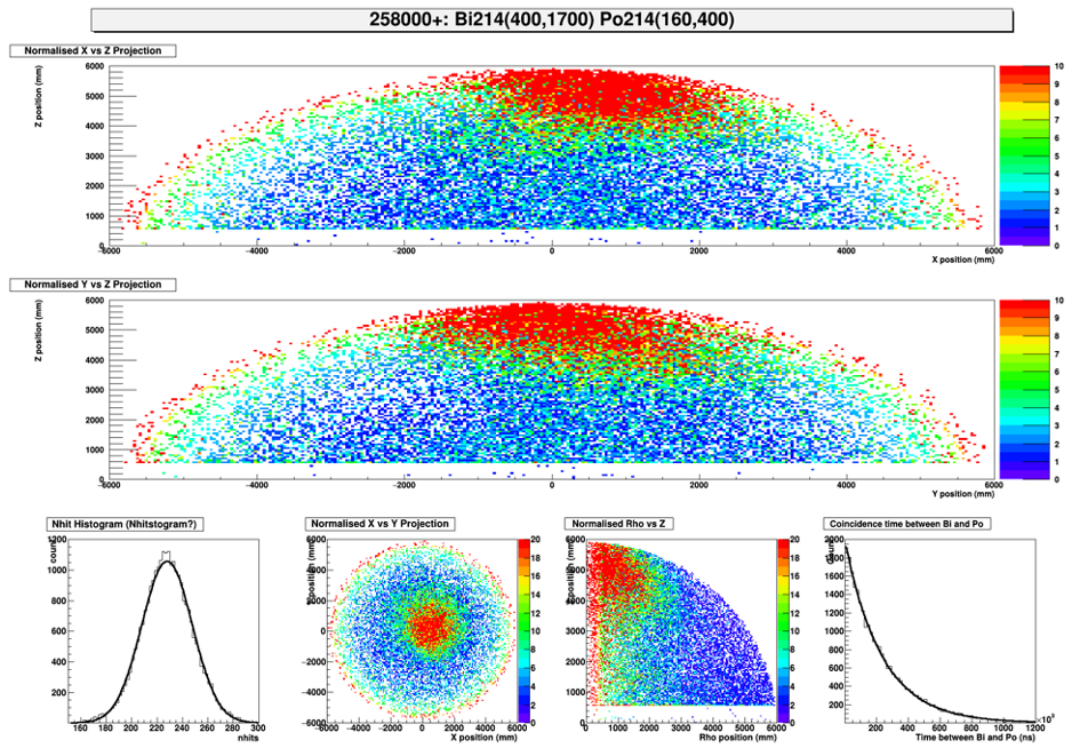


Figure 4.18: The ^{214}Po tagged using the parameters in Table 4.10 for the 1030 runs that occurred in April 2020.

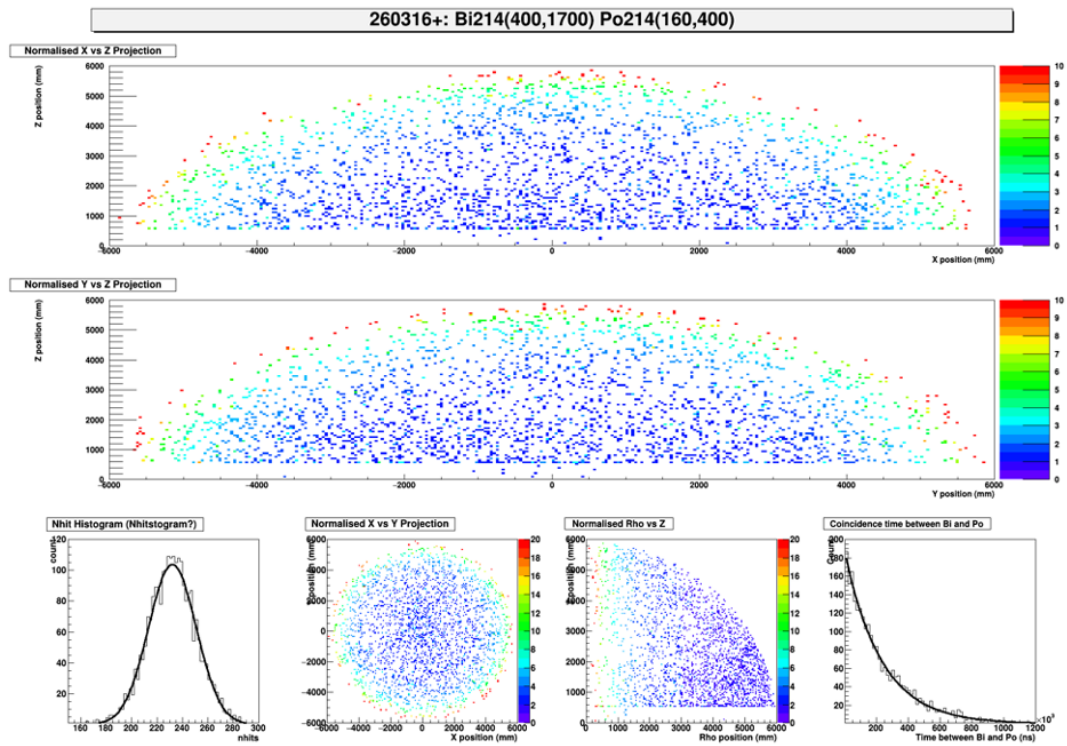


Figure 4.19: The ^{214}Po tagged using the parameters in Table 4.10 for the 950 runs that occurred in July 2020.

That radon infiltrated the detector undetected by measurements targeting *p*-benzoquinone implies that the leak in the system is small and occurs at some point after distillation. The fact that the leak did not bring with it obvious quantities of radioactive contamination also implies that the leak is likely within an enclosed area such as an insulated line, as much of the infrastructure in the scintillator plant are wrapped in layers of insulation for safety and operation consistency. However, the encapsulation of the hardware makes hunting for the leak challenging.

Radon tracking during the recirculation phase showed that the radon ingress remains high if C-300 is operating, even after primary distillation systems were shut down, narrowing the leak to the gas stripping systems. The hunt for this leak is ongoing, well over a year after the end of the fill. For now, usage of the gas stripping systems will regrettably necessitate a “cooldown period” in which the radon brought into the detector must decay away.

4.3.4 End of an Era

The most important cooldown period occurred at the end of the scintillator fill. The last gram of PPO was added on 22 April 2022, and after nearly a week of constant AV recirculation at the maximum rate of 10 tonnes/hr, UV-Vis measurements on the AV scintillator and radon tracking confirmed the homogeneity of the detector on 29 April 2022. The scintillator fill was complete, and the scintillator phase period of quiet^y data-taking had begun.

Hence was accomplished a major SNO+ milestone, as the custom scintillator upon which the experiment was fundamentally based upon was at last deployed. The final

^yIn this context, “quiet” describes a period of time when the detector volume was minimally disturbed by underground activities.

AV volume was 792.4 tonnes of LAB doped with 2.20 g/L of PPO. Notable in the concentration is that the amount of PPO exceeded the original projection of 2 g/L.^z

During the bulk fill, the PPO addition process was lossier than expected; although sufficient PPO was procured to reach a concentration of 2 g/L if an efficiency of 96% was achieved, measurements of the PPO concentration after 32% had been loaded found that the efficiency was closer to 59%. Although many of these unexpected losses were related to commissioning challenges that were later resolved, it was determined that the efficiency could only be optimised to 94%. An analysis of the PPO concentration at various points within the process systems indicated that PPO solids were being deposited across the kilometres of piping that made up the scintillator plant due to the high master solution concentration.

The concentration was decreased and additional PPO was procured to make up for this shortfall. However, when the bulk fill ended, the flow rate within the process systems increased from a maximum of 750 kg/hr (limited by primary distillation) to 10 tonne/hr (nominal recirculation feed). As such, some of the PPO that had deposited in the piping and presumed lost was recovered to reach a final concentration of 2.2 g/L. This final AV concentration measurement was achieved in three ways: a UV-Vis measurement on the AV scintillator, external matrix spike testing, and through integration of the sample concentrations taken throughout the fill.

The accuracy of the integrated technique was reflective of the diligence of the QA programme. Throughout the years-long operation of the scintillator plant, from the first drop on 25 October 2018 until the end of recirculation on 29 April 2023, 4925 samples were analysed from the efforts of 115 individuals who kept watch of the scintillator around the clock; a summary of this expansive QA programme can

^zOne could say 110% effort was put in.

Table 4.12: A summary of the QA programme during the scintillator fill period, which lasted between 24 October 2018 and 29 April 2022.

LAB added to AV	792.4 tonnes
PPO added to AV	2.01 tonnes
Volume recirculated	4376.9 tonnes
Team members	115 individuals
Trained QA analysers	51 individuals
Trained PPO powder loaders	32 individuals
QA analyses performed	4925 samples
Discrete PPO batches added	80 batches

be seen in Table 4.12. Although there are numerous secondary physics goals in the scintillator phase (as discussed in Section 2.5.4), the main objective of the scintillator phase is to determine if the efforts of so many were in vain.

4.4 Scintillator Phase Ex-Situ Measurements

While sufficient data was being collected for such analyses, there were other direct measurements that were made to provide a preliminary characterisation of the quality and properties of the scintillator. To facilitate this, nearly 40 L of the scintillator was removed from the AV in the final moments of recirculation before the scintillator phase began.

4.4.1 Absolute Light Yield

A sample of the completed AV scintillator was deoxygenated for 30 minutes, and inserted into a SNO UVT acrylic block coupled directly to a 1" R7600-200 HQE trigger PMT. A 12" Hamamatsu R11780 HQE detection PMT was placed 0.5 cm away. The same measurement was also performed on two samples of AV scintillator taken earlier in the recirculation process. These measurements were compared to those of 14 standards created with LAB doped to different concentrations of PPO. These measurements were normalised to a sample of 2.0 g/L PPO concentration, previously discussed in Sec 4.2.4.2 to have a measured absolute light yield of 11920 ± 630 photons/MeV.

The measurements of the standards were fitted to a two-exponential fit of $L_R = ae^{bC_B} + ce^{dC_B}$, where L_R is the relative light yield normalised to PPO concentration $C_B = 2.0$. The fit parameters and χ^2 are displayed in Tab. 4.13. As can be seen in Figure 4.20, the samples of the AV scintillator are well described by this fit and align with external measurements from [196] first shown in Figure 4.6. The measurement of the final 2.2 g/L scintillator had a normalisation of 0.990619, corresponding to an absolute light yield of 11808 ± 630 photons/MeV.

Table 4.13: The fit parameters for the two-exponential fit of $L_R = ae^{bC_B} + ce^{dC_B}$ performed on measurements of LAB + PPO, where L_R is the relative light yield normalised to PPO concentration $C_B = 2.0$.

a	b	c	d	χ^2
0.8419	0.0603	-0.5063	-2.696	0.02647




Figure 4.20: The measured light yield for samples of (red) 14 standards of LAB doped with various concentrations of PPO, (blue) measurements published in [196], and (gold stars) the deployed SNO+ scintillator taken during various times in the fill. The deployed SNO+ scintillator at 2.2 g/L is the final scintillator phase material. A two-exponential fit was performed on the standard measurements, with fit parameters shown in Table 4.13.

4.4.2 Density

The temperature-dependent density of the scintillator was also measured with the same Anton Paar DMA 35 portable density and concentration meter used for QA analyses throughout the fill. Prior to measurements, the device was calibrated using UPW at 20°C. Measurements were taken at SNOLAB atmospheric pressures (0.12 MPa), and changes in pressures typical within the laboratory ($\pm 5\%$) are expected to have a negligible effect [212]. To achieve cooler and warmer temperatures,

samples in pre-cleaned vessels were chilled in a refrigerator or placed in a hot bath, respectively. Measurements were then taken every 0.1°C while the aliquot reverted to room temperature.

Table 4.14: The density temperature-dependence parameters for the linear fit of $\rho = \alpha T + \beta$, where ρ is the density and T is the temperature.

α	β
-0.00070831	0.87122

The density temperature dependence can be seen in Figure 4.21, and is well described with a linear fit $\rho = \alpha T + \beta$, with fit parameters shown in Table 4.14. This fit agrees with external measurements published in [212].

4.4.3 Absorption

A 1 cm UV-vis absorption measurement was performed using the same Thermo Scientific™ Orion™ AquaMate 8000 UV-Vis spectrophotometer as was used in QA analyses. The average of three measurements are shown in Figure 4.22 which demonstrate the superior optical clarity in absorbance given as,

$$A = \log_{10} \left(\frac{\phi_e^i}{\phi_e^t} \right), \quad (4.4)$$

where ϕ_e^i and ϕ_e^t are the measured and transmitted fluxes, respectively. When compared to measurements of control LAB, the SNO+ scintillator exhibits superior optical clarity at all wavelengths above the absorbance of PPO at 2.2 g/L.




Figure 4.21: The measured density of the SNO+ scintillator at various temperatures and 0.12 MPa. For cooler measurements, a sample was placed in a refrigerator and measured while it warmed to room temperature; error bands are shown in blue. For warmer measurements, a sample was placed in a hot bath and measured while it cooled; error bands are shown in red. Also shown are measurements of the Day Bay scintillator at various temperatures and pressures, as published in [212]. The temperature dependence follows a linear fit with parameters shown in Table 4.14.




Figure 4.22: The 1 cm UV-Vis absorbance measurements of (blue) the SNO+ scintillator, (green) P-500-Q high purity LAB procured from CEPSA, and (red) the absorption of PPO at 2.2 g/L inferred through control subtraction. As can be seen in the inset plot, the SNO+ scintillator has superior optical clarity when compared to the procured high purity LAB.

Also performed were 10 cm UV-Vis absorption measurements using a PerkinElmer Lambda 800 Spectrometer. As shown in Figure 4.23, it was identified that the absorbance measured at certain wavelengths increased as the fill proceeded. This implies that absorbing agents were introduced into the scintillator at levels below the detection limit of the 1 cm measurements. Through tests performed on the scintillator in closed loop and re-analysis of saved samples from primary and secondary distillation, these contaminants were determined to have entered the AV in two ways:

- Despite nitrogen pump-purging performed on the master solution following open-air loading (series 2 onwards), trace amounts of oxygen remained, producing trace amounts of *p*-benzoquinone in the scintillator when heated through the secondary distillation column. Either the kettle distillation factor or bottoms removal rate may not have been sufficient to fully suppress these contaminants. Such peaks would have been unidentifiable in QA analyses of master solution due to the dominant PPO absorption at high concentrations, but built up over the course of the PPO addition.
- The ongoing leak in C-300 caused a build-up of oxygen in the detector, causing a minor development of *p*-benzoquinone species both in the 100°C of the column. This contribution was confirmed to have been suppressed in closed loop testing by increasing the nitrogen gas extraction rate from 3 kg/hr to 5 kg/hr, which removed oxygen at a faster rate.

Despite the slight increase in absorption, the attenuation lengths at wavelengths above 420 nm are still well over 20 m, sufficient when compared to the 6-m radius AV.





Figure 4.23: (Top) 10 cm UV-Vis baseline-subtracted absorption measurements of the SNO+ scintillator sampled at various times during the fill. The spectrum of the final SNO+ scintillator (2.2 g/L) exhibits reduced optical clarity relative to earlier samples. Through a subtraction against 1 and 5 cm absorbance measurements, the absorbance spectrum was converted into (bottom left) attenuation and (bottom right) attenuation lengths.

4.5 The Enhanced Scintillator Phase

An ingress of trace amounts of oxygen may pose a risk to the detector through the development of *p*-benzoquinone over time. To mitigate this risk, butylated hydroxytoluene (BHT) — a commonly used antioxidant — was procured from Sigma-Aldrich and added to the AV for a final concentration of 6.5 mg/L. The structural formula for BHT is shown in Figure 4.24(a). The BHT was added in powder form through V-501, nitrogen pump-purged, and water extracted. However, the BHT was added directly to the detector without distillation as to minimise both molecular damage to the BHT and further potential *p*-benzoquinone ingresses. The BHT was added to the AV in two rounds, on 29 June 2023 and 17 July 2023.



