

Measuring the ${}^7\text{Be}$ Neutrino Flux From the Sun:
Calibration of the Borexino Solar Neutrino Detector

Steven E. Hardy

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Physics

R. Bruce Vogelaar, Chair

Mark Pitt

Ramaswamy Raghavan

Tatsu Takeuchi

March 31, 2010

Blacksburg, Virginia

Keywords: Neutrino physics, calibration, source development

Copyright 2010, Steven E. Hardy

Measuring the ^7Be Neutrino Flux From the Sun: Calibration of the Borexino Solar Neutrino Detector

Steven E. Hardy

(ABSTRACT)

The Borexino solar neutrino detector is a real-time liquid scintillator detector designed to measure sub-MeV neutrinos. With its unprecedented level of radio-purity, Borexino is poised to provide the most precise measurements to-date of solar neutrino and geo-antineutrino fluxes.

However, in order to reduce the systematic errors to sub-5% levels, the detector must be carefully calibrated to understand, among other things, the position and energy reconstructions. To that end, the Virginia Tech component of the Borexino collaboration has constructed a system for deploying and locating calibration sources within the detector. The system was used in four separate calibration campaigns and deployed numerous sources in almost 300 locations throughout the detector. The data from the calibrations have already resulted in the reduction of several sources of systematic error by a factor of two or more.

With the results from the calibration, the Borexino detector has entered a new era of low-energy, high-precision, neutrino detection.

This work was supported by NSF Grant 0802114

Dedication

To my wife Beth, and parents, Don and Lori. I love you all and will never forget all that you have done for me in my life.

Acknowledgments

I owe a tremendous debt of gratitude to all those that have given their help and support during my graduate career. The work presented in this dissertation would not have been possible without the involvement of a great number of people.

I would first like to thank my advisor, Prof. Bruce Vogelaar for his support and guidance over the last five years. His willingness to let his students explore their own ideas and take on new projects has made my graduate career truly a pleasure. I would also like to thank Prof. Raju Raghavan for the interest he has taken in my professional development and for being a wonderful resource for any physics or chemistry questions.

None of this work could have been accomplished without the staff of the Virginia Tech Physics department. Tina, Betty, Chris, Diane, Tammy, Judy, Nora, and Lisa have always been willing to help, and even forfeit their lunch hour when my plane tickets to New Zealand were mysteriously cancelled. I would also like to thank Roger & Travis for a truly endless stream of computing support over the last six years and for not getting mad when I went through three laptop keyboards in the course of writing this dissertation. I also owe a great deal of thanks to John, Melvin, Scott, Ron, Fred and Norm in the electronics and machine shops for their help and humor. Thanks also to Tom Wortalik in the glass shop for sealing every one of our source vials, and for not losing faith when they started exploding.

I have had the pleasure of working with many wonderful and hilarious students in Room 10 over the years. Thank you to: Brad Williams and Mike Sperry for annoying each other in the most hilarious ways, Matt Joyce for sharing an appreciation of Olympic curling and for always wanting a coffee, Derek Rountree for breaking a radioactive source inside of my lab book, and Russ Mammei for always knowing where to find whatever I was looking for; you have all helped to make days in the lab go by a lot easier. I also owe a great deal of thanks to Henning Back for his help and guidance throughout both my undergraduate and graduate career. A special thank you also to Szymon Manecki for taking the reigns over the past year while I have been writing; you are a pleasure to work with and your ability to develop new ways of ruining rental cars is unparalleled.

I also would like to thank the *Mach4* analysis group, particularly: Richard Saldanha, Ben Loer, and Alvaro Chavarria for always being willing to help and for remaining calm when bugs were pointed out. Thank you also to Kevin McCarty, Prof. Cristiano Galbiati, and Prof. Frank Calaprice for inviting me to join the analysis effort several years ago.

I owe a tremendous thank you to the entire Borexino collaboration for their efforts and help. In particular thank you to Andrea and Aldo Ianni, Augusto Goretti, Yury Suvorov, George Korga, Paolo Lombardi, Sergio Parmeggiano, and Augusto Brigatti for their help during the fabrication and installation of the insertion system. I also would like to express my gratitude to Steve Kidner and Laszlo Papp, without their help, none of this work would have been possible; my stays in Italy would have been a lot more difficult if not for their wives, Cecile and Anita, who helped make Assergi feel more like home.

Thank you to my parents and brother for their support in everything I have done in my life. Many of the figures in this dissertation would look much worse if not for my brother and his Photoshop prowess.

Finally, thank you to my wife Beth for her help and support throughout this entire process. Your patience throughout the *many* trips to Italy and late nights in the lab have made the last six years so much easier.

Contents

1	Introduction	1
2	Neutrino Physics	3
2.1	Neutrino Masses and Mixing	3
2.1.1	Vacuum Oscillations	4
2.1.2	Neutrino Oscillations in Matter	12
2.2	Solar Neutrinos	18
2.2.1	PP chain	21
2.2.2	CNO cycle	25
2.3	Outlook	29
2.3.1	Neutrino Mass	29
2.3.2	Neutrino-less Double Beta Decay	32
2.3.3	The value of θ_{13}	35
2.3.4	Elemental Abundances	36

3	The Borexino Detector	41
3.1	Neutrino Detection & Their Sources in Borexino	42
3.1.1	^7Be Neutrinos	44
3.1.2	CNO and pep neutrinos	50
3.1.3	^8B Neutrinos	51
3.1.4	Geo and supernova neutrinos	53
3.2	Detector Design	61
3.2.1	Structural components	61
3.2.2	Scintillator	67
3.2.3	Photomultiplier Tubes	79
3.2.4	Purification Plants	82
3.3	Data Acquisition and Software	85
3.3.1	DAQ	85
3.3.2	Software Reconstruction	90
3.4	Backgrounds	104
3.4.1	Decay Chains	105
3.4.2	Other isotopes	112
3.4.3	Muons and Cosmogenics	116
3.4.4	CTF	120

4	Motivation For A Calibration	125
4.1	Fiducial Volume Determination	126
4.2	Energy Response	129
4.2.1	Spectral Impacts	131
4.2.2	Electronics Saturation	132
4.3	Other Effects	135
4.3.1	Trigger Threshold	135
4.3.2	Photomultiplier Tube Timing	137
4.4	Design Considerations for a Calibration System	139
4.4.1	Maintaining Purity	139
4.4.2	Accuracy of Position Determination	142
4.4.3	Mechanical Considerations	143
4.4.4	Adopted Design Philosophy	147
5	A Source Imaging System for Borexino	149
5.1	Source Location System Hardware	150
5.1.1	Digital Cameras	150
5.1.2	Camera Housing	152
5.1.3	Cabling	158
5.1.4	Camera System Control	161

5.2	Software Reconstruction of Images	163
5.2.1	Camera System Geometry	164
5.2.2	Transforming Between Simulated and Real Camera Geometries	165
5.2.3	Position Reconstruction From Images	170
5.2.4	Camera Calibration	174
5.3	Analysis of the Vessel Volumes	179
5.3.1	Obtaining the Vessel Shape	180
6	A Calibration Source Deployment System For Borexino	199
6.1	In-Detector Hardware	200
6.1.1	Insertion Rods	201
6.1.2	Tether & Fiber Optic	205
6.1.3	Source Coupler	208
6.2	Glovebox, Cross, and Associated Hardware	210
6.2.1	Glovebox	212
6.2.2	Cross	217
7	Calibration Source Production	225
7.1	Calibration Source Design	225
7.1.1	Containment Vials	226
7.1.2	Design Variations For Different Particle Types	234

7.2	Preparation of the Sources	244
7.2.1	Source Loading Station	244
7.2.2	Production of ^{14}C Sources	249
7.2.3	Production of Rn Sources	255
7.2.4	Gamma Source Production	259
7.3	Characterization of Calibration Sources	262
7.3.1	Quality Control Detector	263
7.3.2	Relative Light Yield Measurement	269
7.3.3	Measurements of Sources Used In Calibrations	278
8	Results From The Calibrations	287
8.1	Performance of The Calibration System	287
8.1.1	On-Axis Calibration	288
8.1.2	Off-Axis Calibrations	291
8.2	Radioactive Contamination Introduced During Calibrations	296
8.2.1	^{238}U Contamination	296
8.2.2	^{232}Th Contamination	304
8.2.3	^{85}Kr Contamination Measurement	309
8.2.4	Contamination In The ^7Be - ^{11}C Valley	318
8.3	Reconstruction of Calibration Data	324

8.3.1	Energy Reconstruction Performance	325
8.3.2	Position Reconstruction Performance	327
8.3.3	Attempts to Resolve the Z offset	332
8.4	Systematic Errors After Calibration	345
8.4.1	Impact on the ${}^7\text{Be}$ Neutrino Analysis	346
8.4.2	Impact on the ${}^8\text{B}$ Analysis	348
8.4.3	Impact on the Geo- $\bar{\nu}_e$ Analysis	349
9	Conclusion	351
A	Process Control for the Source Deployment System	366
A.1	Hardware	366
A.1.1	Electronics	377
A.2	Data Acquisition Hardware	377
A.3	Software	382
B	Electronics Box Routing	393
C	On Axis Calibration Run List	398
D	First Off Axis Calibration Run List	399
E	Second Off Axis Calibration Run List	403