(a) Butylated Hydroxytoluene (BHT) (b) 1,4-Bis(2-methylstyryl)benzene (bis-MSB)

Figure 4.24: Structural formulae for the two reagents added to the scintillator cocktail for the enhanced scintillator phase.

BHT is optically inert, and was observed to have no effect on the absorption of the scintillator or performance of the detector. However, this would not be the case in the addition of another reagent in the cocktail: the wavelength shifter 1,4-Bis(2-methylstyryl)benzene (bis-MSB). The structural formula for BHT is shown in Figure 4.24(b). Added to further enhance the light yield of the detector, bis-MSB has an absorption peak of 347 nm and an emission peak of 423 nm. With absorption and emission spectra that do not strongly overlap, bis-MSB further boosts the light yield of the scintillator by emitting above the absorption range of *p*-benzoquinone, LAB, PPO, and itself, as can be seen in Figure 4.25.




Figure 4.25: The absorption (solid) and emission (dashed) spectra for the components of the enhanced scintillator: (blue) LAB, (green) PPO, and (yellow) bis-MSB. BHT does not to absorb or emit in these wavelengths.

Bis-MSB purchased from Curtiss Laboratories Inc. is thus to be added to the scintillator through the same secondary systems as PPO and BHT. Due to the easy aeration of the powder (thereby presenting a contamination risk if exposed to the laboratory air), the bis-MSB was pre-mixed with LAB into a slurry in a surface laboratory, and was then brought into the underground laboratory and added into V-501. Due to the poor solubility of bis-MSB in LAB when compared to PPO and BHT, additional modifications such as heat tracing were made to V-501 to maximise the stability of the bis-MSB master solution. As with the preparation of BHT, the bis-MSB was water extracted but not distilled. The desired final concentration within the AV is 3–5 mg/L, as optimisation studies have identified strong diminishing returns above this concentration [213]. As the energy transfer from PPO to bis-MSB is radiative and bis-MSB has a quantum yield of ~ 0.95 , the addition of bis-MSB will increase overall detection efficiency without affecting the intrinsic scintillator light yield.

Table 4.15: The components that make up the enhanced SNO+ liquid scintillator. In bold are the components that made up the original SNO+ scintillator cocktail used in the initial scintillator phase.

Material	Amount added (kg)	Concentration in LAB (mg/L)
LAB	792400	–
PPO	2010	2200
Bis-MSB	4.5	4.9
BHT	5	5.4

The first 0.6 mg/L of bis-MSB was added to the AV on 20 July 2023, concluding the initial scintillator phase; a preliminary analysis by looking at the shift in the ^{210}Po nhit

peak implies that the light yield immediately increased by 50–60%. The remainder of the wavelength shifter will be added after the completion of this document, which will herald the beginning of the enhanced scintillator phase, with a completed AV medium summarised in Table 4.15. Following this, all that will be left to add to the scintillator to enable the search for $0\nu\beta\beta$ is the tellurium itself.

Chapter 5

The Loading of Tellurium

Yet hope there is still, if we can but stand unconquered for a little while.


– Gan. 3:6

With the objectives of the scintillator phase under way, attention can be turned to the addition of tellurium into the detector. However, the loading of ^{130}Te into liquid scintillator is a novel process, as is the purification of ^{130}Te to levels necessary to maintain the overall purity of the scintillator. Indeed, the loading of ^{130}Te has never been attempted in any organic liquid scintillator. Any loading would have to be compatible with a final total purity of $\mathcal{O}(10^{-15})$ g/g equivalent for U and Th in order to be competitive. From the backgrounds measured in the partial fill phase (as discussed in Chapter 3), this implies that the target for trace contamination would have to be $< 1 \times 10^{-13}$ g/g U and $< 5 \times 10^{-14}$ g/g Th [131].

Tellurium is a brittle, insoluble metalloid that is naturally found and recovered through mining,^a and can be turned into a water soluble form when reacted with a strong oxidising agent such as hydrogen peroxide to form $\text{Te}(\text{OH})_6$ (orthotelluric acid,

^aThe Creighton Mine within which the experiment is located produces tellurium, though the tellurium used in SNO+ is sourced from elsewhere.

telluric acid, TeA),



The resulting TeA structure seen in Figure 5.1(a) has a net global dipole of zero, but possesses strong local dipoles from the -OH branches and thus soluble in water while immiscible within the non-polar oil-based LAB scintillator. It is thus in developing a method to purify and add the TeA to the scintillator that lay the next grand challenge of the SNO+ experiment.




Figure 5.1: Structural formulae for the two components which form TeBD, an LAB-soluble molecule containing Te.

5.1 The Diolisation Process

The initial attempt at creating a producing Te-loaded scintillator hearkened to the original industrial use case of LAB as an intermediary in the creation of linear alkylbenzene sulfonate (LAS), a critical ingredient in the production of synthetic detergents. As a surfactant, the LAS molecular structure is amphiphilic, consisting of a hydrophilic benzene-based head and a hydrophobic carbon chain tail. This provides the resulting detergent with the capability to mix oil and grease with water, providing the cleansing capabilities often found wanting in individuals in this field. The same mechanism allows for the dispersed suspension of aqueous nano-particulates within an oil-based solvent. Specifically, the aqueous TeA could be wrapped by hydrophilic head groups with hydrophobic tails that shoot out, allowing for dispersion within the non-polar scintillator medium. Although the loading of aqueous TeA was successfully dispersed in LAB using isopropylamine linear alkylbenzene sulfonate developed using the Calimulse[®] proprietary sulfonation (PRS) process, there were challenges in purification at large scales as the surfactant could not be distilled.


Another industrial process was turned to instead. Many pharmaceutical treatments require active medical ingredients to be loaded as aqueous substances into oils to be used as creams. As the polar aqueous and non-polar oil components would typically form unstable dispersions (emulsions) when mixed, aliphatic^b diols^c (glycols) are used to stabilise the dispersion of the aqueous component within the oil [214]. This same mechanism was found to work with stabilising aqueous TeA within LAB [215]. Although several different diols were found to be compatible with the process, 1,2-butanediol (BD) was selected due to its relatively low molecular weight, thereby

^bAliphatic compounds are those that do not contain benzene structures.

^cA chemical compound containing two hydroxyl (-OH) groups.

allowing a higher fractional Te content. The structural formula for BD is shown in Figure 5.1(b).

Both -OH BD moieties react with the TeA structure through a condensation reaction to form a range of similar compounds known as tellurium butanediol (TeBD),



where the resulting C–H bonds in the CH₃ functional group have low dipole moments, thereby improving miscibility in LAB. The resulting class of TeBD molecules is therefore a non-aqueous viscous liquid that dissolves readily in the hydrocarbon matrix that makes up LAB. As the bonds are bidentate,^d there can be 1–3 BD ligands that react with each TeA. However, as seen in Figure 5.2, nuclear magnetic resonance (NMR) measurements indicate a strong preference for the double attachment form.

BD can not only be distilled, but is relatively low cost and can be manufactured in a non-biogenic production process, thus suppressing ¹⁴C levels that can cause backgrounds (as discussed in Section 3.1.1). As such, this novel loading method is compatible with the low-background requirements of the detector. However, this technique is not without challenges.

5.2 Tellurium Stability

As seen in Equation 5.2, the forward condensation reaction is in equilibrium (\rightleftharpoons) with the reverse hydrolysis reaction. However, there is a natural preference towards the latter reaction, as TeA has an inherent acidity (with a pH of ~ 3.5 in a 10%

^dMeaning that the ligands bind through two donor sites.




Figure 5.2: NMR spectrum of TeBD molecules in water, with spectral formulae on top. As can be seen, the spectrum implies that the double-bidentate structure is dominant, with some contributions from the single-bidentate variation and unreacted TeA. The absence of peaks in the dotted region implies that there is a preference against triple-bidentate structures.

w/w aqueous solution) which is further increased by the addition of BD. Furthermore, -OH functional groups of a TeA molecule can dimerise,^e eventually causing an insoluble solid crystalline structure through oligomerisation;^f this is accentuated by the hydrolysis reaction, which unbinds the BD from the TeA -OH groups.

Altogether, the TeBD is not independently stable, especially in the presence of even trace amounts of water. Titration tests using a Rainin EDP-Plus electronic pipette to add distilled water in 1 μL increments into a 30 g sample of 0.5% TeBD in

^eDimerisation is when the unbound -OH groups from two TeA molecules can bind together causing a molecule with two Te centres.

^fOligomerisation is the extension of dimersisation, when multiple TeA molecules continue to bind together.

LAB determined that $< 10 \mu\text{L}$ water (~ 300 ppm) of water causes immediate separation of TeA from LAB. Furthermore, water concentrations of an order of magnitude less were observed to cause TeA crystals to precipitate out (“crash out”) of the solution. As the < 5 ppm humidity in the AV that was ensured at the beginning of the fill is not guaranteed to last due to the diffusion of water from the acrylic of the AV, this instability would imperil the experiment.

Many of these instabilities can be controlled by a Te:BD ratio during synthesis. Insufficient BD will cause an over-abundance of the single attachment form, thus driving oligermisation. However, due to the preference towards the cross-positional double bidentate TeBD form as seen in Fig.5.2, over-addition of BD will increase the acidity of the solution, driving the hydrolysis reaction. In order to provide a first order investigation on the maximum BD ratio, BD was dyed with trace amounts of Rhodamine B Base and titrated in $1 \mu\text{L}$ increments using the same Rainin EDP-Plus electronic pipette into a 30 g TeBD sample created with a 3:1 BD:Te ratio and loaded at 0.5% in LAB. It was determined that adding between 0.07–0.19 g BD corresponding with a final ratio of (3.5–4.3):1 BD:Te was sufficient to cause the TeA to crash out. As can be seen in Figure 5.3, the resulting precipitate was dyed, therefore suggesting that the excess BD was responsible for the instability.

5.2.1 Amine Stabilisation

Although tuning the BD:Te ratio can be used to slow the instability, full suppression of both the hydrolysis reaction and crystal formation could be achieved through the introduction of an amine. Consisting of a long carbon chain attached to a basic nitrogen atom with a lone pair, amines are aliphatic compounds that can act as a



Figure 5.3: A crashed sample of TeBD in LAB, caused by titration with dyed BD.

reverse surfactant. The hydrophilic N centers converge around the polar H_2O , while their long hydrophobic carbon tails improve miscibility in LAB. Furthermore, it is hypothesised that the nitrogen within the amine is also capable of de-protonating the $-\text{OH}$ groups in TeBD; the resulting ionic liquid association between the O^- in TeA and H^+ in the amine would suppress their capability to form crystalline structures. The introduction of the amine is thus capable of confronting the natural instability of TeBD through multiple interactions.

As summarised in Table 5.1, seven amines were identified as potential stabiliser candidates, though only three were identified to suppress instabilities for more than two months when added at a molar ratio of 1:1 between the amine and Te. Of these three remaining candidates (HDA, DMDA, and DDA), another distilled water titration test was performed on 30 g samples of 0.5% TeBD in LAB. In the case

Table 5.1: The seven amines identified as potential stabiliser candidates.

Amine	Formula	Stable	Max. humidity (ppm)
Octylamine (OA)	$C_8H_{19}N$	No	—
1-Ethylhexylamine (EHA)	$C_8H_{19}N$	No	—
N,N-Dimethyldecylamine (DMDA)	$C_{10}H_{21}N$	Yes	>78000
N,N-Dimethyldodecylamine (DDA)	$C_{14}H_{31}N$	Yes	>78000
Tetradecylamine (TDA)	$C_{14}H_{31}N$	No	—
Hexadecylamine (HDA)	$C_{16}H_{35}N$	Yes	900
Dimethyltetradecylamine (DMTDA)	$C_{16}H_{35}N$	No	—

of HDA, immediate insolubility occurred after 30 μ L water (900 ppm) was added. However, DMDA and DDA remained stable even after 2 mL (78,000 ppm) of distilled water was added through titration. Of the two remaining amine candidates, DDA was chosen over DMDA as the preferred amine due to the higher flash point, thereby allowing for safer handling. Furthermore, the addition of DDA allowed for the re-solubilisation of certain samples that had crashed out. Added into the solution as N,N-dimethyldodecylamine with a molecular structure seen in Figure 5.4; if protonated, the amine becomes dimethyldodecylammonium.




Figure 5.4: The structural formula for DDA.

5.2.2 Thousand Standard Test

While DDA was shown to be an effective method to protect the TeBD from trace amounts of H₂O that would inevitably enter the detector, there was still a desire to find optimal synthesis conditions that would minimise the risk of instabilities while explicitly examining the efficacy of the protection bequeathed by the DDA.

To this end, a “Thousand Standard Test” (TST) was designed to served as a standardised and consistent platform to inspect long-term changes across different variations of the TeBD synthesis procedure. Over 150 TeBD synthesis variations were carried out to investigate run conditions that maximised stability in the absence of the amine. These variations differed in BD:Te ratio, synthesis times, LAB preparation, reagent suppliers, humidity exposure, and other environmental conditions or procedural variations.

The TeBD material was synthesised in various different batch sizes^g and diluted 1:1 in LAB for storage and transport. Through a tightly regulated preparatory procedure, this material was then further diluted to 0.5% TeBD in LAB within an oxygen-free nitrogen-sparged glove bag with 30 mL LAB that had been nitrogen sparged for 15 minutes and sent through a 0.2 μ m filter. For each synthesis technique, at least four duplicate samples were created along with an LAB reference. A number of samples were “tortured” through exposure to extreme environmental effects: some were exposed to humidity for 6 hours, while others were produced in an oxygen-rich environment.

A total of 1136 samples were created in this TST, of which 88 were tortured with a 6-hour humidity exposure and 240 tortured through oxygen exposure. The samples

^gThis also allowed this test to examine different scale-up methodologies.

were stored in a light-controlled environment and regularly^h visually inspected for any signs of macroscopic particulates (indicative of a crash-out), discolouration, or another anomalies. A summary of these inspections at the 73-month mark (May 2023) can be seen in Table 5.2 while images of the set-up can be seen in Figure 5.5. As can be seen, stability has been directly confirmed to persist in 86.8% of samples even after 5 years at room temperature if not subjected to torture.

Table 5.2: The status of the Thousand Standard Test after 73 months. Failed samples are those with any visually identifiable crash-outs (precipitate) or discolouration. The standard samples were those not tortured with humidity or oxygen exposure.

Statistic	Standard Samples	Tortured Samples	
		Humidity	Oxygen
Sample Count	808	88	240
Crash Count	107	20	31
Crash rate (%)	13.2	22.7	12.9
Crash Count (with DDA)	0	0	0

Although the TeBD from multiple synthesis conditions were stable in these time frames even without the use of DDA, the most obvious trend was that the most stable solutions were synthesised when the BD:Te ratio was $\sim(2.5-3):1$. Most importantly, no samples with DDA have been seen to exhibit instability in this time frame, regardless of synthesis or torture conditions. In the time since the development of the TST,

^hInspections were originally every 2 months. After 2 years, the inspections were reduced to every 6 months; after 4 years, annually.




Figure 5.5: Photographs of the Thousand Sample Test (left) a rack of samples, showing the four duplicates and LAB for each synthesis variation investigated, and (right) one of several shelves containing TST samples, protected from light exposure with black tarps.

other stability tests using differing techniques have been set up, and all continue to arrive at this same conclusion [215].

As a further benefit to the usage of this amine, the potential de-protonation of the TeBD would lead to a decrease in the fluorescent quenching in the scintillator [215]. As can be seen in Figure 5.6, the DDA:Te molar ratio that maximises the light yield is in the (0.25–0.75):1 range. A DDA:Te molar ratio of 0.25:1 was chosen as the final amine concentration, and demonstrated in small-scale tests to uphold stability while providing an improved light yield.

5.3 The Te Cocktail

The Thousand Standard Test provided an explicit demonstration of the effectiveness of DDA in a 0.2:1 molar ratio with Te towards stabilising TeBD for a time period in excess of five years. The final TeBD molecule to be loaded is now registered as




Figure 5.6: The light yield of 1% TeBD in LAB with 2 g/L PPO. The light yield is normalised to that of unloaded LAB + 2 g/L PPO.