List of Figures

2.1	Depiction of two-neutrino oscillation	8
2.2	Theoretical survival probability of neutrinos in matter as a function of energy	16
2.3	Experimental data points for neutrino survival probability as a function of energy	17
2.4	Graphical representation of the MSW effect	19
2.5	MSW level crossing scheme	20
2.6	Neutrino flux vs. energy for the pp and CNO cycle neutrinos	25
2.7	The pp chain of energy production in the stars	26
2.8	The CNO cycle of energy production in stars	28
2.9	Mass scale of the fermions	30
2.10	Neutrino-less double beta decay Feynman diagram	33
2.11	Comparison of measured ${}^7\text{Be}$ and ${}^8\text{B}$ neutrino fluxes with predictions from different metallicities	40
2.12	Comparison of measured ${}^8\text{B}$ and CNO neutrino fluxes with the predictions from different metallicities	40

3.1	Feynman diagrams for neutral and charged current neutrino-electron scattering	42
3.2	Pictorial representation of the orbit of the earth around the sun	46
3.3	Cyclical variation of the ${}^7\text{Be}$ ν event rate due to the eccentricity of the earth's orbit about the sun	47
3.4	Electron neutrino survival probability data measured by Borexino	53
3.5	Theoretical spectrum of geo-neutrinos incident on terrestrial neutrino experiments	59
3.6	Schematic of the underground portion of the LNGS laboratory	61
3.7	The major components of the Borexino experiment in Hall C	61
3.8	Partial section view of the Borexino experiment	62
3.9	The stainless steel sphere before installation of the PMTs	64
3.10	A view of the PMTs through the three meter port on the SSS	64
3.11	The southern end of the inner vessel assembly	66
3.12	A Teflon diffuser ball used as a point light source for vessel monitoring . . .	67
3.13	A photo of an inner vessel diffuser ball taken by one of the CCD cameras . .	67
3.14	Photo of the vessel assembly process	68
3.15	Test inflation of the inner and outer vessels	68
3.16	Pseudocumene molecule	69
3.17	PPO molecule	69
3.18	DMP molecule	69

3.19	Pictorial representation of σ and π bonding	70
3.20	Energy level scheme for the π electrons of pseudocumene	72
3.21	Absorption and emission spectra of pure pseudocumene	74
3.22	Absorption and emission spectra of a PC + PPO mixture	75
3.23	Emission spectrum of PC, PC+PPO, and the quantum efficiency of the PMTs	76
3.24	Scintillation emission spectra from α and β induced excitation of the Borexino scintillator	77
3.25	Side-view of a Borexino PMT with light concentrator and μ -metal shielding attached	80
3.26	Front view of an 8" ETL-9351 PMT	80
3.27	Drawing and acceptance of the Borexino light concentrators	81
3.28	Schematic of the encapsulation design for the muon veto PMTs in the water tank	82
3.29	Overall view of the purification and filling procedure for Borexino	83
3.30	One of the fourteen inner -detector electronics racks	87
3.31	The dependence of nhits corrected on the number of hit PMTs	96
3.32	Comparison of the tail-to-total ratio and Gatti parameter for source data . .	102
3.33	The ^{238}U decay chain	107
3.34	The decay series of ^{232}Th	110
3.35	Level scheme of the electron capture decay of ^{40}K	115

3.36	Decay scheme of ^{85}Kr	116
3.37	PMT hit pattern for a muon event in Borexino	117
3.38	The PMT support structure for the CTF	122
3.39	Dimensioned section-view of the CTF2	123
4.1	The reconstructed radial distribution of some ^{214}Po events	129
4.2	The status of the PMTs and electronics chains	131
4.3	Plot of the ^{214}Po energy as a function of the Z coordinate	132
4.4	Sample spectral fit of the Borexino data relevant for the ^7Be ν analysis . . .	133
4.5	Results of a simulation to show the effect of the electronics saturation	134
4.6	Simulated decay spectrum from the $Q = 173$ keV β decay of ^{85}Kr	137
4.7	Time alignment of the inner detector electronics after secondary correction .	138
5.1	Location of the seven calibration cameras	152
5.2	Screen shot of the camera configuration page of the calibration software . . .	154
5.3	CAD drawing of the major components of the camera housing	155
5.4	Test photo of the pinhole LEDs	158
5.5	The software interface for controlling the picture-taking operation	162
5.6	Photo of the camera system control box	163
5.7	The ideal camera system geometry	166
5.8	The three possible CCD misalignment angles: roll, pitch, and yaw	167

5.9	The tweak correction of an image	169
5.10	The source diffuser as located by the camera software	174
5.11	A screen shot of the camera calibration process	178
5.12	Comparison of two vessel photos showing the development of haze in October of 2007	182
5.13	Screen shot of the vessel shape analysis process	183
5.14	Exaggerated view of the pitfall of not correcting for the displacement of the vessel origin in the shape analysis	185
5.15	Screen-shot of the shape profile produced by the camera software	186
5.16	Vessel shape data from all of the cameras and the average curve drawn through them	187
5.17	Compilation of vessel shape data	192
5.18	Differential rate of events occurring above the north pole from April of 2008 onwards	193
5.19	The time evolution of the volume of the inner vessel	194
5.20	Simulated stress on an underinflated vessel	196
5.21	Photo of the two points used in the creep analysis	197
6.1	Photo of a rod coupler	202
6.2	CAD drawing of the hinge rod, source, and source coupler	204
6.3	Photo of the source coupler, fiber optic and diffuser	207

6.4	Photo of the tether drum installed on the glovebox	208
6.5	Assembly drawing of the source coupler	210
6.6	The source coupler and neutron source during a source change	211
6.7	CAD drawing of the glovebox panels, top hat, collar, and frame	214
6.8	Photo of the glovebox as seen from the southwest corner of CR4	217
6.9	Exploded CAD view of the cross	218
6.10	Section view of the top flange with the rod, tether, and sweep arm seals . . .	220
6.11	Photo of the internal components of the rod seal installed on an insertion rod	222
6.12	The inside of the cross during a source change	224
7.1	The calibration source attachment used for the CTF 3 calibration	227
7.2	Dimensioned CAD drawing of the calibration source vials	228
7.3	Photo of a calibration source after sealing	231
7.4	Photos of the source attachment process and result	233
7.5	Simulated energy deposit of ^{85}Kr β events in the wall of the vial	236
7.6	Dimensioned CAD drawing of the neutron source holder	240
7.7	Photos of the neutron source loading and attachment	241
7.8	P&ID of the source loading station	247
7.9	Annotated photo of the source loading stations	248
7.10	Results of the energy deposition of ^{14}C events in the vial wall	250

7.11	Photos of the items used to load ^{14}C into the vials	254
7.12	The valve assembly used for loading the radon sources	257
7.13	A closeup view of ^{14}C - ^{222}Rn source before loading	258
7.14	Photos of sources during the loading process	260
7.15	Photo of the source characterization detector	264
7.16	The electronics for the source characterization detector	265
7.17	The vial assembly used for source characterization studies	274
7.18	Energy spectra of quenched and unquenched sources	275
7.19	Fit results of a source exposed to a ^{22}Na source	276
7.20	Experimental data from a radon source exposed to different gamma sources .	278
7.21	The energy spectra of the on and off-axis calibration sources	280
7.22	Comparison of the weak radon source and the off-axis radon source	280
7.23	Energy spectra of sources made with different scintillators	283
7.24	Borexino data acquired from calibration sources	284
7.25	Comparison of the sources used in the first and second off-axis calibrations .	286
8.1	Plot of the source deployment locations	289
8.2	The hardware used for the on-axis test in the CTF	290
8.3	Photo of the on-axis source deployed in the detector	294
8.4	Photo of the AmBe source deployed during the off-axis calibration	295

8.5	Simulated energy spectrum of ^{214}Bi	297
8.6	Plot of the distance between ^{214}Bi and ^{214}Po events from a source run . . .	298
8.7	Plot of the Gatti parameter distribution for ^{214}Po events from a source run .	299
8.8	Energy distribution of ^{214}Po events from the internal contamination of the detector	300
8.9	The rate of ^{214}Bi - ^{214}Po coincidences within the fiducial volume from May 2007 to October 2009	301
8.10	Simulated spectrum of the β decay of ^{212}Bi	305
8.11	The time evolution of the ^{212}Bi - ^{212}Po coincidence rate in the fiducial volume	306
8.12	Energy spectrum of ^{85}Sr source events	311
8.13	TTR parameter distribution of ^{85}Sr source events	311
8.14	Energy distribution of the sixteen candidate ^{85}Kr events	316
8.15	Energy distribution of the sixteen candidate ^{85m}Rb events	316
8.16	Plot of the decay time for ^{85}Kr - ^{85m}Rb coincidence events	317
8.17	Distribution of the coincidence likelihood for the ^{85}Kr - ^{85m}Rb analysis	317
8.18	A <i>simulated</i> energy spectrum and fit for Borexino	320
8.19	Temporal evolution of ^{210}Po and “valley” events inside of the fiducial volume	322
8.20	Relative energy response of Borexino to a ^{222}Rn calibration source as a func- tion of Z	326
8.21	Deviation in X of the reconstructed position of the ^{214}Po events in the ^{222}Rn source from the camera-determined position	329