“tellurium 1,2-butanediol hydroxy oxo complexes” and has been assigned the Chemical Abstracts Service Registry Number (CAS RN) 2173121-84-9. As summarised in Table 5.3, the initial “Te cocktail” to be used in the tellurium phase will consist of this molecule dissolved in the enhanced liquid scintillator at a concentration of 0.5% TeBD by mass, and further loaded with 0.25:1 molar ratio of DDA to Te.

5.3.1 Material Compatibility

Although the Te cocktail was explicitly demonstrated to be stable in solution through the thousand standard test, such stability does not preclude incompatibility with the materials within the AV. Potential incompatibilities could lead to cocktail instabilities or degradation of detector hardware. As such, an exhaustive material compatibility

Table 5.3: The Te cocktail composition to be initially deployed in the tellurium phase.

Component	Concentration
LAB	Base
PPO	2.2 g/L
Bis-MSB	3.5 mg/L
BHT	5.4 mg/L
TeBD	0.5% by mass
DDA	0.25:1 molar ratio to Te

test (MCT) was set up to explicitly determine that both the Te cocktail and materials it could be exposed to would remain stable.

The materials used in the MCT was an exhaustive list of all permanent installations (such as the AV) or temporary items (such as deployed calibration equipment) that the Te cocktail could be exposed to during the tellurium fill and tellurium phase; the list of all materials tested are shown in Table 5.4. These materials were sorted into hardware categories, and placed in eight 500 mL jars and two 40 mL vials which were subsequently filled with the Te cocktail. Regularⁱ visual inspections of these jars have not identified any macroscopic particulates, discolouration, or other anomalies in either the cocktail or materials within the 44 months since the MCT began.

ⁱRoughly annual.

Table 5.4: The list of materials investigated in the MCT, with relevant suppliers linked in the footnotes.

Metals	Polymers and Plastics	Miscellaneous
Anodized aluminum	Acrylic	Activated aluminum oxide
Galvanized steel	Cable ties (nylon 66, white)	Fibre reinforced ECTFE ^j
SNO NCD Nameplate ^k	Delrin (black/white)	Silicon carbide from pump
Silver-plated copper	FEP encapsulated O-ring ^l	Quadrasil-AP ^m
Silver-plated VCR ⁿ	FFKM O-ring ^o (black)	
Stainless steel 316/316L	FFKM O-ring ^p (white)	
	HDPE	
	Nylon compression fitting ^q	
	PFA tubing ^r	
	POM (Delrin, rough/smooth)	
	PTFE (tubing, clear/white)	
	PTFE tape ^s	
	Tensylon ^t ropes	
	Tygothane tubing ^u	

^jJohn Brooks, Innomag TB/U-Mag high purity pump parts.^kNickle name plates during the AV construction; some were not removed due to their location.^lMcMaster-Carr.^mAlfa Aesar, QuadraSil Aminopropyl 46303.ⁿSwagelok, SS-12-VCR-2.^oCOG, perlast-g80a.^pTRP Polymer Solutions Ltd.^qMcMaster-Carr.^rCrist Group, Ametek ultra high purity PFA.^sMcMaster-Carr.^tBAE Systems (DuPont since 2012).^uMcMaster-Carr, Tygothane C-210-A.

5.3.1.1 Acrylic Compatibility

A material of particular importance is the compatibility of the Te cocktail with the acrylic that makes up the AV. As such, extensive additional testing was performed to ensure that the AV would remain undamaged after long-term exposures to the Te-loaded liquid scintillator.

Table 5.5: The TST acrylic compatibility test matrix.

Acrylic Type	Duplicates	DDA	Result
Dry	4	No	Crash after 6 months
Wet	4	No	Crash after 3 months
Dry	4	Yes	No crash
Wet	4	Yes	No crash, crazed (damaged) acrylic

A preliminary small-scale acrylic compatibility test was included within the Thousand Standard Test. Four samples were set up using a relatively unstable version of TeBD with a high 4:1 BD:Te ratio, and prepared while exposed to oxygen. DDA was added to two of these samples. As with all examinations within the TST, each sample was duplicated four times and accompanied with an LAB reference prepared in the same manner. As summarised in Table 5.5, 1-cm diameter extruded acrylic spheres were procured from McMaster-Carr and added to the samples. Before addition, some spheres were saturated with 2000 ppm UPW (“wet acrylic”) while others were heated to 40°C in an oven for several days to remove humidity (“dry acrylic”). Samples in cocktails unprotected by DDA were expectedly observed to have crashed within six months; this was accelerated to three months in samples with wet acrylic. Meanwhile, samples using standard acrylic in DDA protected samples continue to be

undamaged after 61 months of inspection. However, wet acrylic in DDA protected samples appeared to have crazed across the entire bulk structure, with visible particulate contamination — presumably from the damaged acrylic — reminiscent of the initial scintillator solvent compatibility tests run in Section 4.2.1.

Table 5.6: The test matrix for the additional acrylic compatibility test.

Acrylic Type	Duplicates	DDA	Result
New Dry	4	No	No crash
New Wet	4	No	Crash after 3 months
Standard	4	No	No crash
New Wet	4	Yes	No crash ^v
New Dry	4	Yes	No crash
Standard	4	Yes	No crash

To investigate this further, a more extensive additional acrylic compatibility test was set up within the SNOLAB clean lab. In this iteration, 1-cm diameter cast acrylic rods procured from McMaster-Carr were sanded and annealed, and cut into 30 1-cm long cylinders. 20 of the cylinders were heated to 40°C; 10 were saturated within water for 7 days to achieve a humidity of 12000 ppm (“new wet” acrylic) while 10 were not (“new dry” acrylic). Meanwhile, 10 of the cylinders were not subjected to any additional stresses (“standard” acrylic). Differing from the initial TST acrylic test, these were then placed within 6 sets of samples in an oxygen-free environment (each with four duplicates) and loaded with 0.5% TeBD synthesised using a more stable process. As summarised in Table 5.6, only wet acrylic samples without DDA

^vThere is also no apparent crazing visually observed either.

have exhibited signs of degradation when prepared in this way, even after 44 months of observation.

The acrylic compatibility has been further examined through multiple accelerated ageing tests, in which acrylic was exposed to highly concentrated ($\sim 25\%$) TeBD in LAB and subjected to various mechanical stresses. These findings found that some minor deterioration may occur in acrylic over $\mathcal{O}(10)$ years, but at levels insufficient to risk the structural integrity of the AV [193, 215].

5.3.2 Loading Technique

With the Te cocktail formula established to maintain stability in the AV for multiple years, attention can be turned to the loading of Te into LAB. There were two particular techniques identified as capable of synthesizing stable TeBD with DDA into LAB at large scales. Throughout this section, the molecular structures noted were determined through a series of NMR or electrospray ionisation mass spectrometry (ESI-MS) measurements.

Type I Loading The first methodology is synthesis through heating, when aqueous TeA (30–50% w/w) is combined in a 1:3 molar ratio with BD. While the initial structures are nearly all Te monomers with 1 or 2 BD ligands, the removal of water through heating, agitation, and nitrogen sparging drives the forward condensation reaction of Equation 5.2. This leads to the formation of more complex oligomerised structures containing up to three Te centers linked together through BD. Once sufficient heating reduces the water concentration below 1000 ppm, the TeBD becomes soluble in LAB. The TeBD molecules created through this process is known as "TeBD-I"; DDA is then added afterwards to provide stabilisation against the reverse hydrolysis reaction.

Type II Loading The second methodology involves the introduction of DDA into the reaction mixture in a 0.5:1:2 DDA:TeA:BD ratio, which can form immediate monomer structures with two bidentate BD attachments in the absence of heat. These “TeBD-II” molecules formed through this cold synthesis approach need not occur in an aqueous environment and is immediately soluble in LAB, provided residual water from the condensation reaction is removed through nitrogen sparging.

5.3.3 Light Yield Quenching

As the tellurium molecules could be considered a form of molecular contamination within the pure liquid scintillator, the loading of tellurium into the scintillator will reduce the overall light yield. This is due to a number of quenching effects, such as a reduction in FRET due to molecular re-arrangement and the immobilisation of π electrons from the electronegativity of the TeBD cocktail. The light yield reduction from the various quenching processes are totalled in an overall “fluorescence” quenching. Depending on the loading type, this fluorescence quenching effect has been measured to reduce the light yield up to 40%, as seen in Figure 5.7.

The superior light yield of the TeBD-II type loading is expected, as the larger electron cloud that surrounds the oligomer molecules allow for accentuated fluorescence quenching. This is supported in tests when TeBD-I Te Cocktails (with DDA) are converted into TeBD-II cocktails through the intentional separation from LAB through water saturation and re-dissolution through nitrogen sparging. The recombined Te cocktail exhibits the same improved light yield, thereby implying that it is the oligomer size (and not other effects from type-I loading) that drive the difference in light yield between the loading types.




Figure 5.7: The light yield of the Te cocktail depending on TeBD synthesis technique and loading amount in LAB with various concentrations of PPO. The light yield is normalised to that of unloaded LAB + 2 g/L PPO.

Notably, it is unlikely that the reduction in measured light yield is due to increased attenuation from material degradation, such as the yellowing effect discussed in Section 4.3.3.1. Indeed, the 10 cm UV-Vis absorption spectra of TST samples were measured with the same method as the AV scintillator samples discussed in Section 4.4.3. As seen in Figure 5.8, measurements of the TeBD loaded at 10% in LAB imply a consistent attenuation length of over $O(10)$ m at 450 nm, comparable to that of un-loaded liquid scintillator. However, these are lower-bound measurements due to high variability in the quality and oxygen exposure of the LAB used. Nevertheless,

absorption measurements of Te-loaded samples using purified scintillator that were stored in dark conditions showed no change on time scales of over three years.^w

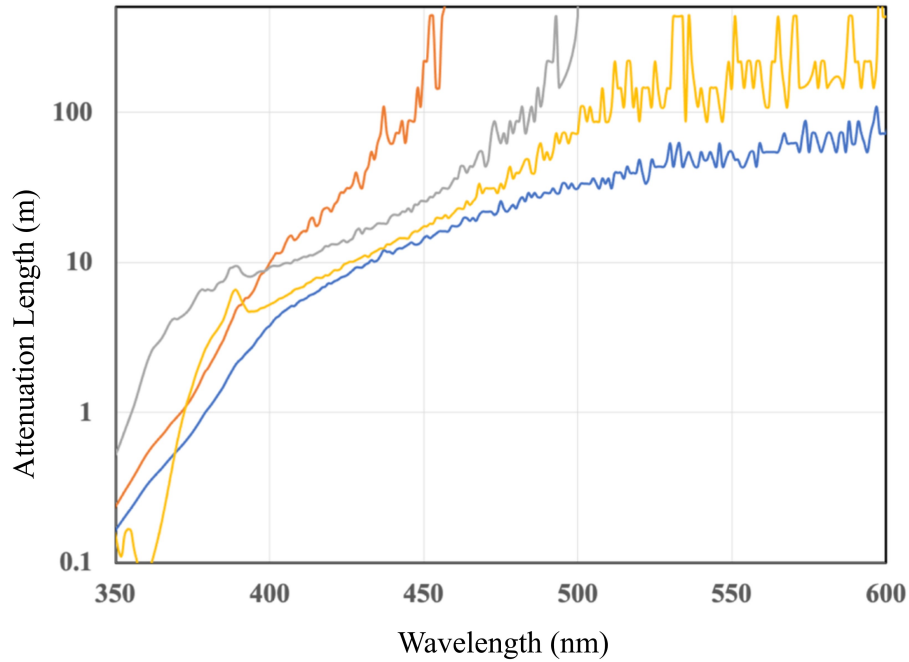


Figure 5.8: The attenuation lengths of various Te cocktail samples loaded at 10% TeBD in LAB with 2 g/L PPO. These measured attenuation lengths are highly variable, presumably based on the quality of the LAB used, and should be understood as lower-bound estimates. In all cases, attenuation length remains over $O(10)$ m at 450 nm.

^wDDA has been observed to possess a UV sensitivity, and degrades in prolonged exposure to UV sources such as sunlight.

5.4 Te Deployment Strategy

As was the case with the pure scintillator, the optical properties are highly dependent on the purity of the Te cocktail. Coupled with the danger posed by the introduction of backgrounds discussed in Chapter 3, the purification and deployment of Te into the liquid scintillator encompasses the final operation that must be conducted on the experiment on the road to $0\nu\beta\beta$.



Figure 5.9: Photographs of (left) the TeA purification plant and (right) the TeBD synthesis plant.

The loading of Te is driven by the tellurium process system, primarily consisting of two chemical plants: the TeA purification plant, and butanediol synthesis plant. Photographs of these plants can be seen in Figure 5.9, and their locations can be seen in Figure 2.2. The TeA is first purified within the TeA purification plant in aqueous

form while the BD is purified in the scintillator plant, before both are combined in the TeBD synthesis plant. A summary of this TeA loading operation is shown in Figure 5.10.

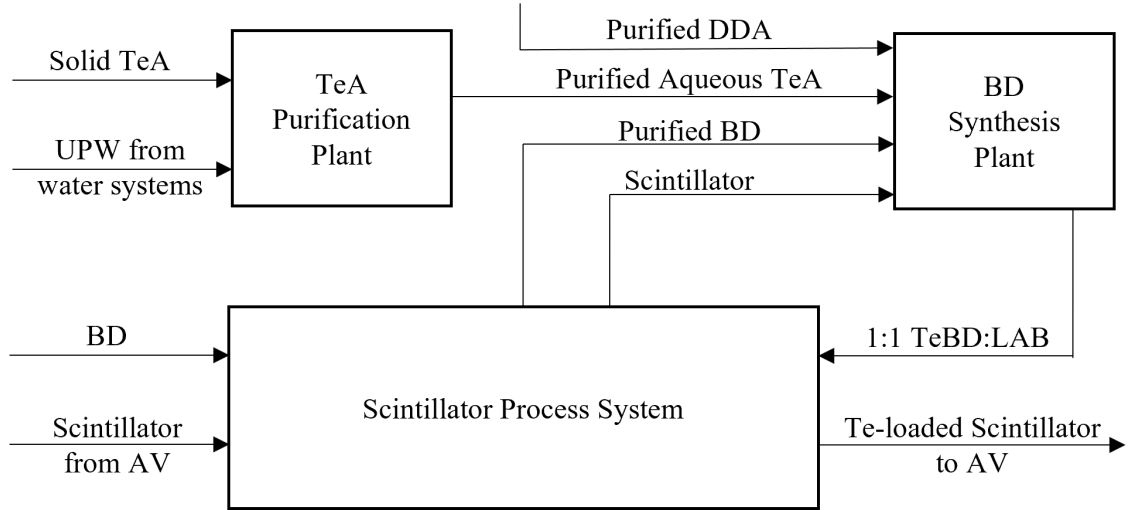


Figure 5.10: Simplified block flow diagram for the Te deployment systems.

The TeA itself was procured in amorphous solid form from ABSCO Ltd and brought underground in 2015; the cosmogenically activated isotopes imparted before arriving in the laboratory as discussed in Section 3.1.1 have since been decaying away, with the rest to be purified away. The inherent U and Th contamination has been measured at the $\mathcal{O}(10^{-11})$ g/g level using inductively coupled plasma mass spectrometry (ICP-MS). Thus, further purification of at least a factor of 10^2 must thus be implemented to bring U to and Th to acceptable levels, and more when also considering cosmogenics.

This purification is performed in 200 kg batches in the underground TeA purification plant; 52 batches are thus required to achieve a concentration of 0.5% within the detector. Within the TeA plant, TeA is first dissolved in hot UPW and filtered

to remove insoluble contamination. As the solubility of TeA in UPW is based on both temperature and pH, the solution can then be recrystallised through the addition of HNO_3 (nitric acid) (“acid recrystallisation”) or reducing the temperature (“thermal recrystallisation”). In either case, the contamination remains in residual supernatant water above the re-solidified TeA, which is removed through vacuum filtration. A single-pass acid recrystallation has been measured to reduce radiogenic and cosmogenic isotopes by a factor of $\mathcal{O}(10^2 - 10^3)$ with only $\mathcal{O}(1)\%$ losses in tellurium content [216]. Two acid recrystallisation passes are expected to achieve the desired background level, which would then be followed by a thermal recrystallisation pass to remove residual nitric acid content.