8.22	Deviation in Y of the reconstructed position of the ^{214}Po events in the ^{222}Rn source from the camera-determined position	330
8.23	Deviation in Z of the reconstructed position of the ^{214}Po events in the ^{222}Rn source from the camera-determined position	331
8.24	Deviation in Z of the reconstructed position of the ^{214}Bi events in the ^{222}Rn source from the camera-determined position	332
8.25	Deviation in Z of the reconstructed position of the ^{14}C events in the ^{222}Rn source from the camera-determined position	333
8.26	Histograms and fits of the X, Y, Z and energy distributions obtained from a ^{222}Rn source at the center	334
8.27	Plots of the dead PMT distribution in the symmetric PMT study	335
8.28	Z coordinate distribution for a ^{222}Rn source using the standard position reconstruction compared with a reconstruction using a symmetric distribution of PMTs	336
8.29	Vessel shape profile before the first calibration and after the last calibration .	338
8.30	Time-of-flight-subtracted arrival times for PMTs grouped by distance	340
8.31	Distribution and fit of the relative weights of the exponentials in the distance dependent PDF	341
8.32	Distribution and fit of the two exponential lifetimes in the distance dependent PDF	342
8.33	Distribution and fit of the resolution and offset for the distance dependent PDF	343
8.34	Results of a comparison of the normal PDF to the experimental PDF	344

A.1	P&ID of the gas control system	370
A.2	The inside of the gas box	376
A.3	Photo of the electronics box	383
A.4	Screen shot of the gas control software	391
A.5	Screen shot of the alarm and operations preset configuration for the gas control system	392

List of Tables

2.1	Sensitivities of neutrino experiments to the squared mass splittings	9
2.2	Best-fit values of the neutrino oscillation parameters	11
2.3	PP and CNO cycle fluxes under the assumption of different solar abundance data sets	37
3.1	CNO neutrino fluxes and errors	50
3.2	Abundances of geo-neutrino producing isotopes in the area around LNGS . .	60
3.3	Physical properties of PC, PPO, and DMP	68
3.4	Parameters of the light produced by the Borexino scintillator for α and β induced excitation	78
3.5	Parameters of the scintillator PDF used by <i>Mach4</i>	99
3.6	Sources of backgrounds and tolerable levels in Borexino	106
3.7	Subset of the ^{238}U decay chain below ^{222}Rn	109
3.8	Isotopes, energies, and half lives of the triple alpha coincidence in the ^{232}Th chain	112

3.9	Half lives, energies, and cross sections for the short lived cosmogenics in Borexino	121
3.10	Radiopurity results from the CTF	124
4.1	Reconstructed radius of the inner vessel as determined by different sources .	130
5.1	Coordinates of the pinholes of the seven calibration cameras	151
5.2	Parameters of the Kodak DC290 cameras used in the source location system	153
5.3	Connection list for the camera control wiring	160
5.4	Calibration parameters for each of the cameras	179
6.1	Characteristics of the calibration system fiber optic	209
7.1	List of the gamma sources used in the calibrations and their energies	261
8.1	^{238}U contamination levels inside of the FV for some relevant time periods . .	302
8.2	^{232}Th contamination levels inside of the FV during some relevant time periods	307
8.3	Energies, locations, time and spatial separation for the nine ^{85}Kr - ^{85m}Rb candidate events	318
8.4	Position reconstruction resolution of ^{14}C , ^{214}Po , and ^{214}Bi events in the calibration sources	333
8.5	The parameters for the distance dependent fits to the distance dependent PDF	343
8.6	Reconstructed positions and widths for a calibration source using the standard and experimental PDFs	345

C.1	Sources, DAQ run numbers, and positions of the sources during the on axis calibration from October 5 through October 10, 2008.	398
D.1	List of DAQ runs and source positions for the off axis calibration in January, 2009	399
E.1	List of DAQ runs and source positions for the second off axis calibration in June, 2009	403
F.1	List of DAQ runs and source positions for the third off axis calibration in July, 2009	407

Chapter 1

Introduction

The work presented herein is focused on the calibration of the Borexino solar neutrino detector using radioactive calibration sources. The dissertation is organized as follows:

- Chapter 2 gives an overview of the physics of solar neutrinos, their oscillations, and the mechanisms of neutrino production in the sun. The chapter concludes with an overview of some of the questions that still remain in the field.
- The Borexino solar neutrino detector is introduced in Chapter 3. The physics potential of Borexino is presented followed by an overview of the detector components, the data acquisition, and the *Mach4* analysis software. A discussion of the main sources of background for the experiment conclude the chapter.
- In Chapter 4, some arguments in favor of performing a calibration are presented. The chapter concludes with a discussion of the design goals and constraints for such a calibration system.
- Chapter 5 introduces the source location system used for finding a source deployed

inside of the detector. The hardware of the system is discussed, followed by an overview of the algorithm used for processing and calibrating the images. The chapter also explains how the cameras are used for evaluating the shape and volume of the inner vessel, a feature that was critically important in determining how to deal with the leak in the inner vessel.

- Chapter 6 presents the hardware used for deploying the calibration sources inside of the detector.
- The development, production and quality control of the radioactive calibration sources is discussed in Chapter 7. Some comparisons between the different sources used in the calibration campaigns conclude the chapter.
- Chapter 8 presents some results from the calibration campaigns. After a review of the performance of the system, results are presented from a detailed analyses of the contamination introduced to the detector during, and after, the calibrations. The impact that the calibrations have had to-date on the systematic errors of the various neutrino measurements is also discussed.
- The dissertation concludes with a summary of the achievements of the experiment and an outlook for its future. Several appendices at the end of this work describe the process control of the source deployment system and several tables give the locations where the sources have been deployed in the calibration campaigns.

Chapter 2

Neutrino Physics

2.1 Neutrino Masses and Mixing

The standard model (SM) of particle physics at the time of this writing assumes neutrinos to be massless, and therefore, can be described as pure flavor states — in other words, an arbitrary neutrino mixture $|\nu\rangle = A_e|\nu_e\rangle + A_\mu|\nu_\mu\rangle + A_\tau|\nu_\tau\rangle$ has only one nonzero A_α coefficient. However, experimental evidence collected in the last half century (first in the solar arena [1], [2], and then from atmospheric sources[3]) revealed a large difference in the neutrino detection rates relative to theoretical calculations — this difference was irreconcilable with the assumption that the neutrino flavor eigenstates were the eigenstates of their propagation Hamiltonian. A “way out” of these contrasting results was provided by Bruno Pontecorvo in the 1960s ([4], [5]) in which he postulated the existence of at least one non zero neutrino mass (and therefore mass eigenstate). If the mass eigenstates form the basis set for the propagation Hamiltonian instead of the flavor eigenstates, then coherent mixing, or *oscillation*, between the flavor eigenstates can occur. If our neutrino detectors on earth are only sensitive to certain neutrino flavors, then this apparent discrepancy can be resolved if the neutrinos have

oscillated to a flavor which is undetectable.

2.1.1 Vacuum Oscillations

For a neutrino produced in a flavor state $|\nu_\alpha\rangle$ where $\alpha = (e, \mu, \tau)$, we can expand it in terms of the mass eigenstates ν_i with the aid of a unitary mixing matrix $U_{\alpha i}$ as:

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^\dagger |\nu_i\rangle \quad \Leftrightarrow \quad |\nu_i\rangle = \sum_\alpha U_{\alpha i} |\nu_\alpha\rangle \quad (2.1)$$

From standard quantum mechanics, the survival probability for mass state i is given by $|U_{\alpha i}|^2$, and for flavor state α it is $|U_{\alpha i}^\dagger|^2$.

In weak interactions, neutrinos are initially produced in a pure flavor state $|\nu_\alpha\rangle$ and we wish to determine the probability that this neutrino changes flavor to state $|\nu_\beta\rangle$ after travelling a distance L . In performing this calculation we expand the flavor states α, β into their mass eigenstates i . Each mass eigenstate contributes its own term, A_i to the overall probability, where this term is given by the product:

$$A_i = U_{\alpha i}^\dagger \Gamma(\nu_i) U_{\beta i} \quad (2.2)$$

The first term represents the probability that the neutrino produced with its analogous lepton α be in the mass state ν_i . The last term is analogous to this one, but for the neutrino detected with production of its lepton β in the mass state i . The middle term $\Gamma(\nu_i)$ is the propagation of the mass eigenstate $|\nu_i\rangle$ travelling the distance L — this is just the quantum mechanical factor from the time dependent Schrödinger equation: $\exp(-iEt/\hbar)$. In the rest frame of the neutrino, this factor is $\exp(-im_i c^2 \tau_i / \hbar)$ — transforming this back to the

laboratory frame, the argument $m_i c^2 \tau_i$ is $Et - p_i L$. We can resolve this further by exploiting the relativistic energy momentum relationship and the fact that neutrino masses are very small ($m_i c^2 \ll E$):

$$Et - p_i L = Et - L \left\{ \sqrt{\frac{E^2}{c^2} - m_i^2 c^2} \right\} \quad (2.3)$$

$$\begin{aligned} &= Et - L \left\{ \frac{E}{c} \sqrt{1 - \left(\frac{m_i c^2}{E} \right)^2} \right\} \\ &= Et - L \left\{ \frac{E}{c} - \frac{m_i^2 c^4}{2E} \right\} \\ Et - p_i L &= E \left(t - \frac{L}{c} \right) + \frac{m_i^2 c^4}{2E} L \end{aligned} \quad (2.4)$$

Note that the first factor in equation 2.4 does not affect the overall phase (all of the terms are constants), so, we neglect it. Inserting the phase into Γ in equation 2.2 and yields:

$$P(|\nu_\alpha\rangle \rightarrow |\nu_\beta\rangle) = \left| \sum_i U_{\alpha i}^\dagger e^{-i \frac{m_i^2 c^4}{2E} L} U_{\beta i} \right|^2 \quad (2.5)$$

Multiplying out equation 2.5 we obtain:

$$\begin{aligned} P(|\nu_\alpha\rangle \rightarrow |\nu_\beta\rangle) &= \delta_{\alpha\beta} - 4 \sum_{i>j} \Re \left(U_{\alpha i}^\dagger U_{\beta i} U_{\alpha j} U_{\beta j}^\dagger \right) \sin^2 \left(\Delta m_{ij}^2 \frac{Lc^3}{4\hbar E} \right) \\ &\quad + 2 \sum_{i>j} \Im \left(U_{\alpha i}^\dagger U_{\beta i} U_{\alpha j} U_{\beta j}^\dagger \right) \sin \left(\Delta m_{ij}^2 \frac{Lc^3}{2\hbar E} \right) \end{aligned} \quad (2.6)$$

Where, in equation 2.6, the value Δm_{ij}^2 is defined to be $m_i^2 - m_j^2$.

An important feature to note is that under the transformation $|\nu_\alpha\rangle \rightarrow |\bar{\nu}_\alpha\rangle$ and $|\nu_\beta\rangle \rightarrow |\bar{\nu}_\beta\rangle$ (i.e. the mirror image of the oscillation of $|\nu_\alpha\rangle$ to $|\nu_\beta\rangle$ under a CP transformation), equation 2.6 maintains the same form, except that the sign of the third argument is reversed. If the mixing matrix U has only real components, then $P(|\nu_\alpha\rangle \rightarrow |\nu_\beta\rangle) = P(|\bar{\nu}_\alpha\rangle \rightarrow |\bar{\nu}_\beta\rangle)$ and CP is conserved. If, on the other hand, U has a complex component to it, then the sign flip for the third component in equation 2.6 means that the survival probabilities for $|\nu_\alpha\rangle \rightarrow |\nu_\beta\rangle$ and $|\bar{\nu}_\alpha\rangle \rightarrow |\bar{\nu}_\beta\rangle$ are different and we have CP violation in the neutrino sector.

Two-Neutrino Mixing

The usual treatment of vacuum oscillations is done for the case of just two neutrino species (the abstraction to three species is trivial) in which we deal with flavor eigenstates (ν_e, ν_x), mass eigenstates (ν_1, ν_2) and the mixing is described by one mixing angle θ . In this regime, the mass eigenstates are given by:

$$(|\nu_e\rangle, |\nu_x\rangle) = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix} \quad (2.7)$$

Using the expression for the mixing matrix, $U_{\alpha i}$, presented in equation 2.7 and inserting it into equation 2.6 we get the oscillation probability for $\alpha \neq \beta$:

$$P(|\nu_\alpha\rangle \rightarrow |\nu_\beta\rangle) = \sin^2(2\theta) \sin^2\left(\Delta m^2 \frac{Lc^3}{4\hbar E}\right) \quad (2.8)$$

$$= \sin^2(2\theta) \sin^2\left(1.27\Delta m^2 \frac{L}{E}\right) \quad (2.9)$$

In equation 2.9, Δm^2 is in units of eV^2/c^4 , L is in km and E in GeV. The first factor in

equation 2.9 sets the overall depth of the oscillation and depends only on the mixing angle θ , the second factor controls the frequency of the oscillation with a characteristic wavelength of:

$$\lambda_{vac} = \frac{4\pi\hbar E}{\Delta m^2 c^3} \quad (2.10)$$

The smallness of θ_{13} allows us to study the oscillation of solar neutrinos on their way to the earth as oscillation between just two neutrino states; we can obtain useful relationships for oscillation probabilities in this case by taking $|\nu_\alpha\rangle \rightarrow |\nu_e\rangle$ and $|\nu_\beta\rangle \rightarrow |\nu_\mu\rangle$:

$$P(|\nu_e\rangle \rightarrow |\nu_\mu\rangle) = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L c^3}{4\hbar E}\right) \quad (2.11)$$

$$P(|\nu_e\rangle \rightarrow |\nu_e\rangle) = 1 - P(|\nu_e\rangle \rightarrow |\nu_\mu\rangle) = 1 - \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L c^3}{4\hbar E}\right) \quad (2.12)$$

Thus, the phenomenon of neutrino oscillations (in vacuum at least) boils down to a quantum interference effect which is governed by the second \sin^2 term in the above equations. The effect is caused by the relative phase difference of the mass eigenstates (which are induced by the relative squared mass differences) as they propagate in time. Neutrinos are born in weak interactions in pure flavor eigenstates whose admixture of $|\nu_1\rangle$ and $|\nu_2\rangle$ mass eigenstates is determined by the mixing parameter θ — *these admixtures do not change*. The term which oscillates is the probability to observe a given neutrino state as a $|\nu_e\rangle$ for instance: $|\langle\nu_e|\nu\rangle|^2$.

Pictorially, this process of two-neutrino mixing is represented in Figure 2.1

For the 862 keV ${}^7\text{Be}$ ν 's and assuming $\Delta m^2 = \Delta m_{12}^2 = 7.59 \times 10^{-5} \text{eV}^2$ [7], equation 2.10 gives a vacuum oscillation length of approximately $\lambda_{vac} \approx 28.1$ km. From examining equation 2.9, we see that the oscillation (or survival if $\alpha = \beta$) probability depends on the distance from




Figure 2.1: Graphical representation of two-neutrino mixing of $|\nu_e\rangle \rightarrow |\nu_\mu\rangle$. The electron neutrino begins with the two mass eigenstates ($|\nu_1\rangle$ and $|\nu_2\rangle$) in phase, however, during propagation, the mass eigenstates lose their coherence due to the mass difference between them. As the coherence is lost, it becomes increasingly probable to observe the initial $|\nu_e\rangle$ as a $|\nu_\mu\rangle$; eventually the coherence is regained and the process repeats. Figure from [6].

the source to the detector. For terrestrially based solar neutrino experiments, the distance L is not a constant for two reasons:

1. The distance from the earth to the sun changes over the course of a year (see Figure 3.2 and the accompanying discussion) — the total variation in distance is roughly 5×10^6 km.
2. Neutrinos are produced not from a point like source, but rather, inside of a shell of the sun’s core — the distribution is that of a radial Gaussian which, for ${}^7\text{Be}$ neutrinos, peaks at $0.06R_\odot$ with half heights occurring at $0.03R_\odot$ and $0.1R_\odot$ respectively[8]. Thus, a neutrino produced at the “earth-facing” edge of the ${}^7\text{Be}$ ν production region travels almost 14×10^5 km less than one produced on the diametrically opposite side.

The net result is that interference between neutrinos emitted at various locations inside of the sun tends to smooth out the vacuum oscillation caused by the phase difference in the mass eigenstates. More precisely, the second \sin^2 in equation 2.9 averages to $\frac{1}{2}$ and the oscillation probability is a constant determined only by the value of θ :

$$P(|\nu_\alpha\rangle \rightarrow |\nu_\beta\rangle) \approx \frac{1}{2} \cdot \sin^2(2\theta_{ij}) \quad (2.13)$$

We see from 2.11 that the survival probabilities for a given Δm^2 are determined by the ratio $\frac{L}{E}$, thus, if we want to work backwards and measure Δm^2 we must choose our source-detector pairing such that the ratio $\frac{L}{E}$ will afford sensitivity to the range of Δm^2 we wish to study. In order for the argument of the second \sin^2 function in the aforementioned equation to be relevant, it should be approximately of order 1, or $\Delta m^2 \approx \frac{E}{L}$. A list of the lower limits of the experimental sensitivities to Δm^2 is presented below in Table 2.1 — it is important to note the extremely high sensitivity of solar neutrino experiments to squared mass splittings.