The TeA is then flowed to the second TeBD synthesis plant, where it is combined with BD purified in the scintillator purification plant C-100 column to form TeBD. This synthesis is a hybridisation of Type I and Type II loading techniques described in Section 5.3.2. The TeA is sent to the TeBD plant from the TeA plant in aqueous form, as the fluidity accommodates transportation. Water is then driven off using a flash chamber (Type I), when DDA — distilled on surface, brought underground, and added directly to the TeBD plant — is applied to complete the solubilisation (Type II). The resulting TeBD is diluted to a 1:1 ratio with LAB (by mass) within the TeBD plant, then sent into the scintillator process system where it can be further diluted through in-line mixing prior to deployment into the detector.

Due to the acidity of TeA, all surfaces that could come into contact with the acid in both plants are constructed with either polypropylene (PP), polytetrafluoroethylene (PTFE), or PFA to minimise leaching of metals into the process fluid. Furthermore, all major process equipment has been pre-leached with acid and rinsed

with UPW. Both spike-tests and pilot-plant runs of the tellurium process system have demonstrated the capability of achieving concentrations of $< 1 \times 10^{-13}$ g/g U, $< 5 \times 10^{-14}$ g/g Th, and $< 7.5 \times 10^{-13}$ Bq/kg ^{60}Co [131], sufficient for the SNO+ $0\nu\beta\beta$ programme.

At the time of this writing, both the TeA purification and TeBD synthesis plants are fully constructed underground. Both plants are now prepared for final stages of commissioning, and SNO+ will soon be ready to start the tellurium fill: the final step towards the search for $0\nu\beta\beta$.

Chapter 6

The Target Out Analysis

Strike wherever the enemy gathers. Other plans we cannot make, for we know not yet how things stand upon the field.

– Theo. 5:5

Through the efforts of many — including the work hitherto discussed within these pages — the SNO+ detector hardware has been fully commissioned, the liquid scintillator deployed, and the tellurium hardware in position. The $0\nu\beta\beta$ search capabilities of the SNO+ experiment have been made ready.

Arrayed against this effort are the various radiogenic backgrounds which threaten the ROI, with each additional unaccounted event setting back the search. Yet here, the SNO+ $0\nu\beta\beta$ search strategy affords the experiment yet another advantage.

As discussed in Section 5.4, the tellurium will be added into the very scintillator that is now deployed into the AV. As such, a direct quantification of the scintillator cocktail backgrounds can be performed prior to the addition of the tellurium target (“target out”). This allows for the $0\nu\beta\beta$ background model described in Chapter 3 to be directly compared against the data, albeit without contributions from the $2\nu\beta\beta$

background and the $0\nu\beta\beta$ signal itself. This confers several advantages not borne in other $0\nu\beta\beta$ search techniques:

- The validity of the background model can be directly tested.
- Unforeseen liquid scintillator backgrounds can be identified independent of the tellurium deployment systems or Te itself.
- A baseline background level for the ROI can be established, and can thus inform on changes made during the addition of Te.
- The lack of $2\nu\beta\beta$ contribution allow backgrounds with energies below the ROI such as the ^{234}Pa to be more easily observed.
- This provides an exercise in analysis techniques that will be used for the upcoming $0\nu\beta\beta$ search, allowing for fine tuning of analysis techniques while “blinded” to the signal, thereby removing the potential for bias.

The target out strategy was thus to use a model of the detector to simulate the energy spectrum of known backgrounds, which were then converted into probability density functions (PDFs). The PDFs of every simulated background component were then fitted onto scintillator phase data through an extended binned maximum likelihood estimation. This allowed for a quantification of the most likely combination of background contributions that corresponded to data taken using the detector.

6.1 Analysis Techniques

In order to perform this search, the extensive simulation and data processing techniques that been developed by the SNO+ collaboration was utilised in the analysis of detector data.

6.1.1 Simulations

Simulations of events within the SNO+ detector are generated using Monte Carlo (MC) experiments: repeated (pseudo-)random sampling to generate estimates of distributions based on input parameters. The parameters themselves depend on an accurate detector model under constant development by the SNO+ collaboration. Informed by studies of the detector hardware, knowledge of underlying physics, and characteristics of the active media (discussed in Chapters 4 and 5), this model provides the bedrock upon which simulations of the detector can be constructed. The simulations can then be compared directly to the data, thus allowing for the capacity to convert raw data into an inference of physical properties.

The MC for SNO+ is performed using a software known as the Reactor Analysis Tool (RAT), a C++ framework initially developed by S. Seibert for the Braidwood Reactor Antineutrino Experiment [217]. Simulations of particles are performed by interfacing with the GEANT4 simulation toolkit [218] and based on the GenericLAND Geant4 simulation (GLG4SIM), a generic liquid scintillator simulation package originally developed by G. Horton-Smith for KamLAND [219].

A version of RAT built for SNO+ is under constant development by the SNO+ collaboration, with simulations hereafter produced using SNO+ RAT v7.0.9^a. Detector parameters are input into RAT using RATDB tables. Some of these RATDB tables are static, as they contain information such as the run number or PMT positions. Others are dynamically adjusted to reflect changing detector state conditions, such as a change in trigger settings or disconnection of a damaged PMT.

^aSimulations used in previous Chapters use a variety of other RAT versions that were appropriate depending on the detector studies and conditions at that time.

MC events are created using generators, which allows for user-defined inputs on parameters including particle types, energies, momenta, positions, and timing. Custom generators have been developed by the SNO+ collaboration for several types of physics processes, including each of the expected backgrounds discussed in Chapter 3. Once a set of initial particle events have been created within the simulation framework using these generators, the information is then passed to GEANT4, which simulates the propagation and decays of the particles and their daughters through the detector medium.

The propagation of every resultant photon is calculated within the simulation. This is dependent on precise geometries of the detector and a full characterisation of the detector media, including the measurements discussed throughout Chapters 4 and 5. Each photon is tracked until they are fully extinguished within the detector medium or impacts a PMT. While a three-dimensional model of the PMTs was developed, an empirical “grey-disc” simplification is used instead, modelling each PMT as a flat circle at the opening of the PMT concentrator. Optical calibrations described in Sec 2.4 are then used to assign reflection and absorption probabilities based on the contact point, angles, and wavelengths for the grey-disc of each PMT.

Within the simulated environment, photons absorbed by the PMT grey-disc generates a PMT pulse, which traverses a simulation of the entire DAQ structure discussed in Section 2.3 to create simulated events. Dynamic detector state parameters from the RATDS tables can be used as an additional input, thus allowing for simulations that are reflective of real detector conditions on a run-by-run basis.

From this point, the simulated “production” data undergoes the same data processing treatment as actual data taken by the detector. However, the production data

retains information on the “true” generated information, which can then be used to inform on analyses or tune the processing techniques.

6.1.2 Data Processing

Whether an event is taken from the detector or generated from MC, it is subjected to a suite of algorithms which convert the raw PMT hits of each event into timing, position, energy, and classification information. These algorithms are dependant on the detector medium due to differences in light propagation; the suite of algorithms used hereafter is called `ScintFitter`, though other suites (such as `PartialFitter`) were used in work discussed in other sections.

Broadly, the `ScintFitter` reconstructs parameters by fitting raw event information (such as `nhits`) or previously reconstructed parameters (such as position) to iteratively tune probability density functions (PDFs) through log likelihood maximisation. Re-tuning and re-creating of these PDFs (“recoordination”) is required when changing the detector model, and done through iterative reconstruction on simulated events in which “true” generated information is already known.

Since the reconstruction of some parameters depends on that of others, the following reconstruction chain is followed:

1. The relative timing of each PMT is used to calculate a global offset between the event trigger time and true event time.
2. The timing information is then correlated against the trigger time of each PMT to geometrically infer the position of the originating event. The `ScintFitter` maximises the timing and position of the event together.

3. The energy of each event is then inferred based on the number of triggered PMTs after accounting for position dependence. The attenuation length, scattering, reflections, and scintillator light yield are folded into the energy calculation.
4. Specific “classification algorithms” can be used to attempt identification of the type of event that has occurred. An example of a classifier used for the work described hereafter is the `alphaBeta212` algorithm developed by Z. Yi [220], which uses hit time information to identify in-window $^{212}\text{BiPo}$ events.

In general, higher energy events correspond with higher detector nhits, which in turn improves the log likelihood calculation for time, position, and energy. Therefore, the accuracy of the reconstruction — and thereby the resolution of the detector — is energy-dependent. In some cases, events can be erroneously mis-reconstructed. Along with instrumental (electronic) noise, mis-reconstructed events present a background on all data analyses. As such, a number of low-level tags are applied to the data in an attempt to reject these backgrounds, which can be flagged during data analysis to cut these events. While a full description of these tags can be found in [221, 222], a few commonly occurring ones include,

- **Fit Validity:** to ensure that all likelihood fits within `ScintFitter` converged appropriately,
- **Cleaned Nhits:** which removes faulty channel and PMT hits from the total event nhits,
- **Corrected Nhits:** which normalises `nhitsCleaned` to account for the number of online channels in each particular run, thus allowing for a direct comparison between runs, and

- Data Cleaning: a suite of cuts based on PMT hit and timing information, applied before reconstruction in an attempt to mitigate instrumental backgrounds. Each event is subjected to a suite of tests, which iterate a bit number when passed. The combined bits then form a hexadecimal number which can be used with a bitmask to classify events during data analysis; the bitmask used in the analyses hereafter is `0x2100000042C2`, though others were used in the analyses discussed in other sections.

MC (“production”) and real (“processed”) data that have undergone reconstruction and data cleaning are saved in identical ROOT file formats [223]. This allows for analysis techniques and infrastructure using the ROOT data analysis framework to be interchangeable between MC and data events.

6.2 Data Selection

The processed data used within this analysis encompasses the first set of scintillator phase data taken by SNO+, with a run range between 300000–306498 corresponding to the time period between 29 April 2022 – 27 November 2022. Although 155.95 days of data were taken, only “gold quality” runs were used in the analysis. The full run selection criteria are found in [224], with the most important being the use of nominal “physics” trigger settings while all electronics crates are operational. The live-time from the dataset was thus reduced to 97.72 days across 2364 runs.

Run number 300000 started at 3:33 pm, 3 minutes after the last drop of scintillator recirculated into the detector. Although a poetic coincidence to start the scintillator phase, the data was freshly contaminated with ^{222}Rn from C-300, as discussed in Section 4.3.3.2. However, with recirculation complete and the detector sealed off,

the ^{222}Rn began to cool at a half-life of 3.82 days until it returned to a rate that is supported by secular equilibrium.

This ^{222}Rn contamination within the detector was tracked by using tagged $^{214}\text{BiPo}$ using the coincidence cuts found in Table 6.1. When used on run-by-run MC without a Δt cut, a cut efficiency of 91.1% was achieved for 1.5–5 MeV range, but 99.96% for the ROI, as can be seen in Figure 6.1. The Δt cut itself has an efficiency of 97.03% for a total $^{214}\text{BiPo}$ cut efficiency of 96.99% in the ROI.

Table 6.1: The cuts used to tag $^{214}\text{BiPo}$ coincidences during the scintillator phase.

Parameter	Cut
^{214}Bi correctedNhits	[400, 1700]
^{214}Po correctedNhits	[160, 400]
Δr	[0, 1.3] m
Δt	[3690, 1×10^6] ns
fitValid	==1
DC bitmask	0x2100000042C2




Figure 6.1: The results of applying $^{214}\text{BiPo}$ coincidence tagging algorithm to MC. The full $^{214}\text{BiPo}$ MC is shown in black, with tagged ^{214}Po in red and tagged ^{214}Bi in blue.

The run-by-run $^{214}\text{BiPo}$ count rate for this dataset can be seen in Figure 6.2. The first 368 runs were cut from the dataset; the $^{214}\text{BiPo}$ rate of the remaining 1996 runs is shown in the inset of Figure 6.2 and indicate a relatively stable ^{222}Rn rate between runs. By removing these runs, the live-time of the dataset was further cut to 78.86 days. A summary of the data selection can be found in Table 6.2.




Figure 6.2: The run-by-run tagged $^{214}\text{BiPo}$ counts during the scintillator phase dataset. Only counts on “gold” runs are shown. The region shown in red features a $^{214}\text{BiPo}$ rate dominated by decays from ^{222}Rn and have been excluded from the dataset. As shown in the inset plot, the remaining runs have a relatively constant rate of $^{214}\text{BiPo}$.

Table 6.2: The live times following each data selection cut for a detector uptime of 155.95 days. The extended muon follower cut was only implemented in the final fit described in Section 6.5.

Cut	Live time remaining (days)
Gold Quality Runs	97.72
Radon cool-off period	78.86
Extended muon follower cut	73.31

After data cleaning was performed on the dataset, the energy scale of the MC was tested by comparing the nhit spectrum of tagged ^{214}Po from the entire dataset to ^{214}Po generated run-by-run using MC. As can be seen in Figure 6.3, both the raw nhits (before data cleaning) and the corrected nhits (after data cleaning) were compared against simulation. The important correction applied was to match the number of operating PMTs between data and MC on a run-by-run basis, which rescaled the raw nhits into corrected nhits. Following this correction, the offset between the simulation and data peaks was 0.22%. After implementing a calibration factor of 0.85 informed by the optical model of the detector constructed by S. Ricetto [225], the data and MC agree on an nhit-energy conversion of roughly 273.5 nhit/MeV.

6.3 Analysis Infrastructure

With the full dataset now selected, simulations can be generated on a run-by-run basis for each of the expected backgrounds discussed in Chapter 3. The list of backgrounds generated are seen in Table 6.3.




Figure 6.3: A preliminary verification of the energy scale performed on this dataset through the comparison between (red) tagged ^{214}Po in the data and (blue) the true ^{214}Po nhit spectrum generated from MC using (left) raw nhits without data cleaning, and (right) corrected nhits after full data cleaning. After all corrections are applied, the data-MC offset is less than 0.22%.

With sufficient MC events, it would be possible to build a smooth PDF from an MC histogram with increasingly small bins or through averaged shifted histogram techniques. However, considering that a PDF is desired for each background contribution for each of the 1996 runs, generating sufficient events would be computationally resource limited.

Table 6.3: The backgrounds for which MCs were generated for use in the fit model.

Background Class	Generated MC
Solar Neutrinos	ν_e ES
Solar Neutrinos	ν_μ ES
^{238}U	$^{214}\text{BiPo}$
^{238}U	^{234m}Pa
^{232}Th	$^{212}\text{BiPo}$
^{232}Th	^{226}Ac
^{232}Th	^{208}Tl
Alpha-n	$^{13}\text{C}(\alpha, n)^{16}\text{O}$
Externals	^{208}Tl (Ropes)
Externals	^{208}Tl (Water)
Externals	^{214}Bi (Ropes)
Externals	^{214}Bi (Water)

6.3.1 Kernel Density Estimation

One way to create a smooth PDF from limited MC events is through Kernel Density Estimation (KDE). A KDE considers each independent event as an individual distribution; a smoothed function is then extrapolated by stacking the distributions on top of each other. This is illustrated in Figure 6.4. As the desired function is univariate, the KDE function built out of n events is explicitly defined as [226],

$$\hat{f}_0(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - t_i}{h}\right), \quad (6.1)$$

where h is the smoothing parameter (“bandwidth”) and t_i is the central value of data point i . The “Kernel” K is itself the distribution used to describe each individual event, of which a Gaussian is the most accepted choice for univariate functions [227]. In the Gaussian kernel choice,

$$K(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}, \quad (6.2)$$

which has an optimal bandwidth of,

$$h = \left(\frac{4}{3}\right)^{1/5} \sigma n^{-1/5}, \quad (6.3)$$

corresponding to the value when the mean integrated squared error is minimised.




Figure 6.4: An illustration of the KDE smoothing technique. By turning 6 independent points (black dots) into individual distributions, the total distribution was summed to form a smoothed PDF.

A KDE PDF was constructed for each background on a run-by-run basis. The run-by-run background model is thus the combined PDF of the various background contributions.