Neutrino Source	$L(\text{m})$	E (MeV)	Δm^2 (eV) ²
Solar	10^{10}	≈ 1	$10^{-10} - 10^{-11}$
Atmospheric	$10^4 - 10^7$	$10^2 - 10^5$	$10^{-1} - 10^{-4}$
Reactor	$10^2 - 10^5$	1	$10^{-2} - 10^{-5}$
Accelerator	$10^2, 10^5 - 10^6$	$10^3 - 10^4$	$> 0.1, 10^{-2} - 10^{-3}$

Table 2.1: This table presents the lower limits of sensitivity to the squared mass splitting Δm^2 for the relevant neutrino sources and baselines. The presence of two entries in the last row is due to the separate treatment of short and long baseline experiments. The ability of a solar neutrino experiment to resolve the differences between squared neutrino masses is a rather remarkable effect considering the scales of the masses involved.[9]

Three-Neutrino Mixing

The generalization of two-neutrino mixing into the three-neutrino basis is a simple, albeit tedious, exercise. The argument proceeds in the same way, and the resulting mixing matrix is now 3×3 and is commonly known as the MNSP (Maki-Nakagawa-Sakata-Pontecorvo) matrix:

$$\begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \\ |\nu_\tau\rangle \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \cdot \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \\ |\nu_3\rangle \end{pmatrix} \quad (2.14)$$

The MNSP matrix is defined by three mixing angles $\theta_{12}, \theta_{23}, \theta_{13}$ and one CP violating phase factor δ and can be represented as shown below using the shorthand $c_{\gamma\kappa} = \cos(\theta_{\gamma\kappa})$ and likewise for sin:

$$U_{MNSP} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \quad (2.15)$$

A more useful representation of equation 2.15 can be obtained by expanding it into a product of three matrices, resembling standard (x, y, z) rotation matrices:

$$U_{MNSP} = \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{\text{Atmospheric}} \cdot \underbrace{\begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix}}_{\text{Long Baseline}} \cdot \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{Solar}} \quad (2.16)$$

From equation 2.16 we can clearly see that solar neutrino oscillation is primarily oscillation of ν_e to ν_μ ; atmospheric oscillation occurs primarily between ν_μ and ν_τ ; and the long baseline experiments primarily observe oscillation from ν_e to ν_τ . In addition to the three mixing angles, the complete neutrino oscillation picture requires knowledge of the Δm^2 's; only two of which are independent because $\Delta m_{12}^2 + \Delta m_{23}^2 = \Delta m_{13}^2$. The best-fit values from a global analysis of neutrino data are shown in Table 2.2.

Two important questions remaining in neutrino physics revolve around the neutrino masses. First, we currently have only a very weak limit on what the absolute neutrino masses are, because our oscillation experiments are only sensitive to the squared mass splittings. The absolute neutrino mass is accessible by experiments which measure the deficit in the beta

Oscillation Parameter	Value
Δm_{21}^2	$7.59_{-0.21}^{+0.20} \times 10^{-5} \text{ eV}^2$
$ \Delta m_{31}^2 $	$2.40_{-0.11}^{+0.12} \times 10^{-3} \text{ eV}^2$
$\sin^2(\theta_{12})$	$0.314_{-0.014}^{+0.018}$
$\sin^2(\theta_{23})$	$0.50_{-0.06}^{+0.07}$
$\sin^2(\theta_{13})$	$0.01_{-0.011}^{+0.016}$

Table 2.2: The current best-fit values of the five neutrino oscillation parameters (neglecting the CP violating phase factor δ). The values for Δm_{12}^2 and $\sin^2(\theta_{12})$ are from the recent low energy neutrino analysis from SNO ([10]), the rest of the values are from the global analysis performed by Schwetz, Tórtola and Valle ([11]). Note the extremely small value (and large error) for $\sin^2(\theta_{13})$, which makes it possible to neglect $|\nu_e\rangle \rightarrow |\nu_\tau\rangle$ transitions in our discussion of solar neutrino oscillations.

decay endpoint energy (on a Curie plot) due to a non-zero neutrino mass. Also sensitive to the neutrino mass, albeit in a less direct way, are neutrino-less double-beta decay experiments (discussed further in Section 2.3.2). Second, we do not know the ordering of the masses: there is the “normal” mass hierarchy where $m_1 < m_2 < m_3$, yet an equally plausible alternative is provided by the “inverted” hierarchy in which $m_3 < m_1 < m_2$ — both are completely consistent with our knowledge of the squared mass splittings. Effectively, the latter question can be resolved by examining the sign of Δm_{13}^2 (or equivalently Δm_{23}^2) — a method is afforded by the MSW effect (see section 2.1.2) which is only possible for a positive value of the squared mass splitting.

2.1.2 Neutrino Oscillations in Matter

The formalism that has been presented thus far only holds for neutrinos which are travelling through a vacuum (such as space). As a neutrino of any flavor passes through matter, it can undergo neutral current interactions mediated by the Z^0 boson which will add a potential energy, V_{NC} , to the interaction Hamiltonian: $\mathcal{H}_M = \mathcal{H}_{vac} + V_{NC}$. Since all active neutrino flavors undergo neutral-current scattering, V_{NC} is diagonal and an integer multiple of the identity matrix. The subject of neutrino oscillations — in vacuum or matter — is a quantum interference phenomenon, the relevant quantity is the relative phase of the eigenstates which, as was shown in Section 2.1.1, is given by the relative energies of the eigenstates. Thus, we can safely neglect (or just subtract) from the Hamiltonian any factors which are multiples of the identity matrix since they have no bearing on the relative phases and therefore the oscillation / survival probabilities. With that being said, the neutral current interactions with matter can be neglected and we do not enter them into the matter Hamiltonian.

The case of charged current interactions¹ mediated by the W^\pm bosons are a more complicated matter however. Since the sun and earth² do not contain any μ 's or τ 's, the charged current process is only possible for $|\nu_e\rangle$. Therefore, V_{CC} — the interaction potential due to the charged current process — is not diagonal, cannot be subtracted from the matter Hamiltonian, and therefore *does* cause a different oscillation scenario than the vacuum oscillations of the previous section. This oscillation enhancement for electron flavor neutrinos is known as the MSW — Mikheyev, Smirnov, Wolfenstein — effect [12], [13]. The effective mass acquired through the forward scattering of electron neutrinos is given by $\sqrt{2}G_F n_e$ in the energy basis. The matter Hamiltonian in the flavor basis is thus:

¹First order Feynman diagrams for the charged and neutral current interactions can be found in the following chapter in Figure 3.1.

²We restrict the discussion to the case of neutrinos that travel either through the sun, the earth, or both. The argument holds for any collection of electronic matter — be it in other stars, planets, galaxies, etc.

$$\mathcal{H}_M = U_{MNSP} \mathcal{H}_{Vac} U_{MNSP}^\dagger + \sqrt{2} G_F n_e \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (2.17)$$

In 2.17, G_F is the Fermi constant ($8.96 \times 10^{-32} \text{MeVcm}^3$) and n_e is the local electron density, and is in general a function of position. In the regime of a very small value of θ_{13} , the situation is again easily described by two-neutrino mixing. The difference of the energy eigenvalues of \mathcal{H}_M , which is proportional to the phase difference, is given by:

$$E_2 - E_1 = \sqrt{\left(\frac{\Delta m^2}{2E} \cos(2\theta_V) - \sqrt{2} G_F n_e\right)^2 + \left(\frac{\Delta m^2}{2E}\right)^2 \sin^2(2\theta_V)} \quad (2.18)$$

Apart from factors of \hbar and c , the effective squared mass splitting is:

$$\Delta m_m^2 = \sqrt{\left(\Delta m^2 \cos(2\theta_V) - 2\sqrt{2} G_F n_e E\right)^2 + (\Delta m^2 \sin(2\theta_V))^2} \quad (2.19)$$

Finally, the equation for the matter-enhanced oscillation amplitude is given by:

$$\sin^2(2\theta_m) = \frac{\sin^2(2\theta_V)}{\left(\cos(2\theta_V) - \frac{2\sqrt{2} G_F n_e E}{\Delta m^2}\right)^2 + \sin^2(2\theta_V)} \quad (2.20)$$

The survival probability for an electron flavor neutrino traveling through matter has the same form as for the vacuum case (assuming two-neutrino mixing):

$$P(|\nu_e\rangle \rightarrow |\nu_\mu\rangle)_m = \sin^2(2\theta_m) \sin^2\left(\frac{\Delta m_m^2 L c^3}{4\hbar E}\right) \quad (2.21)$$

$$P(|\nu_e\rangle \rightarrow |\nu_e\rangle)_m = 1 - \sin^2(2\theta_m) \sin^2\left(\frac{\Delta m_m^2 L c^3}{4\hbar E}\right) \quad (2.22)$$

Resonance Enhancement of Matter Oscillations

An important factor which affects the survival probabilities is that the matter through which a given neutrino may pass will not have a constant density and will vary with position (and therefore time), thus, θ_m and Δm_m^2 are functions of time. An interesting phenomenon occurs if a neutrino traverses a region of matter with a specific value of the local electron density. In equation 2.20 it is easy to see that for a special electron density it is possible to have the first term in the denominator equal to zero if:

$$\cos(2\theta_V) = \frac{2\sqrt{2}G_F n_e E}{\Delta m^2} \quad (2.23)$$

When this condition is met, equation 2.20 becomes:

$$\sin^2(2\theta_m)_{res} = \frac{\sin^2(2\theta_V)}{(0)^2 + \sin^2(2\theta_V)} = 1 \quad (2.24)$$

Thus, the oscillation amplitude is equal to unity *regardless of the value of θ_V* . This phenomenon is known as the resonance enhancement of neutrino oscillations in matter because its net effect is to amplify a potentially negligible value for θ_V into maximal mixing if the resonance electron density is reached. When the resonance condition is reached it is easy to check that equation 2.20 yields $\theta_M = \frac{\pi}{4}$, thus, the admixtures of $|\nu_1\rangle$ and $|\nu_2\rangle$ are equal and we have maximal mixing for any value of the vacuum mixing angle θ_{12} .

The minimum energy for which resonance occurs in the sun can be obtained by solving equation 2.23 for E after inserting an appropriate figure for the solar electron density. The

electron density n_e in a material is, on average, given by $n_e = \frac{\rho}{m_p} \frac{n_e}{n_n + n_p}$, the second factor can be taken to be \bar{Z}/\bar{A} which is approximately equal to 0.7 in the neutrino producing regions in the core of the sun. Thus, the minimum energy is given by (using the approximation $G_F = 10^{-5}/m_p^2$ with $\hbar = c = 1$):

$$E_{min}^r = \frac{\cos(2\theta_{12}) \Delta m_{12}^2 m_p \bar{A}}{2\sqrt{2} G_F \rho \bar{Z}} \quad (2.25)$$

$$\approx 6.6 \cos(2\theta_{12}) \left(\frac{\Delta m_{12}^2}{10^{-4}} \right) \quad (2.26)$$

$$(2.27)$$

With the values for Δm_{12}^2 and θ_{12} given in Table 2.2, the minimum energy for resonance enhancement of solar neutrinos is approximately 1.86 MeV. Therefore, for solar neutrinos with energies below ~ 1.9 MeV, the MSW effect does not introduce an appreciable effect, however, above this energy the oscillation probability is significantly altered as can be seen in Figure 2.2. Figure 2.3 shows experimental data superimposed on an averaged survival probability curve; Borexino has already provided data points in the vacuum and matter dominated oscillation regimes, and can — via the pep neutrinos — obtain a data point in the transition region.

It is important to note that for a neutrino born at a distance $r > r_{res}$ (where r_{res} is the location where the electron density is at its resonant value) it will either pass through zero or two resonance layers if its momentum vector points approximately in the $-\hat{r}$ direction. Consider the two separate cases of a neutrino born at $r > r_{res}$ on the side of the sun closest to the earth, and another neutrino born on the side of the sun opposite the earth; the former neutrino will not pass through a resonance layer at all and will not experience significant matter enhancement of its oscillation probability. However, a neutrino born on the opposite

side of the sun with a momentum direction approximately in the $-\hat{r}$ direction will pass through *two* resonance layers as it travels first towards $r = 0$ and then back out towards R_\odot . This effect will act to smear the oscillation probability and is taken into account with the simulation codes used by theoretical physicists when generating neutrino survival probability plots like that shown in Figure 2.3.




Figure 2.2: The survival probability of ν_e in matter as a function of energy. For low neutrino energies, the matter effect is negligible because the denominator of equation 2.20 is close to 1. At energies above the minimum resonant energy of ~ 2 MeV (equation 2.25), matter oscillations are dominant and the suppression can be seen. At much higher energies, the Hamiltonian is no longer diagonal and transitions from $|\nu_{1m}\rangle \rightarrow |\nu_{2m}\rangle$ are possible — however, this is irrelevant for solar neutrino energies. Figure from reference [14].

The resonance condition (equation 2.23) is only reached if $\Delta m^2 > 0$ (otherwise the first term in the denominator of the matter enhanced oscillation amplitude can not go to zero). Since it is defined as $\Delta m_{21}^2 = \Delta m_2^2 - \Delta m_1^2$, the MSW effect will only be observed if $m_2 > m_1$ — the fact that it has been observed in solar neutrino experiments ([16]) means that we do know that the sign of Δm_{21}^2 is positive. If we could produce a very long baseline accelerator experiment, we could in theory look for the resonance enhancement of the oscillation of μ to τ flavor neutrinos; non-observation of the MSW effect in the earth would exclude a negative Δm_{32} .




Figure 2.3: Experimental data points and their errors superimposed on the predicted survival probability curve obtained with the global fit parameters for the MSW-LMA solution. To date, the Borexino experiment has performed measurements in both the vacuum oscillation region as well as the matter dominated region. A critical area to examine is that of the transition region between the two oscillation regimes — the 1.44 MeV pep neutrinos, to which Borexino is sensitive, provide an opportunity to perform this measurement. Figure from [15].

Adiabatic Flavor Conversion

A special condition of the MSW effect occurs if a neutrino traverses a medium in which the electron density changes slowly or *adiabatically* — in other words, a small $\frac{d\theta_m}{dt}$. In this regime, the propagation eigenstates $|\nu_{1m}\rangle$ and $|\nu_{2m}\rangle$ have time to adjust to the density change and therefore propagate independently without $|\nu_{1m}\rangle \leftrightarrow |\nu_{2m}\rangle$ transitions. The negligible transitions between the instantaneous mass eigenstates means that the admixtures of the eigenstates do not change — they are fixed at the value of θ_m at their production point: θ_m^0 .

For an electron neutrino produced in the core of the sun, the initial admixture of $|\nu_{1m}\rangle$ and $|\nu_{2m}\rangle$ depends on the value of θ_m , which, for the very large electron density in the core of the sun means that $\theta_m \sim \frac{\pi}{2}$. Thus, at birth, the neutrino flavor state is given by:

$$|\nu_e\rangle = \cos(\theta_m)|\nu_{1m}\rangle + \sin(\theta_m)|\nu_{2m}\rangle \quad (2.28)$$

$$|\nu_e\rangle_0 \approx \sin(\theta_m)|\nu_{2m}\rangle \quad (2.29)$$

Thus, an electron neutrino produced at the center of the sun is born predominantly in the $|\nu_{2m}\rangle$ eigenstate. Now, as the neutrino propagates outwards through the sun, the density reduces approximately monotonically with a shallow slope, and meets the condition for an adiabatic density change. This slow density change means that the neutrino stays in the $|\nu_{2m}\rangle$ mass eigenstate and upon exiting the sun, the density (and the $|\nu_{2m}\rangle$ eigenstate) transitions smoothly and the neutrino now exists predominantly in the $|\nu_2\rangle$ vacuum eigenstate. Once the neutrino reaches the earth it is observed as a ν_e with probability:

$$P(|\nu(t)\rangle \rightarrow |\nu_e\rangle) = |\langle\nu_e|\nu(t)\rangle|^2 \approx \sin^2(\theta_V)_{12} \sim \frac{1}{3} \quad (2.30)$$

Since the density at the core is so high, the initial mixing is very minimal, and the adiabatic density change means that the conversion can be very close to maximal. An excellent representation of this effect can be seen in Figure 2.4. Another representation of this level crossing effect is provided in Figure 2.5.

2.2 Solar Neutrinos

The fusion process occurring in the core of the sun proceeds by fusing four protons to yield: one helium atom, two positrons, and two electron neutrinos. The reaction has a total energy release given by the mass difference between the initial and final state particles of 26.7 MeV:





Figure 2.4: Pictorial representation of the adiabatic flavor conversion caused by the MSW effect for an electron neutrino produced in the sun as a function of electron density. The upper box represents $|\nu_2\rangle$, and the lower one represents $|\nu_1\rangle$ — the colors represent the flavors (red for $|\nu_e\rangle$ and green for $|\nu_\mu\rangle$), and the yellow line indicates the resonant electron density. At production, the neutrino is predominantly in the $|\nu_2\rangle$ state and as the density decreases (note that at the resonant electron density the mixing is maximal), it stays in this state and becomes increasingly more probable to observe $|\nu_\mu\rangle$. The adiabatic change in the density means that there are no $|\nu_{1m}\rangle \leftrightarrow |\nu_{2m}\rangle$, thus, the lengths of the boxes do not change. Figure from [17].



For stars with masses less than $\sim 1.3M_\odot$, the elemental burning proceeds via the pp (proton-proton) cycle, and for larger stars, the core temperature is high enough to burn heavier elements and proceeds via the CNO (Carbon-Nitrogen-Oxygen) cycle. The method by which the burning proceeds determines the energy spectrum of the neutrinos which we observe in our detectors and will be discussed further in Sections 2.2.1 and 2.2.2 for the pp and CNO chains respectively.