6.3.2 Binned Extended Maximum Likelihood Estimation

The individual PDFs that make up the combined model are each scaled based on a relative parameter. For a model constructed of m PDFs, this set of parameters can be listed as,

$$\vec{\theta} = \{\theta_1, \theta_2, \dots, \theta_m\}. \quad (6.4)$$

Maximum Likelihood Estimation (MLE) is a statistical method that can be used to find the values of each parameter θ such that the combined background model most closely matches the dataset \vec{x} comprised of N independent data points,

$$\vec{x} = \{x_1, x_2, \dots, x_i, \dots, x_N\}. \quad (6.5)$$

In an unbinned MLE, each data point x_i is tested against the combined PDF comprised of the set of parameters that make up $\vec{\theta}$. The probability that the data point corresponds with the fit is,

$$f(x_i; \vec{\theta}). \quad (6.6)$$

The overall likelihood function of the entire dataset is recovered by multiplying the individual probabilities of each event,

$$\mathcal{L}(\vec{x}|\vec{\theta}) = \prod_{i=1}^N f(x_i; \vec{\theta}). \quad (6.7)$$

The parameters can be normalised against the total number of observed events N . In doing so, the normalised parameters are not scaling factors but the expected event contributions from each component PDF. However, normalising in this way turns N into a Poisson random variable that fluctuates around an expected value ν . Thus, an extension to Equation 6.7 must be applied to account for this,

$$\mathcal{L}(\vec{x}|\nu, \vec{\theta}) = e^{-\nu} \frac{\nu^N}{N!} \prod_{i=1}^N f(x_i; \vec{\theta}). \quad (6.8)$$

Thus, the parameters that make up $\vec{\theta}$ can be iteratively changed in an attempt to find the set that maximises $\mathcal{L}(\vec{x}, \vec{\theta})$.

This process is computationally expensive when used in large datasets as in this analysis. There are two improvements that can be immediately applied to the MLE to increase the process feasibility.

First, the computation can be performed in the logarithm of Equation 6.8, thereby converting exponentials into factors and multiplicative operations into summations,

$$\ln \mathcal{L}(\vec{x}|\nu, \vec{\theta}) = -\nu_N \ln \nu + \sum_{i=1}^N \ln f(x_i; \vec{\theta}). \quad (6.9)$$

Second, the data can be binned into a histogram to expedite this process, reducing the number of summations from N observed events to the number of bins B . In this case, it is not the set of \vec{x} that is tested but the set of \vec{n} ,

$$\vec{n} = \{n_1, n_2, \dots, n_i, \dots, n_B\}, \quad (6.10)$$

where n_i is the number of events within bin i . In this binned case, the probability of an event to appear in bin i is,

$$P_i(\vec{\theta}) = \int_{x_i^\downarrow}^{x_i^\uparrow} f(x_i; \vec{\theta}), \quad (6.11)$$

where x_i^\downarrow and x_i^\uparrow are the lower and upper boundaries of bin x_i , respectively. Assuming a fixed N , the likelihood follows a multinomial distribution,

$$\mathcal{L}(\vec{n}|\vec{\theta}) = N! \prod_{i=1}^B \frac{P_i(\vec{\theta})^{n_i}}{n_i!}. \quad (6.12)$$

As in the unbinned case, Equation 6.12 can be extended to be,

$$\mathcal{L}(\vec{n}|\nu, \vec{\theta}) = e^{-\nu} \frac{\nu^N}{N!} N! \prod_{i=1}^B \frac{P_i(\vec{\theta})^{n_i}}{n_i!}. \quad (6.13)$$

Taking $N = \sum n_i$ and $\nu = \sum \nu_i = \sum P_i(\vec{\theta})\nu$, Equation 6.13 can be reduced to,

$$\mathcal{L}(\vec{n}|\nu) = \prod_{i=1}^B e^{-\nu_i} \frac{\nu_i^{n_i}}{n_i!}. \quad (6.14)$$

This form is convenient in understanding that the binned likelihood is a product of the bin-by-bin Poisson probabilities of observing n_i events in bin i given an expected value of ν_i . Maximising this probability is then computationally determined. As before, this can be converted to a log likelihood to further improve computational performance.

6.3.3 Fit Verification

In order to determine whether the analysis framework has been properly implemented, a model comprised of every background KDE was implemented using conditions from ten arbitrary scintillator phase runs. An arbitrary number of background counts were then generated from each MC background for the same ten runs. The model was then fit to the data on a run-by-run basis using the binned extended MLE method discussed.^b The fits were then summed, as is shown in Figure 6.5. As can be seen, the fit was able to converge on the dataset while recovering the expected number of counts.




Figure 6.5: A verification of the analysis framework, performed by applying the fit to a fake MC dataset.

^bHereafter, all fits should be assumed to use the binned extended MLE methods unless otherwise specified.

6.4 The Fit

With the dataset chosen and analysis infrastructure verified, the fit can now be performed on the data itself. For the first look of this scintillator phase data, a 1–7 MeV energy region and 4-m FV was chosen. Recalling that the ROI is defined as [2.42–2.45] MeV in a 3.3-m FV, these parameters extend outside the ROI, allowing for the first look to encompass a broad snapshot that including the surrounding position and energy regions. In order to reduce covariances in the fit, the MC distributions from the various external backgrounds were scaled to their relative predicted rates and combined into a single fit. This treatment was also performed on the ν_e and ν_μ components of the ^8B solar neutrino fit.

The data was sorted into 0.1 MeV bins prior to fitting, with the original parameter assumptions being the predicted rate contributions taken from [166] and summarised in Chapter 3. The contributions from the solar and alpha-n rates were fixed to expected values; the former is constrained by known physics discussed in Chapter 3, while the latter from an accounting of ^{210}Po decays within the detector provided by S. Riccetto.

The first look at the scintillator fitted with the background model is shown in Figure 6.6.




Figure 6.6: The first look at the scintillator phase data between 1.5–7 MeV, fitted with the background model. The data shown is within a 4-m FV.

From this fit, there are three unexpected anomalies:

- There is an abnormally high rate of events at high energies unexplained by the solar fit.
- The fit is in excess at 2.5–3 MeV while lacking in the 2–2.5 MeV region. This implies that there is an excess in the 2–2.5 MeV region, for which the fit is trying to compensate by boosting the $^{214}\text{BiPo}$ contribution.
- The ^{234m}Pa rate within this energy region is higher than that of the $^{214}\text{BiPo}$, despite the boosted contribution.

6.4.1 High Energy Events

The tagged $^{214}\text{BiPo}$ events were then removed from the data; the new fit can be seen in Figure 6.7.




Figure 6.7: The 1.5–7 MeV 4-m FV data fitted after removal of tagged $^{214}\text{BiPo}$.

Many of the high energy events were caught within this cut despite the fact that the $^{214}\text{BiPo}$ spectrum ends under 4 MeV. This implies that the high energy events are due to either a mis-reconstruction of the event energy or incomplete flagging of instrumental events during the data cleaning process. Ultimately, it was discovered that the latter was true; some instrumental events had passed the data cleaning regimen and were being reconstructed despite the fact that a majority of their raw

nhits were flagged as erroneous and removed. These excess events therefore had a high ratio between the nhit and nhitCleaned counts. To discriminate against them, a “hit-ratio” classification can be defined as,

$$R_H = \frac{\text{nhits} - \text{nhitsCleaned}}{\text{nhits}}. \quad (6.15)$$

As seen in Figure 6.8, numerous events in the $^{214}\text{BiPo}$ spectrum had a hit ratio of $R_H > 0.2$, including every event with energies over 4 MeV. When implementing the $R_H > 0.2$ cut on the fit, the high energy anomalies were removed from the dataset as shown in Figure 6.9; remaining high energy events are well described by the expected solar neutrino rate. This hit-ratio cut was used in all subsequent investigations.

6.4.2 2.2 MeV Excess

The 2.0–2.5 MeV excess appears to persist following the cut, indicating that the R_H was not responsible for this unexpected background source.

To remove events in the region that could hinder this investigation, the $^{214}\text{BiPo}$ events were removed from the dataset. For this cut, the coincidence window was increased from $[3.6 \mu\text{s}, 1 \text{ ms}]$ to $[3.6 \mu\text{s}, 1.3 \text{ ms}]$. Furthermore, a $^{212}\text{BiPo}$ cut was also implemented, including both an out-of-window coincidence cut and in-window classifier cut; the cut parameters are shown in Table 6.4. It was determined that 3.7% of $^{214}\text{BiPo}$ events were tagged in both $^{214}\text{BiPo}$ and $^{212}\text{BiPo}$ cuts; the $^{214}\text{BiPo}$ cuts were performed first and removed from the dataset prior to implementation of $^{212}\text{BiPo}$ to ensure that the data would not be over-cut.

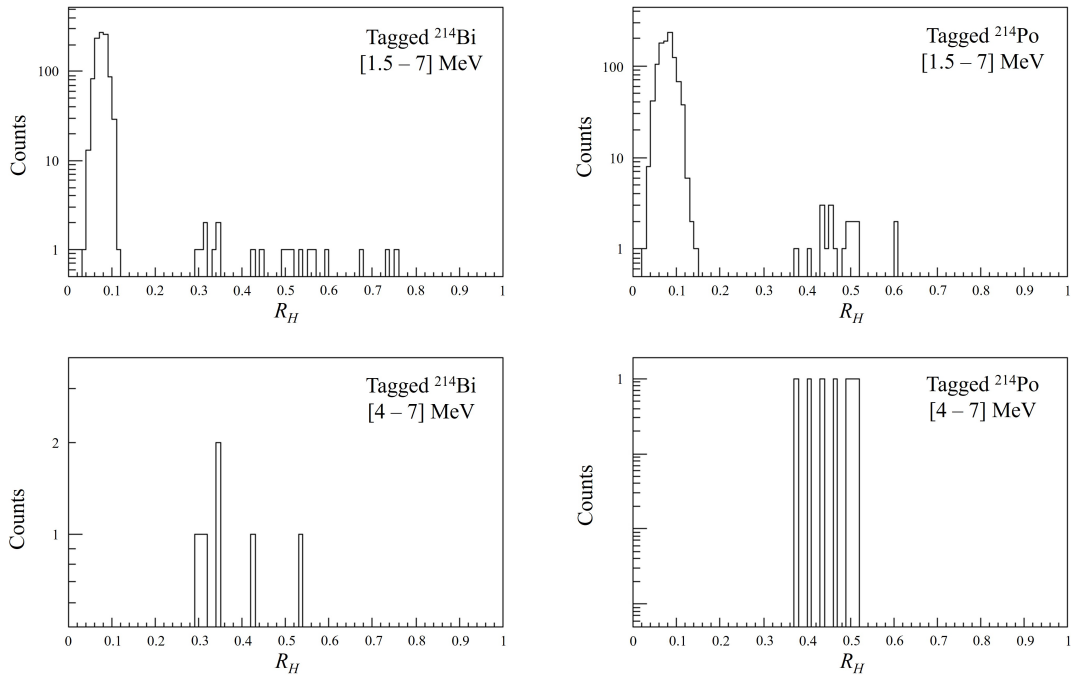


Figure 6.8: The hit-ratio of (left) tagged ^{214}Bi and (right) ^{214}Po . The top plots show the 1.5–7 MeV range, while the bottom show a 4–7 MeV range where no $^{214}\text{BiPo}$ events should occur. All high energy events have a hit-ratio of over 0.2.

The excess is reminiscent of the 2.2 MeV signal that occurs when a neutron is thermalised and captured by hydrogen. To investigate this possibility, the prompt and delayed components of the alpha-n PDFs were separated into independent fits in the model. While the prompt event is a variety of scattering recoils between the neutron and surrounding particles, the follower event corresponds with a 2.2 MeV peak.

The fit was performed on the 1.5–7 MeV 4-m dataset after removal of tagged $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ events. The fit kept the alpha-n prompt fixed at the expected rate, but allowed the alpha-n neutron follower fit to float; the results are shown in Figure 6.10. As can be seen, allowing the neutron rate to float results in an excellent




Figure 6.9: The 1.5–7 MeV 4-m FV data fitted after implementation of the hit-ratio cut. In this fit, the $^{214}\text{BiPo}$ events were not removed to demonstrate the R_H cut effectiveness

fit to the data. However, this implies that the neutron rate required to reconcile this fit was $\mathcal{O}(100)$ higher than expected.

6.5 Improved Fit Results

In order to mitigate contributions from externals to take a closer look at this outstanding 2.2 MeV mystery, a tighter FV was used. As such, the fit was performed on the same 1.5–7 MeV dataset with the tighter canonical ROI FV of 3.3-m. In this fit, both the $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ cuts were applied to the data. Additionally, as IBD

Table 6.4: The cuts used to tag $^{212}\text{BiPo}$ events during the scintillator phase.

Parameter	Cut
^{212}Bi correctedNhits	[400, 810]
^{212}Po correctedNhits	[230, 830]
Δr	[0,1.3] m
Δt	[400, 2600] ns
fitValid	==1
DC bitmask	0x2100000042C2
alphaBeta212	< -5

Table 6.5: The high sacrifice cut used to tag IBD events during the scintillator phase, with a cut efficiency of 100%.

Parameter	Cut
Prompt correctedNhits	[240, 2150]
Delayed correctedNhits	[480, 650]
Δr	None
Δt	< 2 ms
fitValid	==1
DC bitmask	0x2100000042C2

events in the scintillator can also cause a 2.2 MeV neutron signal, an IBD coincidence cut with an efficiency of 100% and parameters in Table 6.5 was also implemented. This coincidence cut was so strong that it should not be used to select IBD events, but instead reject any possible coincidence event that could result in a neutron follower.




Figure 6.10: The 1.5–7 MeV 4-m FV data fitted after removal of tagged $^{214}\text{BiPo}$ and $^{212}\text{BiPo}$ events. The alpha-n fit was separated into prompt and follower events. The latter, being a neutron that thermalises and produces a signature 2.2 MeV peak, reconciles the fit well at high rates.

This cut removed 922 events across the dataset, but only 25 in the 1.5–3 MeV range. Furthermore, to fully eliminate the chance of muon spallation events, a 20 second dead time^c was applied after any event with ≥ 3000 corrected nhits (corresponding to ~ 11.1 MeV). This muon cut was not limited by the FV; 23978 high energy events were identified, resulting in a further live time reduction to a 73.31 days.

With no more data cuts to make in the 2–3 MeV region, a final iteration of the fit was performed on the data. This fit is shown in Figure 6.11, with fit parameters in Table 6.6.

^cData taken during the dead time period is cut from the analysis.




Figure 6.11: The 1.5–7 MeV 3.3-m FV data fitted after removal of tagged $^{214}\text{BiPo}$ and $^{212}\text{BiPo}$ events. The alpha-n fit was separated into prompt and follower events. The latter, being a neutron that thermalises and produces a signature 2.2 MeV peak, reconciles the fit well at high rates.

6.6 The Target Out Results

A majority of the events in the 1.5-7 MeV energy region within the 3.3-m-FV data set is well understood by the expected background contributions. The anomalous high energy events above 4 MeV were defeated through the introduction of a hit-ratio cut of $R_H > 0.2$. In the final fit, high energy events are well described by the expected solar neutrino rate. Similarly, the 3–5 MeV range is well described by the ^{208}Tl spectrum.

Table 6.6: Target Out fit parameters for the 1.5–7 MeV energy range within a 3.3-m FV.

Background	Fit Component	Parameter
Alpha-n	Fixed	1.80
Unknown Neutrons	Floated	124 ± 19
^8B	Multi-component, Fixed	19
$^{212}\text{BiPo}$	Floated	8.2 ± 77.5
$^{214}\text{BiPo}$	Floated	0.0013 ± 0.056
Externals	Multi-component, Floated	6.8 ± 10.8
^{234m}Pa	Floated	1763 ± 52
^{208}Tl	Floated	88.0 ± 9.9

On the low energy side, the 1.5–2 MeV range is dominated by ^{234m}Pa . The relatively high ratio between the ^{234m}Pa and $^{214}\text{BiPo}$ counts imply that the ^{238}U chain may not be in secular equilibrium. Indeed, the rates further down the chain appear to be lower, suggesting that the purification of the scintillator may have been more effective at removing ^{226}Ra than ^{238}U .