Since the same reactions which produce the neutrinos we observe also produce the energy which eventually gets emitted as the light that we see (albeit the latter is delayed by $\sim 10^5$ years) we can obtain a constraint on the total neutrino flux by measuring the light output of the sun. This constraint, known as the luminosity constraint[19], assumes that we know all of the reactions which produce neutrinos, and that the fusion rate of the sun is stable over the time it takes for the energy produced in the core to be emitted as photons at the surface of the sun. The neutrino flux is given by:

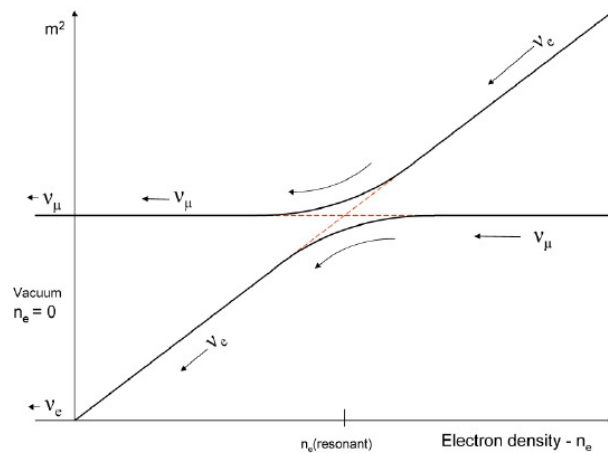


Figure 2.5: The level crossing scheme for the MSW-LMA effect for solar neutrinos. The lines are proportional to the energy level of the mass eigenstates; the solid ones are for an adiabatic electron density change, whereas the dashed lines represent non-adiabatic mixing or no mixing. Produced in the solar center, the initial electron neutrino is predominantly in the heavier mass eigenstate, and if the density changes slowly, it stays in the heavier mass eigenstate all the way until it exits the sun. If the neutrino is in the heavier mass eigenstate when it arrives at the earth, it has only a 33% probability to be observed as an electron neutrino (see equation 2.30). Figure from [18].

$$\Phi_\nu = \frac{2L_\odot}{26.7\text{MeV} \cdot 4\pi r^2} \quad (2.32)$$

The factor of 2 in the numerator in equation 2.32 is due to the fact that two neutrinos are emitted for each reaction which releases the 26.7 MeV, and the distance from the sun to the earth, r , is not a constant due to the elliptical orbit. Averaging the earth-sun distance we obtain an average neutrino flux of $6.6 \times 10^{10}/(\text{cm}^2 \cdot \text{s})$ — almost *70 billion* neutrinos per second pass through a square centimeter every second. With so many neutrinos, one would naively think they would be easy to detect...

The actual luminosity constraint is given by the simple relation:


$$\frac{L_{\odot}^{\gamma}}{4\pi r^2} = \sum_i \alpha_i^{\nu} \Phi_i^{\nu} \quad (2.33)$$

In equation 2.33 we have define L_{\odot}^{γ} to be the solar luminosity, and again r is the earth-sun distance, the factors α_i^{ν} is the average energy released by the neutrino producing reaction i , and Φ_i^{ν} is the neutrino flux from the reaction i . A precise test of whether or not $L_{\odot}^{\gamma}/L_{\nu} = 1$ is of paramount importance for the neutrino community in order to vet our current understanding of solar energy production. Since the pp neutrinos comprise the vast majority of the flux, a truly rigorous test of the luminosity constraint requires a precise measurement of the pp neutrino flux.

Combining the Borexino ${}^7\text{Be}$ ν measurement with the Gallium and Chlorine results allows one to obtain a value for f_{pp} which is the ratio of the model prediction to the measured value of $f_{pp} = 1.005_{-0.020}^{+0.008}$ [20]. While the central value is very close to 1, the large error bars make a more precise measurement highly desirable and is one of the primary goals of the LENS experiment[21].

2.2.1 PP chain

The first comprehensive paper in which the mechanisms for stellar energy production were first laid out was by Bethe in his 1939 article “Energy Production in Stars” [22]. The pp chain provides four possible terminations for the fusion of four protons. However, lepton number conservation dictates that two neutrinos must be produced by the end of whichever reactions occur to produce the helium nuclei. The first step of the chain is one of the two following reactions with the relative branching ratios indicated:




This first step of the chain is the slowest step in the process because it is governed by the weak interaction. The second of the two interactions which produces the so-called “pep” neutrinos is an even slower process because it requires three initiating particles as opposed to two — the ratio of initiating reactions which produce “pp” neutrinos to those which produce “pep” neutrinos is approximately 400 : 1. From Table 2.3, we can see that the energy spectrum for the pp neutrinos is a continuum with an endpoint at about 420 keV, however the direct measurement of such a low energy neutrino is very complicated due to the presence of naturally occurring background in the energy region of interest. By contrast, the pep neutrinos are mono-energetic at 1.4 MeV, and therefore lie above the vast majority of naturally occurring background, and a direct measurement is a slightly easier task. Since the two reactions are effectively identical (they are transformed into each other by adding an electron to both sides), their matrix elements are identical and they describe effectively the same physics — therefore, the easier measurement of the pep neutrinos can constrain the pp interaction.

Regardless of the method by which the deuterium nucleus is formed, it will quickly capture a free proton via:



From this point, there are three possible reactions that are collectively known as the ppII branch:

1. In approximately 85% of the cases, two such ${}^3\text{He}$ nuclei will undergo:




The chain is complete if it reaches this point, which it usually does.

2. Another possibility which occurs $\sim 15\%$ of the time is for a ${}^3\text{He}$ to find a ${}^4\text{He}$ produced by reaction 2.36 (or 2.38) and undergo:



3. The final possible variation is for a single ${}^3\text{He}$ to undergo proton capture:



However, this interaction requires a flavor change and is thus weak and very slow compared to the other two possible interactions — the branching ratio is $\sim 10^{-8}$. The chain would also conclude at this point since a helium nucleus is produced.

After the conclusion of the ppII branch, only one possible reaction (2.37) has not led to a termination. If a ppI branch results in ${}^7\text{Be}$, then the chain continues with the ppIII branch with two possible interactions to burn the ${}^7\text{Be}$:

1. The vast majority of the time that the reaction gets to this point in the chain (roughly 15% of all the reactions), the ${}^7\text{Be}$ encounters an electron and captures it in the reaction



The electron capture can occur to the ground state as well as the first excited state of ${}^7\text{Li}$, the former gives the full energy 0.862 MeV to the neutrino, and the latter gives only 0.383 MeV to the neutrino — the branching ratio to the excited state of ${}^7\text{Li}$ is roughly 10%. Once the Li has been produced it quickly captures a free proton in the reaction:



2. The remaining 0.016% of the time, the ${}^7\text{Be}$ captures a proton in the reaction:



The ${}^8\text{B}$ produced is unstable and splits into two helium nuclei via:



A plot of the neutrino fluxes versus energy is given in Figure 2.6, and a cartoon drawing of the pp chain is given in Figure 2.7.

One might wonder why the production of He does not proceed immediately after the first step in the chain by merging two deuterium nuclei via: ${}^2\text{H} + {}^2\text{H} \rightarrow {}^4\text{He}$. While in theory this is possible, the production of ${}^3\text{He}$ via the reaction depicted in equation 2.35 occurs so rapidly, and the first step is so slow, that the local abundance of deuterium is never (statistically speaking) high enough for the fusion of two deuterium nuclei into helium to occur.




Figure 2.6: Plot of neutrino flux versus energy for the pp cycle (black) neutrinos as well as the CNO cycle (blue) neutrinos. Model uncertainties in the fluxes are indicated for each component of the pp chain. The data used to draw the picture are from the solar model produced by Bahcall and collaborators in 2005 — so, the fluxes in the plot are mildly different from the present day (2010) solar models. Figure from [23].