The most notable mystery uncovered with the target out analysis is the presence of the 2.2 MeV peak. This peak is well described by neutrons within the detector at a rate > 68.9 times greater than expected. An extensive side-band analysis in the 1.9–2.3 MeV range has been performed to investigate the validity of this peak. The events in this range were further subdivided into a 1.9–2.1 MeV bin and a 2.1–2.3 MeV bin. As the lower bin is dominated by ^{234m}Pa while the upper bin dominated by this neutronesque peak, any mis-reconstructions, data cleaning issues, or discriminatory

parameters should be apparent by comparing the event parameters between the two bins. Thus far, every potential raw and reconstructed parameter has been investigated towards this end to no avail. Additionally, this excess appears to be stable over time and uniformly distributed in radius.

Efforts undertaken in collaboration with S. Naugle continue to be pursued to order to understand this mysterious signal that appears precariously close to the ROI. At the time of this writing, a recent summary of these efforts has been collected in [228].

6.7 Projected SNO+ $0\nu\beta\beta$ Sensitivity

In developing the final fit for the Target Out analysis, signals near the ROI capable of being tagged were removed, leaving only four events in the ROI. These four events are summarised in Table 6.7. Two complementary analyses using different event tagging techniques by S. Naugle and D. Gooding have also converged on these same four events within the ROI.

Table 6.7: List of events remaining in the ROI after implementation of analysis cuts.

Run Number	Event ID	Energy (MeV)
301512	7974238	2.52544
301817	3403600	2.42882
304166	2011530	2.54876
305071	10574794	2.42841

For the observed live time of 73.31 days (0.2008 years), this ROI background count rate is doped with the expected $2\nu\beta\beta$ contribution of 0.2430 and the cosmogenic background count rate of 0.02170 to achieve an expected count rate of 4.2647 events

within the ROI for a period of 0.2008 years, corresponding to a rate of 21.24 events/yr.

Using Equation 3.17, this results in a SNO+ experimental sensitivity of,

$$\mathcal{S}_{1/2}^{0\nu} > \begin{cases} 6.1484 \times 10^{25} & \text{after 1 year live time (90\% C.L.),} \\ 1.3749 \times 10^{26} & \text{after 5 year live time (90\% C.L.).} \end{cases} \quad (6.16)$$

Using the background count informed by the quantification of scintillator data, the SNO+ $0\nu\beta\beta$ programme will surpass the current world-leading ^{130}Te $0\nu\beta\beta$ sensitivity limit of $\mathcal{S}_{1/2}^{0\nu} > 2.2 \times 10^{25}$ [99] within 47 days of detector live time with 0.5% loading of Te in the liquid scintillator.

The half-life sensitivity calculated from this analysis demonstrates that the ROI background levels within SNO+ are sufficient to initiate a competitive search for $0\nu\beta\beta$. However, the nature of the target out analysis — an investigation of the backgrounds prior to the addition of Te — inherently excludes a formidable source of expected backgrounds: the trace amounts of U and Th that will enter the detector alongside the Te. In order to suppress the ROI counts these backgrounds would add, the critical task ahead will be to ensure the extreme purity of the Te as it is loaded into the scintillator. Although this first look at the SNO+ scintillator phase data projects an optimistic future for the SNO+ $0\nu\beta\beta$ search, the final sensitivity will ultimately depend on the deployment of Te.

Chapter 7

Concluding Remarks

Don't adventures ever have an end? I suppose not.

– Bil. 2:1

The search for neutrinoless double beta decay continues to be among the most pressing scientific objectives of our time. As there has yet been no hint of $0\nu\beta\beta$ uncovered by those hunting for it, the global community is now rallying behind a new suite of Next Generation experimental programmes in the race to find the elusive signal.

Foremost among those upcoming is the SNO+ experiment, born of the legacy of SNO. The experimental hardware has been improved, revamped, and forged anew to hunt for $0\nu\beta\beta$ signals in ^{130}Te . The experiment has since completed an initial phase filled with ultrapure water, which successfully achieved preparatory goals including the measurement of external background sources, tuning of the detector response, stable operation of all detector systems, and evaluation of the calibration systems.

The water within the detector was then replaced with a newly designed liquid scintillator to boost the light yield of the experiment. Comprised of linear alkylbenzene

and 2.2 g/L diphenyloxazole, this liquid scintillator has been extensively characterised, allowing for the construction of a precise detector model; further improvements to the liquid scintillator are also under way. Furthermore, a novel chemical loading technique allowing for the dissolution of ^{130}Te into the SNO+ liquid scintillator was developed and verified. This technique of loading the tellurium into the deployed scintillator affords the experiment a unique advantage through the target out analysis, which has informed on the detector model and will be used to understand changes to the experiment as the tellurium is deployed. These studies have also been used to optimise the analyses in the upcoming $0\nu\beta\beta$ search.

This search will not be without challenges. Due to the rarity of the signal, even trace contamination may introduce sufficient backgrounds to eliminate the experiment from competition. Significant efforts were thus undertaken to ensure the exceptional purity of the liquid scintillator during deployment. The first look at the fully deployed liquid scintillator has confirmed that these efforts were not in vain. Everything now hinges on the deployment of tellurium.

The studies and efforts recounted within in this thesis have enabled the SNO+ experiment to stand on the precipice of a new $0\nu\beta\beta$ search. The next step will be to take this plunge.

Bibliography

- [1] D. Laërtius and translated by C. D. Yonge, *The Lives and Opinions of Eminent Philosophers*. G. Bell and Sons, LTD., 3rd century CE, translated 1853.
- [2] W. Pauli and translated by L. M. Brown, *Dear radioactive ladies and gentlemen*, Phys. Today **31N9** (1930, translated 1978) 27.
- [3] S. Weinberg, *The Making of the Standard Model*, Eur. Phys. J. C **34** (2004) 5–13, [hep-ph/0401010].
- [4] C.-N. Yang and R. L. Mills, *Conservation of Isotopic Spin and Isotopic Gauge Invariance*, Phys. Rev. **96** (1954) 191–195.
- [5] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes and R. P. Hudson, *Experimental Test of Parity Conservation in β Decay*, Phys. Rev. **105** (1957) 1413–1414.
- [6] S. L. Glashow, *Partial Symmetries of Weak Interactions*, Nucl. Phys. **22** (1961) 579–588.
- [7] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, Phys. Rev. Lett. **13** (1964) 508–509.
- [8] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, Phys. Rev. Lett. **13** (1964) 321–323.
- [9] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, Phys. Rev. Lett. **13** (1964) 585–587.
- [10] A. Salam, *Weak and Electromagnetic Interactions*, Conf. Proc. C **680519** (1968) 367–377.
- [11] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19** (1967) 1264–1266.
- [12] H. Fritzsch, M. Gell-Mann and H. Leutwyler, *Advantages of the Color Octet Gluon Picture*, Phys. Lett. B **47** (1973) 365–368.

- [13] P. A. M. Dirac, *The Principles of Quantum Mechanics*. Clarendon Press, 1930.
- [14] ATLAS collaboration, G. Aad et al., *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. B **716** (2012) 1–29, [1207.7214].
- [15] CMS collaboration, S. Chatrchyan et al., *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, Phys. Lett. B **716** (2012) 30–61, [1207.7235].
- [16] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse and A. D. McGuire, *Detection of the free neutrino: A Confirmation*, Science **124** (1956) 103–104.
- [17] M. Goldhaber, L. Grodzins and A. W. Sunyar, *Helicity of Neutrinos*, Phys. Rev. **109** (1958) 1015–1017.
- [18] L. D. Landau, *On the conservation laws for weak interactions*, Nucl. Phys. **3** (1957) 127–131.
- [19] T. D. Lee and C.-N. Yang, *Parity Nonconservation and a Two Component Theory of the Neutrino*, Phys. Rev. **105** (1957) 1671–1675.
- [20] A. Salam, *On parity conservation and neutrino mass*, Nuovo Cim. **5** (1957) 299–301.
- [21] J. N. Bahcall, *Beta Decay in Stellar Interiors*, Phys. Rev. **126** (1962) 1143–1149.
- [22] J. N. Bahcall, *Solar neutrinos. I: Theoretical*, Phys. Rev. Lett. **12** (1964) 300–302.
- [23] J. N. Bahcall, *Neutrino Opacity 1. Neutrino-Lepton Scattering*, Phys. Rev. **136** (1964) B1164–B1171.
- [24] J. N. Bahcall, *Solar neutrino cross sections and nuclear beta decay*, Phys. Rev. **135** (1964) B137–B146.
- [25] R. Davis, Jr., D. S. Harmer and K. C. Hoffman, *Search for neutrinos from the sun*, Phys. Rev. Lett. **20** (1968) 1205–1209.
- [26] J. N. Bahcall, N. A. Bahcall and G. Shaviv, *Present status of the theoretical predictions for the Cl-36 solar neutrino experiment*, Phys. Rev. Lett. **20** (1968) 1209–1212.

- [27] B. T. Cleveland, T. Daily, R. Davis, Jr., J. R. Distel, K. Lande, C. K. Lee et al., *Measurement of the solar electron neutrino flux with the Homestake chlorine detector*, *Astrophys. J.* **496** (1998) 505–526.
- [28] SAGE collaboration, J. N. Abdurashitov et al., *Solar neutrino flux measurements by the Soviet-American Gallium Experiment (SAGE) for half the 22 year solar cycle*, *J. Exp. Theor. Phys.* **95** (2002) 181–193, [astro-ph/0204245].
- [29] GALLEX collaboration, W. Hampel et al., *GALLEX solar neutrino observations: Results for GALLEX IV*, *Phys. Lett. B* **447** (1999) 127–133.
- [30] KAMIOKANDE-II collaboration, K. S. Hirata et al., *Observation of B-8 Solar Neutrinos in the Kamiokande-II Detector*, *Phys. Rev. Lett.* **63** (1989) 16.
- [31] B. Pontecorvo, *Mesonium and anti-mesonium*, *Sov. Phys. JETP* **6** (1957) 429.
- [32] S. Esposito and N. Tancredi, *Pontecorvo neutrino - anti-neutrino oscillations: Theory and experimental limits*, *Mod. Phys. Lett. A* **12** (1997) 1829–1838, [hep-ph/9705351].
- [33] Z. Maki, M. Nakagawa and S. Sakata, *Remarks on the unified model of elementary particles*, *Prog. Theor. Phys.* **28** (1962) 870–880.
- [34] L. Wolfenstein, *Neutrino Oscillations in Matter*, *Phys. Rev. D* **17** (1978) 2369–2374.
- [35] S. P. Mikheyev and A. Y. Smirnov, *Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos*, *Sov. J. Nucl. Phys.* **42** (1985) 913–917.
- [36] S. P. Mikheev and A. Y. Smirnov, *Resonant amplification of neutrino oscillations in matter and solar neutrino spectroscopy*, *Nuovo Cim. C* **9** (1986) 17–26.
- [37] S. P. Mikheev and A. Y. Smirnov, *Neutrino Oscillations in a Variable Density Medium and Neutrino Bursts Due to the Gravitational Collapse of Stars*, *Sov. Phys. JETP* **64** (1986) 4–7, [0706.0454].
- [38] SUPER-KAMIOKANDE collaboration, Y. Fukuda et al., *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (1998) 1562–1567, [hep-ex/9807003].

- [39] SNO collaboration, 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, [nucl-ex/0204008].
- [40] GNO collaboration, M. Altmann et al., *Complete results for five years of GNO solar neutrino observations*, Phys. Lett. B **616** (2005) 174–190, [hep-ex/0504037].
- [41] F. Kaether, W. Hampel, G. Heusser, J. Kiko and T. Kirsten, *Reanalysis of the GALLEX solar neutrino flux and source experiments*, Phys. Lett. B **685** (2010) 47–54, [1001.2731].
- [42] SAGE collaboration, J. N. Abdurashitov et al., *Measurement of the solar neutrino capture rate with gallium metal. III: Results for the 2002–2007 data-taking period*, Phys. Rev. C **80** (2009) 015807, [0901.2200].
- [43] SUPER-KAMIOKANDE collaboration, K. Abe et al., *Solar Neutrino Measurements in Super-Kamiokande-IV*, Phys. Rev. D **94** (2016) 052010, [1606.07538].
- [44] SNO collaboration, B. Aharmim et al., *Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory*, Phys. Rev. C **88** (2013) 025501, [1109.0763].
- [45] BOREXINO collaboration, G. Bellini et al., *Precision measurement of the ^7Be solar neutrino interaction rate in Borexino*, Phys. Rev. Lett. **107** (2011) 141302, [1104.1816].
- [46] SUPER-KAMIOKANDE collaboration, K. Abe et al., *Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV*, Phys. Rev. D **97** (2018) 072001, [1710.09126].
- [47] ICECUBE collaboration, R. Abbasi et al., *First all-flavor search for transient neutrino emission using 3-years of IceCube DeepCore data*, JCAP **01** (2022) 027, [2011.05096].
- [48] DAYA BAY collaboration, F. P. An et al., *Precision Measurement of Reactor Antineutrino Oscillation at Kilometer-Scale Baselines by Daya Bay*, Phys. Rev. Lett. **130** (2023) 161802, [2211.14988].
- [49] RENO collaboration, G. Bak et al., *Measurement of Reactor Antineutrino Oscillation Amplitude and Frequency at RENO*, Phys. Rev. Lett. **121** (2018) 201801, [1806.00248].

- [50] DOUBLE CHOOZ collaboration, T. Abrahão et al., *Reactor rate modulation oscillation analysis with two detectors in Double Chooz*, JHEP **01** (2021) 190, [2007.13431].
- [51] KAMLAND collaboration, S. Abe et al., *Precision Measurement of Neutrino Oscillation Parameters with KamLAND*, Phys. Rev. Lett. **100** (2008) 221803, [0801.4589].
- [52] K2K collaboration, M. H. Ahn et al., *Measurement of Neutrino Oscillation by the K2K Experiment*, Phys. Rev. D **74** (2006) 072003, [hep-ex/0606032].
- [53] MINOS+ collaboration, P. Adamson et al., *Precision Constraints for Three-Flavor Neutrino Oscillations from the Full MINOS+ and MINOS Dataset*, Phys. Rev. Lett. **125** (2020) 131802, [2006.15208].
- [54] T2K collaboration, K. Abe et al., *Improved constraints on neutrino mixing from the T2K experiment with 3.13×10^{21} protons on target*, Phys. Rev. D **103** (2021) 112008, [2101.03779].
- [55] NOvA collaboration, M. A. Acero et al., *Improved measurement of neutrino oscillation parameters by the NOvA experiment*, Phys. Rev. D **106** (2022) 032004, [2108.08219].
- [56] S. Gariazzo et al., *Neutrino mass and mass ordering: no conclusive evidence for normal ordering*, JCAP **10** (2022) 010, [2205.02195].
- [57] JUNO collaboration, A. Abusleme et al., *Sub-percent precision measurement of neutrino oscillation parameters with JUNO*, Chin. Phys. C **46** (2022) 123001, [2204.13249].
- [58] T2K collaboration, K. Abe et al., *Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations*, Nature **580** (2020) 339–344, [1910.03887].
- [59] HYPER-KAMIOKANDE collaboration, J. Bian et al., *Hyper-Kamiokande Experiment: A Snowmass White Paper*, in *Snowmass 2021*, 3, 2022. 2203.02029.
- [60] DUNE collaboration, B. Abi et al., *Long-baseline neutrino oscillation physics potential of the DUNE experiment*, Eur. Phys. J. C **80** (2020) 978, [2006.16043].
- [61] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz and A. Zhou, *The fate of hints: updated global analysis of three-flavor neutrino oscillations*, JHEP **09** (2020) 178, [2007.14792].

- [62] “NuFit 5.2: Three-neutrino fit based on data available in november 2022.” www.nu-fit.org.
- [63] P. F. de Salas, D. V. Forero, S. Gariazzo, P. Martínez-Miravé, O. Mena, C. A. Ternes et al., *2020 global reassessment of the neutrino oscillation picture*, JHEP **02** (2021) 071, [2006.11237].
- [64] KATRIN collaboration, M. Aker et al., *Analysis methods for the first KATRIN neutrino-mass measurement*, Phys. Rev. D **104** (2021) 012005, [2101.05253].
- [65] PLANCK collaboration, N. Aghanim et al., *Planck 2018 results. VI. Cosmological parameters*, Astron. Astrophys. **641** (2020) A6, [1807.06209].
- [66] T. Asaka and M. Shaposhnikov, *The ν MSM, dark matter and baryon asymmetry of the universe*, Phys. Lett. B **620** (2005) 17–26, [hep-ph/0505013].
- [67] A. M. Knee, D. Contreras and D. Scott, *Cosmological constraints on sterile neutrino oscillations from Planck*, JCAP **07** (2019) 039, [1812.02102].
- [68] L. Feng, R.-Y. Guo, J.-F. Zhang and X. Zhang, *Cosmological search for sterile neutrinos after Planck 2018*, Phys. Lett. B **827** (2022) 136940, [2109.06111].
- [69] ICECUBE collaboration, M. G. Aartsen et al., *eV-Scale Sterile Neutrino Search Using Eight Years of Atmospheric Muon Neutrino Data from the IceCube Neutrino Observatory*, Phys. Rev. Lett. **125** (2020) 141801, [2005.12942].
- [70] MINIBOONE collaboration, A. A. Aguilar-Arevalo et al., *MiniBooNE and MicroBooNE Combined Fit to a 3+1 Sterile Neutrino Scenario*, Phys. Rev. Lett. **129** (2022) 201801, [2201.01724].
- [71] V. V. Barinov et al., *Search for electron-neutrino transitions to sterile states in the BEST experiment*, Phys. Rev. C **105** (2022) 065502, [2201.07364].
- [72] STEREO collaboration, H. Almazán et al., *STEREO neutrino spectrum of ^{235}U fission rejects sterile neutrino hypothesis*, Nature **613** (2023) 257–261, [2210.07664].
- [73] E. Majorana and translated by Luciano Maiani, *Teoria simmetrica dell’elettrone e del positrone*, Nuovo Cim. **14** (1937, translated 2006) 171–184.
- [74] T. Yanagida, *Horizontal gauge symmetry and masses of neutrinos*, Conf. Proc. C **7902131** (1979) 95–99.

- [75] S. Weinberg, *Baryon and Lepton Nonconserving Processes*, Phys. Rev. Lett. **43** (1979) 1566–1570.
- [76] C. M. Ho and R. J. Scherrer, *Anapole Dark Matter*, Phys. Lett. B **722** (2013) 341–346, [1211.0503].
- [77] M. Fukugita and T. Yanagida, *Baryogenesis Without Grand Unification*, Phys. Lett. B **174** (1986) 45–47.
- [78] S. M. Bilenky, J. Hosek and S. T. Petcov, *On Oscillations of Neutrinos with Dirac and Majorana Masses*, Phys. Lett. B **94** (1980) 495–498.
- [79] C. Giunti, *No Effect of Majorana Phases in Neutrino Oscillations*, Phys. Lett. B **686** (2010) 41–43, [1001.0760].
- [80] E. Ma, M. Raidal and U. Sarkar, *Phenomenology of the Neutrino-Mass-Giving Higgs Triplet and the Low-Energy Seesaw Violation of Lepton Number*, Nucl. Phys. B **615** (2001) 313–330, [hep-ph/0012101].
- [81] B. J. P. Jones, *The Physics of Neutrinoless Double Beta Decay: A Primer*, in *Theoretical Advanced Study Institute in Elementary Particle Physics: The Obscure Universe: Neutrinos and Other Dark Matters*, 8, 2021. 2108.09364.
- [82] E. Fermi and translated by F. L. Wilson, *Fermi’s Theory of Beta Decay*, Am. J. Phys. **36** (1934, translated 1968) 1150–1160.
- [83] C. F. V. Weizsacker, *Zur Theorie der Kernmassen*, Z. Phys. **96** (1935) 431–458.
- [84] B. Povh, K. Rith, C. Scholz, F. Zetsche and translated by M. Lavelle., *Particles and Nuclei: An Introduction to the Physical Concepts, Sixth Edition*. Springer-Verlag Berlin Heidelberg, Germany, 2008.
- [85] M. Goeppert-Mayer, *Double beta-disintegration*, Phys. Rev. **48** (1935) 512–516.
- [86] F. Šimkovic, V. Rodin, A. Faessler and P. Vogel, *$0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements, quasiparticle random-phase approximation, and isospin symmetry restoration*, Phys. Rev. C **87** (2013) 045501, [1302.1509].
- [87] A. Barabash, *Precise Half-Life Values for Two-Neutrino Double- β Decay: 2020 Review*, Universe **6** (2020) 159, [2009.14451].
- [88] PARTICLE DATA GROUP collaboration, R. L. Workman et al., *Review of Particle Physics*, PTEP **2022** (2022) 083C01.

- [89] J. Schechter and J. W. F. Valle, *Neutrinoless Double beta Decay in $SU(2) \times U(1)$ Theories*, Phys. Rev. D **25** (1982) 2951.
- [90] G. Benato, *Effective Majorana Mass and Neutrinoless Double Beta Decay*, Eur. Phys. J. C **75** (2015) 563, [1510.01089].
- [91] M. J. Dolinski, A. W. P. Poon and W. Rodejohann, *Neutrinoless Double-Beta Decay: Status and Prospects*, Ann. Rev. Nucl. Part. Sci. **69** (2019) 219–251, [1902.04097].
- [92] M. Berglund and M. E. Wieser, *Isotopic compositions of the elements 2009 (iupac technical report)*, Pure and Applied Chemistry **83** (2011) 397–410.
- [93] H. V. Klapdor-Kleingrothaus, A. Dietz, H. L. Harney and I. V. Krivosheina, *Evidence for neutrinoless double beta decay*, Mod. Phys. Lett. A **16** (2001) 2409–2420, [hep-ph/0201231].
- [94] C. E. Aalseth, F. T. Avignone, A. Barabash, F. Boehm, R. L. Brodzinski, J. I. Collar et al., *Comment on "evidence for neutrinoless double beta decay"*, Modern Physics Letters A **17** (2002) 1475–1478.
- [95] A. M. Bakalyarov, A. Y. Balysh, S. T. Belyaev, V. I. Lebedev and S. V. Zhukov, *Results of the experiment on investigation of Germanium-76 double beta decay: Experimental data of Heidelberg-Moscow collaboration November 1995 - August 2001*, Phys. Part. Nucl. Lett. **2** (2005) 77–81, [hep-ex/0309016].
- [96] GERDA collaboration, M. Agostini et al., *Final Results of GERDA on the Search for Neutrinoless Double- β Decay*, Phys. Rev. Lett. **125** (2020) 252502, [2009.06079].
- [97] CUORE collaboration, C. Arnaboldi et al., *CUORE: A Cryogenic underground observatory for rare events*, Nucl. Instrum. Meth. A **518** (2004) 775–798, [hep-ex/0212053].
- [98] D. J. Fixsen, *The temperature of the cosmic microwave background*, The Astrophysical Journal **707** (nov, 2009) 916, [0911.1955].
- [99] CUORE collaboration, D. Q. Adams et al., *Search for Majorana neutrinos exploiting millikelvin cryogenics with CUORE*, Nature **604** (2022) 53–58, [2104.06906].
- [100] MAJORANA collaboration, N. Abgrall et al., *The Majorana Demonstrator Neutrinoless Double-Beta Decay Experiment*, Adv. High Energy Phys. **2014** (2014) 365432, [1308.1633].

- [101] GERDA collaboration, K. H. Ackermann et al., *The GERDA experiment for the search of $0\nu\beta\beta$ decay in ^{76}Ge* , Eur. Phys. J. C **73** (2013) 2330, [1212.4067].
- [102] EXO-200 collaboration, M. Auger et al., *The EXO-200 detector, part I: Detector design and construction*, JINST **7** (2012) P05010, [1202.2192].
- [103] EXO-200 collaboration, N. Ackerman et al., *The EXO-200 detector, part II: auxiliary systems*, JINST **17** (2022) P02015, [2107.06007].
- [104] KAMLAND-ZEN collaboration, S. Abe et al., *Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen*, Phys. Rev. Lett. **130** (2023) 051801, [2203.02139].
- [105] CUPID collaboration, O. Azzolini et al., *Final result of CUPID-0 phase-I in the search for the ^{82}Se Neutrinoless Double- β Decay*, Phys. Rev. Lett. **123** (2019) 032501, [1906.05001].
- [106] CUPID collaboration, E. Armengaud et al., *New Limit for Neutrinoless Double-Beta Decay of ^{100}Mo from the CUPID-Mo Experiment*, Phys. Rev. Lett. **126** (2021) 181802, [2011.13243].
- [107] R. Arnold et al., *Technical design and performance of the NEMO 3 detector*, Nucl. Instrum. Meth. A **536** (2005) 79–122, [physics/0402115].
- [108] NEMO-3 collaboration, R. Arnold et al., *Final results on ^{82}Se double beta decay to the ground state of ^{82}Kr from the NEMO-3 experiment*, Eur. Phys. J. C **78** (2018) 821, [1806.05553].
- [109] NEMO-3 collaboration, R. Arnold et al., *Results of the search for neutrinoless double- β decay in ^{100}Mo with the NEMO-3 experiment*, Phys. Rev. D **92** (2015) 072011, [1506.05825].
- [110] NEMO-3 collaboration, J. Argyriades et al., *Measurement of the two neutrino double beta decay half-life of Zr-96 with the NEMO-3 detector*, Nucl. Phys. A **847** (2010) 168–179, [0906.2694].
- [111] NEMO-3 collaboration, R. Arnold et al., *Measurement of the $2\nu\beta\beta$ decay half-life of ^{150}Nd and a search for $0\nu\beta\beta$ decay processes with the full exposure from the NEMO-3 detector*, Phys. Rev. D **94** (2016) 072003, [1606.08494].
- [112] NEMO-3 collaboration, R. Arnold et al., *Measurement of the $2\nu\beta\beta$ Decay Half-Life and Search for the $0\nu\beta\beta$ Decay of ^{116}Cd with the NEMO-3 Detector*, Phys. Rev. D **95** (2017) 012007, [1610.03226].

- [113] NEMO-3 collaboration, R. Arnold et al., *Measurement of the double-beta decay half-life and search for the neutrinoless double-beta decay of ^{48}Ca with the NEMO-3 detector*, Phys. Rev. D **93** (2016) 112008, [1604.01710].
- [114] A. S. Barabash et al., *Final results of the Aurora experiment to study 2β decay of ^{116}Cd with enriched $^{116}\text{CdWO}_4$ crystal scintillators*, Phys. Rev. D **98** (2018) 092007, [1811.06398].
- [115] S. Umehara et al., *Neutrino-less double-beta decay of Ca-48 studied by Ca F(2)(Eu) scintillators*, Phys. Rev. C **78** (2008) 058501, [0810.4746].
- [116] MAJORANA collaboration, I. J. Arnuist et al., *Final Result of the Majorana Demonstrator's Search for Neutrinoless Double- β Decay in Ge76*, Phys. Rev. Lett. **130** (2023) 062501, [2207.07638].
- [117] EXO-200 collaboration, G. Anton et al., *Search for Neutrinoless Double- β Decay with the Complete EXO-200 Dataset*, Phys. Rev. Lett. **123** (2019) 161802, [1906.02723].
- [118] CUPID collaboration, O. Azzolini et al., *Final Result on the Neutrinoless Double Beta Decay of ^{82}Se with CUPID-0*, Phys. Rev. Lett. **129** (2022) 111801, [2206.05130].
- [119] C. Augier et al., *Final results on the $0\nu\beta\beta$ decay half-life limit of ^{100}Mo from the CUPID-Mo experiment*, Eur. Phys. J. C **82** (2022) 1033, [2202.08716].
- [120] LEGEND collaboration, N. Abgrall et al., *The Large Enriched Germanium Experiment for Neutrinoless Double Beta Decay (LEGEND)*, AIP Conf. Proc. **1894** (2017) 020027, [1709.01980].
- [121] NEXO collaboration, G. Adhikari et al., *nEXO: neutrinoless double beta decay search beyond 10^{28} year half-life sensitivity*, J. Phys. G **49** (2022) 015104, [2106.16243].
- [122] NEXT collaboration, J. J. Gomez-Cadenas et al., *Present status and future perspectives of the NEXT experiment*, Adv. High Energy Phys. **2014** (2014) 907067, [1307.3914].
- [123] X. Chen et al., *PandaX-III: Searching for neutrinoless double beta decay with high pressure ^{136}Xe gas time projection chambers*, Sci. China Phys. Mech. Astron. **60** (2017) 061011, [1610.08883].
- [124] SUPERNEMO collaboration, R. Arnold et al., *Probing New Physics Models of Neutrinoless Double Beta Decay with SuperNEMO*, Eur. Phys. J. C **70** (2010) 927–943, [1005.1241].

- [125] A. Jeremie, *The SuperNEMO demonstrator double beta experiment*, Nucl. Instrum. Meth. A **958** (2020) 162115.
- [126] S. D. Biller, *Combined constraints on Majorana masses from neutrinoless double beta decay experiments*, Phys. Rev. D **104** (2021) 012002, [2103.06036].
- [127] SNO collaboration, B. Aharmim et al., *Low Energy Threshold Analysis of the Phase I and Phase II Data Sets of the Sudbury Neutrino Observatory*, Phys. Rev. C **81** (2010) 055504, [0910.2984].
- [128] N. D. Scielzo et al., *Double-beta decay Q values of $Te-132$, $Te-128$, and $Te-120$* , Phys. Rev. C **80** (2009) 025501, [0902.2376].
- [129] SNO collaboration, J. Boger et al., *The Sudbury neutrino observatory*, Nucl. Instrum. Meth. A **449** (2000) 172–207, [nucl-ex/9910016].
- [130] SNO+ collaboration, S. Andringa et al., *Current Status and Future Prospects of the SNO+ Experiment*, Adv. High Energy Phys. **2016** (2016) 6194250, [1508.05759].
- [131] SNO+ collaboration, V. Albanese et al., *The SNO+ experiment*, JINST **16** (2021) P08059, [2104.11687].
- [132] T. Kajita, *Atmospheric neutrinos*, Advances in High Energy Physics **2012** (01, 2012) .
- [133] O. C. Allkofer, K. Carstensen, W. D. Dau and H. Jokisch, *The Absolute Cosmic Ray Flux at Sea Level*, J. Phys. G **1** (1975) L51–L52.
- [134] J. L. Autran, D. Munteanu, T. Saad Saoud and S. Moindjie, *Characterization of atmospheric muons at sea level using a cosmic ray telescope*, Nucl. Instrum. Meth. A **903** (2018) 77–84.
- [135] SNO collaboration, B. Aharmim et al., *Measurement of the Cosmic Ray and Neutrino-Induced Muon Flux at the Sudbury Neutrino Observatory*, Phys. Rev. D **80** (2009) 012001, [0902.2776].
- [136] F. Duncan, A. J. Noble and D. Sinclair, *The construction and anticipated science of SNOLAB*, Ann. Rev. Nucl. Part. Sci. **60** (2010) 163–180.
- [137] JNE collaboration, Z. Guo et al., *Muon flux measurement at China Jinping Underground Laboratory*, Chin. Phys. C **45** (2021) 025001, [2007.15925].
- [138] J.-P. Cheng et al., *The China Jinping Underground Laboratory and its Early Science*, Ann. Rev. Nucl. Part. Sci. **67** (2017) 231–251, [1801.00587].

- [139] Y. Suzuki and K. Nakamura, eds., *Frontiers of neutrino astrophysics. Proceedings, International Symposium, Takayama, Japan, October 19-22, 1992*, vol. 5 of *Frontiers science series*, (Tokyo, Japan), Univ. Acad. Pr., 1993.
- [140] SNO collaboration, E. Hallman, “Sudbury Neutrino Observatory Underground Spray Testing Urylon Polyurethane Coatings.” SNO STR 93.
- [141] B. M. Ovchinnikov and V. V. Parusov, *Low-background shielding*, *Instruments and Experimental Techniques* **43** (2000) 691–693.
- [142] A. Bialek et al., *A rope-net support system for the liquid scintillator detector for the SNO+ experiment*, *Nucl. Instrum. Meth. A* **827** (2016) 152–160.
- [143] P. Jagam and J. Simpson, *Measurements of th, u and k concentrations in a variety of materials*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **324** (1993) 389–398.
- [144] G. Doucas, N. A. Jelley, M. E. Moorhead, N. W. Tanner, S. Gil, L. McGarry et al., *Light concentrators for the Sudbury Neutrino Observatory*, *Nucl. Instrum. Meth. A* **370** (1996) 579–596.
- [145] T. Kaptanoglu, *Reactor Antineutrinos in the SNO+ Water Phase and Detector R&D for Large-Scale Neutrino Detectors*. PhD thesis, University of Pennsylvania, 2020.
- [146] D. Cowen, T. Ekenberg, J. Klein, F. Newcomer, R. Van Berg, R. Van de Water et al., *The sudbury neutrino observatory electronics chain*, *IEEE Transactions on Nuclear Science* **42** (1995) 925–932.
- [147] C. J. Jillings, R. J. Ford, A. L. Hallin, P. Harvey, R. W. MacLeod, H. B. Mak et al., *The photomultiplier tube testing facility for the Sudbury Neutrino Observatory*, *Nucl. Instrum. Meth. A* **373** (1996) 421–429.
- [148] SNO+ collaboration, M. R. Anderson et al., *Optical calibration of the SNO+ detector in the water phase with deployed sources*, *JINST* **16** (2021) P10021, [2106.03951].
- [149] M. R. Dragowsky et al., *The N-16 calibration source for the Sudbury Neutrino Observatory*, *Nucl. Instrum. Meth. A* **481** (2002) 284–296, [nucl-ex/0109011].
- [150] SNO+ collaboration, M. R. Anderson et al., *Measurement of neutron-proton capture in the SNO+ water phase*, *Phys. Rev. C* **102** (2020) 014002, [2002.10351].

- [151] N. J. Tagg, A. Hamer, B. Sur, E. D. Earle, R. L. Helmer, G. Jonkmans et al., *The Li-8 calibration source for the Sudbury Neutrino Observatory*, Nucl. Instrum. Meth. A **489** (2002) 178–188, [nuc1-ex/0202024].
- [152] SNO+ collaboration, M. Anderson et al., *Search for invisible modes of nucleon decay in water with the SNO+ detector*, Phys. Rev. D **99** (2019) 032008, [1812.05552].
- [153] SNO+ collaboration, A. Allega et al., *Improved search for invisible modes of nucleon decay in water with the SNO+ detector*, Phys. Rev. D **105** (2022) 112012, [2205.06400].
- [154] KAMLAND collaboration, T. Araki et al., *Search for the invisible decay of neutrons with KamLAND*, Phys. Rev. Lett. **96** (2006) 101802, [hep-ex/0512059].
- [155] SNO+ collaboration, A. Allega et al., *Evidence of Antineutrinos from Distant Reactors using Pure Water at SNO+*, Phys. Rev. Lett. **130** (2023) 091801, [2210.14154].
- [156] SNO+ collaboration, M. Anderson et al., *Measurement of the ^8B solar neutrino flux in SNO+ with very low backgrounds*, Phys. Rev. D **99** (2019) 012012, [1812.03355].
- [157] SNO+ collaboration, B. Krar, “Solar Low Background Water Analysis.” SNO+ Internal Document 7401.
- [158] C. Mills, *Improved sensitivity to Δm_{21}^2 by the classification of $^{13}\text{C}(\alpha, n)^{16}\text{O}$ background in the SNO+ antineutrino analysis*. PhD thesis, University of Sussex, 2022.
- [159] W. McDonough and S. s. Sun, *The composition of the earth*, Chemical Geology **120** (1995) 223–253.
- [160] KAMLAND collaboration, A. Gando et al., *Reactor On-Off Antineutrino Measurement with KamLAND*, Phys. Rev. D **88** (2013) 033001, [1303.4667].
- [161] BOREXINO collaboration, M. Agostini et al., *Comprehensive geoneutrino analysis with Borexino*, Phys. Rev. D **101** (2020) 012009, [1909.02257].
- [162] KAMLAND collaboration, S. Abe et al., *Abundances of uranium and thorium elements in earth estimated by geoneutrino spectroscopy*, Geophysical Research Letters **49** (2022) , [2205.14934].

- [163] J. Paton, *Directional Reconstruction in Liquid Scintillator Neutrino Detectors*. PhD thesis, University of Oxford, 2023.
- [164] H.-T. Janka, *Neutrino Emission from Supernovae*, pp. 1575–1604. Springer International Publishing, Cham, 2017. 10.1007/978-3-319-21846-5.
- [165] SNEWS collaboration, S. Al Kharusi et al., *SNEWS 2.0: a next-generation supernova early warning system for multi-messenger astronomy*, New J. Phys. **23** (2021) 031201, [2011.00035].
- [166] SNO+ collaboration, “Expected radioactive backgrounds in SNO+.” SNO+ Internal Document 507.
- [167] SNO+ collaboration, “Double Beta Official Plots. Using Te+DDA Optics.” SNO+ Internal Document 5681.
- [168] Q. R. Ahmad, *Muon Correlated Background at The Sudbury Neutrino Observatory*. PhD thesis, Brown U., 2002.
- [169] S. W. Li and J. F. Beacom, *First calculation of cosmic-ray muon spallation backgrounds for MeV astrophysical neutrino signals in Super-Kamiokande*, Phys. Rev. C **89** (2014) 045801, [1402.4687].
- [170] V. Lozza and J. Petzoldt, *Cosmogenic activation of a natural tellurium target*, Astropart. Phys. **61** (2015) 62–71, [1411.5947].
- [171] T. Kroupová, *Improving the Sensitivity to Neutrinoless Double Beta Decay in SNO+*. PhD thesis, University of Oxford, 2020.
- [172] SNO collaboration, B. Aharmim et al., *Search for hep solar neutrinos and the diffuse supernova neutrino background using all three phases of the Sudbury Neutrino Observatory*, Phys. Rev. D **102** (2020) 062006, [2007.08018].
- [173] BOREXINO collaboration, M. Agostini et al., *Comprehensive measurement of pp-chain solar neutrinos*, Nature **562** (2018) 505–510.
- [174] C. Giunti et al., *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, Oxford, 2007.
- [175] J. N. Bahcall, M. Kamionkowski and A. Sirlin, *Solar neutrinos: Radiative corrections in neutrino - electron scattering experiments*, Phys. Rev. D **51** (1995) 6146–6158, [astro-ph/9502003].

- [176] E. Vitagliano, I. Tamborra and G. Raffelt, *Grand Unified Neutrino Spectrum at Earth: Sources and Spectral Components*, Rev. Mod. Phys. **92** (2020) 45006, [1910.11878].
- [177] H. Ejiri and S. R. Elliott, *Charged current neutrino cross section for solar neutrinos, and background to $\beta\beta(0\nu)$ experiments*, Phys. Rev. C **89** (2014) 055501, [1309.7957].
- [178] CUORE collaboration, D. Q. Adams et al., *Measurement of the $2\nu\beta\beta$ Decay Half-Life of ^{130}Te with CUORE*, Phys. Rev. Lett. **126** (2021) 171801, [2012.11749].
- [179] B. von Krosigk, *Measurement of proton and α particle quenching in LAB based scintillators and determination of spectral sensitivities to supernova neutrinos in the SNO+ detector*. PhD thesis, Technische Universität Dresden, 2015.
- [180] E. Rutherford, *The scattering of alpha and beta particles by matter and the structure of the atom*, Phil. Mag. Ser. 6 **21** (1911) 669–688.
- [181] S. Curran and J. Craggs, *Counting Tubes: Theory and Applications*. Laboratory technique monograph. Academic Press, 1949.
- [182] R. Fernow, *Introduction to experimental particle physics*. Cambridge University Press, Cambridge, 1986, 10.1017/CBO9780511622588.
- [183] J. B. Birks, *The Theory and Practice of Scintillation Counting*. Pergamon Press, Oxford, 1964.
- [184] DAMA collaboration, R. Bernabei et al., *The DAMA/LIBRA apparatus*, Nucl. Instrum. Meth. A **592** (2008) 297–315, [0804.2738].
- [185] D. J. Hart, C. M. Hadad, L. E. Carine and H. Hart, *Organic Chemistry: A Short Course*. Cengage Learning, Boston, 13th ed., 2011.
- [186] F. Elisei et al., *Measurements of liquid scintillator properties for the Borexino detector*, Nucl. Instrum. Meth. A **400** (1997) 53–68.
- [187] S. Yoshida et al., *Light output response of KamLAND liquid scintillator for protons and C-12 nuclei*, Nucl. Instrum. Meth. A **622** (2010) 574–582.
- [188] H. O. Back et al., *Study of phenylxylylethane (PXE) as scintillator for low energy neutrino experiments*, Nucl. Instrum. Meth. A **585** (2008) 48–60.
- [189] E. Cartlidge, *Gran sasso lab to shut down controversial experiments*, Physics World **32** (feb, 2019) 12.

- [190] G. Goldstein and W. Lyon, *Liquid scintillators using 1-methylnaphthalene*, The International Journal of Applied Radiation and Isotopes **15** (1964) 133–137.
- [191] S. H. Song, K. K. Joo, S. H. So and I. S. Yeo, *Feasibility study of a gadolinium-loaded DIN-based liquid scintillator*, J. Korean Phys. Soc. **63** (2013) 970–974.
- [192] A. J. Wright, *Robust Signal Extraction Methods and Monte Carlo Sensitivity Studies for the Sudbury Neutrino Observatory and SNO+ Experiments*. PhD thesis, Queen’s University, 2009.
- [193] D. J. Bartlett, *Quality Assurance Testing, and Sensitivity Studies for the SNO+ Experiment*, Master’s thesis, Queen’s University, 2018.
- [194] C. Buck, B. Gramlich and S. Wagner, *Light propagation and fluorescence quantum yields in liquid scintillators*, JINST **10** (2015) P09007, [1509.02327].
- [195] SNO+ collaboration, M. R. Anderson et al., *Development, characterisation, and deployment of the SNO+ liquid scintillator*, JINST **16** (2021) P05009, [2011.12924].
- [196] J. B. Cumming, S. Hans and M. Yeh, *Improving Light Yield Measurements for Low-Yield Scintillators*, Nucl. Instrum. Meth. A **925** (2019) 1–5, [1810.02885].
- [197] M. C. Chen, *SNO and SNO+*, AIP Conf. Proc. **944** (2007) 25–30.
- [198] A. Bonhomme, C. Buck, B. Gramlich and M. Raab, *Safe liquid scintillators for large scale detectors*, JINST **17** (2022) P11025, [2205.15046].
- [199] H. M. O’Keeffe, E. O’Sullivan and M. C. Chen, *Scintillation decay time and pulse shape discrimination in oxygenated and deoxygenated solutions of linear alkylbenzene for the SNO+ experiment*, Nucl. Instrum. Meth. A **640** (2011) 119–122, [1102.0797].
- [200] B. von Krosigk et al., *Measurement of α -particle quenching in LAB based scintillator in independent small-scale experiments*, Eur. Phys. J. C **76** (2016) 109, [1510.00458].
- [201] H. Wan Chan Tseung, J. Kaspar and N. Tolich, *Measurement of the dependence of the light yields of linear alkylbenzene-based and EJ-301 scintillators on electron energy*, Nucl. Instrum. Meth. A **654** (2011) 318–323, [1105.2100].

- [202] B. von Krosigk, L. Neumann, R. Nolte, S. Röttger and K. Zuber, *Measurement of the proton light response of various LAB based scintillators and its implication for supernova neutrino detection via neutrino-proton scattering*, Eur. Phys. J. C **73** (2013) 2390, [1301.6403].
- [203] E. J. Callaghan, B. L. Goldblum, J. A. Brown, T. A. Laplace, J. J. Manfredi, M. Yeh et al., *Measurement of proton light yield of water-based liquid scintillator*, Eur. Phys. J. C **83** (2023) 134, [2210.03876].
- [204] S. Asahi, *The Rayleigh ratio measurement for SNO+ liquid scintillator and background studies of the internal ropes*, Master's thesis, Queen's University, 2012.
- [205] M. Yu, W. Wu, N. Peng, T. Yu, Y. Ding, Q. Liu et al., *Measurements of Rayleigh ratios in linear alkylbenzene*, Rev. Sci. Instrum. **93** (2022) 063106, [2203.03126].
- [206] A. Shokri and S. Karimi, *A review in linear alkylbenzene (lab) production processes in the petrochemical industry*, Russian Journal of Applied Chemistry **94** (2021) 1546–1559.
- [207] SNO+ collaboration, R. J. Ford, *A Scintillator Purification Plant and Fluid Handling System for SNO+*, AIP Conf. Proc. **1672** (2015) 080003, [1506.08746].
- [208] SNO+ collaboration, J. D. Wilson, *Thermally-driven scintillator flow in the SNO+ neutrino detector*, Nucl. Instrum. Meth. A **1055** (2023) 168430, [2212.00251].
- [209] J. Benziger et al., *The Scintillator Purification System for the Borexino Solar Neutrino Detector*, Nucl. Instrum. Meth. A **587** (2008) 277–291, [0709.1503].
- [210] K. Niaz, F. Khan and M. Ajmal Shah, *Chapter 25 - Analysis of quinonoids*, pp. 749–766. Elsevier, 2020.
<https://doi.org/10.1016/B978-0-12-816455-6.00025-1>.
- [211] A. S. C. Inácio, *Data Analysis of the Water and Scintillator Phases of SNO+ : from Solar Neutrino Measurements to Double Beta Decay Sensitivity Studies*. PhD thesis, Universidade de Lisboa, 2022.
- [212] X. Zhou, Q. Zhang, Q. Liu, Z. Zhang, Y. Ding, L. Zhou et al., *Densities, isobaric thermal expansion coefficients and isothermal compressibilities of linear alkylbenzene*, Phys. Scripta **90** (2015) 055701, [1408.0877].

- [213] JUNO, DAYA BAY collaboration, A. Abusleme et al., *Optimization of the JUNO liquid scintillator composition using a Daya Bay antineutrino detector*, Nucl. Instrum. Meth. A **988** (2021) 164823, [2007.00314].
- [214] B. P. Binks, P. D. Fletcher, M. A. Thompson and R. P. Elliott, *Effect of added diols (glycols) on the emulsion properties of oil, water and surfactant mixtures*, Colloids and Surfaces A: Physicochemical and Engineering Aspects **390** (2011) 67–73.
- [215] D. J. Auty et al., *A method to load tellurium in liquid scintillator for the study of neutrinoless double beta decay*, Nucl. Instrum. Meth. A **1051** (2023) 168204, [2212.12444].
- [216] S. Hans et al., *Purification of telluric acid for SNO+ neutrinoless double-beta decay search*, Nucl. Instrum. Meth. A **795** (2015) 132–139.
- [217] T. Bolton, *The Braidwood reactor antineutrino experiment*, Nucl. Phys. B Proc. Suppl. **149** (2005) 166–169.
- [218] GEANT4 collaboration, S. Agostinelli et al., *GEANT4—a simulation toolkit*, Nucl. Instrum. Meth. A **506** (2003) 250–303.
- [219] G. Horton-Smith, “GLG4sim page.”
<https://www.phys.ksu.edu/personal/gahs/GLG4sim/>.
- [220] SNO+ collaboration, Z. Yi, “In-Window BiPos Backgrounds.” SNO+ Internal Document 7738.
- [221] N. McCauley, *Producing a Background Free Data Set for Measurement of the Charge Current Flux and Day-Night Asymmetry at the Sudbury Neutrino Observatory*. PhD thesis, University of Oxford, 2001.
- [222] J. M. G. Walker, *Study of Invisible Mode Nucleon Decay in the SNO+ Detector*. PhD thesis, University of Liverpool, 2016.
- [223] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, Nucl. Instrum. Meth. A **389** (1997) 81–86.
- [224] SNO+ collaboration, J. Page and E. Caden, “How To Run Selection.” SNO+ Internal Document 7753.
- [225] SNO+ collaboration, S. Riccetto, “RAT optics 2.2g/L.” SNO+ Internal Document 7728.

-
- [226] K. S. Cranmer, *Kernel estimation in high-energy physics*, Comput. Phys. Commun. **136** (2001) 198–207, [hep-ex/0011057].
- [227] B. Silverman, *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, 1986.
- [228] SNO+ collaboration, S. Naugle, “2.2 MeV Bump Update.” SNO+ Internal Document 7911.