2.2.2 CNO cycle

The sun produces its energy primarily through the pp chain reactions — approximately 98.5% of its luminosity is from the pp chain, and the remaining 1.5% is generated by the reactions in the CNO cycle. In terms of the solar neutrino flux, the CNO neutrinos contribute about 0.76% of the total solar neutrino flux with the pp chain taking up the rest. The basic process of fusing four protons into a helium nucleus, two positrons, and two neutrinos is still the same, however, in this case, nuclei of carbon, nitrogen, and oxygen participate as spectators which catalyze the reactions. In stars significantly larger than the sun, the core temperature is high enough for the participating particles to overcome the Coulomb barrier and thus for the CNO cycle to become the dominant energy production scheme — it is this high degree of correlation to the core temperature which introduces significant uncertainty

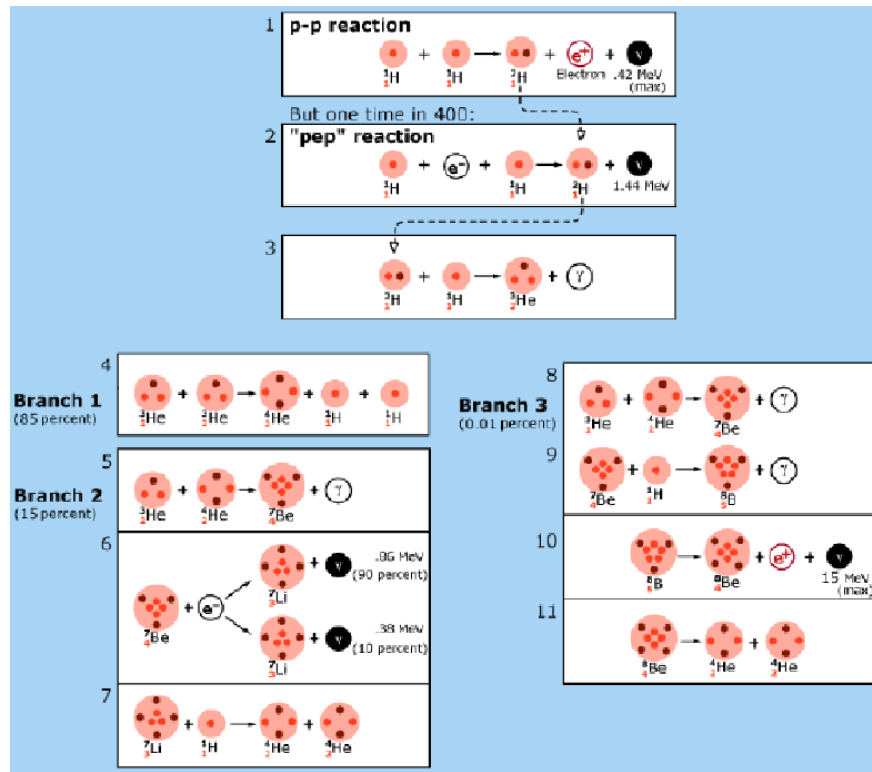



Figure 2.7: The pp chain for producing neutrinos. The overall reaction is $4p \rightarrow {}^4\text{He} + 2e^+ + 2\nu_e$, however, several different paths are available (the relative branching ratios of each reaction are indicated). The very rare hep neutrinos are not shown. Figure from [24]. © Nobel e-Museum June, 2000.

in the CNO fluxes.


The chain of reactions is shown in Figure 2.8, and the fluxes and energies are given in the lower half of Table 2.3. The cycle begins³ with a ${}^{12}\text{C}$ nucleus capturing a free proton to form radioactive ${}^{13}\text{N}$ which in turn produces a neutrino:



³The definition of where it “begins” is rather arbitrary since the reaction is a loop.

$$(2.45)$$

The neutrino produced in reaction 2.44 has a continuous spectrum up to $2.221 - 2 \cdot 0.511 = 1.199$ MeV. The reaction continues with the decay of ^{13}C until it reaches a branching point after the production of ^{15}N :



Again, the neutrino produced in reaction 2.48 has a continuous energy with an endpoint given by the Q value minus the energy of the annihilation gammas: 1.732 MeV. At this point, the ^{15}N can decay by one of two channels — the most probable one has a branching ratio of 99.9% and proceeds via:



The ^{12}C produced in reaction 2.49 is now available to begin the chain again at reaction 2.43. The other possible method of burning the ^{15}N proceeds at much lower rate because the temperature in the core is not sufficient to appreciably burn it.





Only in 0.01% of the cases does the CNO chain produce a neutrino from the decay of ${}^{17}\text{F}$, thus the flux is much lower — and this is confirmed by the data in Table 2.3.




Figure 2.8: The CNO cycle of solar neutrino generation in the sun. The net reaction for the chain is still $4p \rightarrow {}^4\text{He} + 2e^+ + 2\nu_e$, however, the carbon, nitrogen, and oxygen nuclei participate as spectators. The CNO cycle contributes about 1.5% to the total solar luminosity, but only 0.76% to the total solar neutrino flux. This figure only shows the dominant set of reactions in the chain — see reactions 2.50 through 2.53 for the sub-dominant alternative method for burning ${}^{15}\text{N}$. Figure from [25].

2.3 Outlook

A detailed account of the history of solar neutrino physics and experiments is beyond the scope of this dissertation and does not play an important role in its contents; for such material see, for example, reference [26]. The more relevant question regarding neutrino history is: what will be added to it in the coming years? There is no shortage of topics in neutrino physics to study, however, we restrict our discussion to: what are the neutrino masses and why are they so different from the other leptons? What is the value of θ_{13} ? Are neutrinos Dirac or Majorana particles? And what are the solar elemental abundances?

2.3.1 Neutrino Mass

One of the more curious aspects of neutrinos is why their masses appear to be so *vastly* smaller than the rest of the fundamental particles. A plot of the fermion masses is presented in Figure 2.9 and shows this rather disparaging difference. The many orders of magnitude spacing between neutrino masses and the rest of the fermions hints towards a possible “chink in the armor” of our theory of fermion masses.

The formalism for masses in the standard model has two separate treatments: one for Dirac particles, and one for Majorana particles. Specializing to the case of neutrinos, a Dirac mass term would transform a neutrino to a neutrino, or an anti-neutrino to an anti-neutrino ($\nu \rightarrow \nu'$ or $\bar{\nu} \rightarrow \bar{\nu}'$); whereas, a Majorana mass term transforms a neutrino to an anti-neutrino ($\nu \rightarrow \bar{\nu}'$) and vice versa. The latter then is only applicable for particles with zero electric charge: if an electron obtained its mass through a Majorana interaction, it would be converted to a positron — in clear violation of the conservation of electric charge. If we consider the Majorana interaction which transforms a ν to a $\bar{\nu}$, we can see that such an interaction does not conserve lepton number; a slightly disappointing result, but one which




Figure 2.9: A 1-D plot showing the masses of the fermions in the standard model. The red dots indicate the leptons, green and blue dots for the quarks. The large order of magnitude difference between the neutrino masses and rest of the fermions is a bit surprising and suggests that we do not completely understand the method by which neutrinos gain their masses. Figure from [27] and reprinted with permission of the author.

appears to be ultimately required.

In terms of the weak Lagrangian, a Dirac mass interaction can be included with a term given by:

$$\mathcal{L}_D = -m^D (\bar{\nu}_R \nu_L + \nu_R \bar{\nu}_L) \quad (2.54)$$

And the Majorana mass terms (for both left- and right-handed states) are given by:

$$\mathcal{L}_M = -\frac{m_L^M}{2} \left\{ \overline{(\nu_L)^C} \nu_L + \text{h.c.} \right\} - \frac{m_R^M}{2} \left\{ \overline{(\nu_R)^C} \nu_R + \text{h.c.} \right\} \quad (2.55)$$

In equation 2.55, h.c. means the hermitian conjugate of the preceding term, the ν terms indicate the neutrino *fields* (not the physical particles), the L/R subscripts indicate the chirality of the field, and the superscript C indicates the charge conjugate of the field. Combining the

two mass terms gives the generalized neutrino mass Lagrangian:

$$\mathcal{L}_{m_\nu} = -\frac{1}{2} \left(\overline{(\nu_L)^C}, \overline{(\nu_R)} \right) \begin{pmatrix} m_M^L & m^D \\ m^D & m_M^R \end{pmatrix} \begin{pmatrix} \nu_L \\ (\nu_R)^C \end{pmatrix} + \text{h.c.} \quad (2.56)$$

We can transform the equation above into the mass basis $|\nu_i\rangle$ through multiplication by an appropriate unitary matrix. Doing so under the approximation that $m_L^M \ll (m^D)^2/m_R^M$ and that $m_R^M \gg m^D$ allows the mass matrix to be put into the form:

$$\mathcal{M}_\nu = \begin{pmatrix} \frac{(m^D)^2}{m_R^M} & 0 \\ 0 & m_R^M \end{pmatrix} \quad (2.57)$$

This matrix multiplies the two-component neutrino mass state defined by⁴:

$$\nu = \nu_L + (\nu_L)^C = \begin{pmatrix} \nu_\ell \\ \nu_r \end{pmatrix} \quad (2.58)$$

So, we see that in the end we have two neutrino mass states, one with a mass given by $(m^D)^2/m_R^M$ and the other with a mass of m_R^M . This mechanism (known as the “seesaw” mechanism) allows one to escape from the apparent problem that the neutrino masses are so much smaller than the rest of the fundamental fermions. If we take m^D to be on the order of that for the rest of the charged fermions, and if the right-handed Majorana mass term, m_R^M , is large enough, then we can achieve a mass for the ν_ℓ state that is on par with what we expect for ν_e, ν_μ , and ν_τ . On the other hand, the ν_r state has an extremely large mass

⁴Note that the state described by equation 2.58 is identical to itself under charge conjugation — this is precisely the definition of a Majorana neutrino: one that goes into itself under charge conjugation. Thus, *any* neutrino described by a Lagrangian that includes both Dirac and Majorana mass terms is a Majorana neutrino.

(predicted to be near the GUT scale of $\sim 10^{15}$ GeV), and is the reason why we have not yet detected it⁵.

The “seesaw” mechanism does provide an attractive solution by allowing a right-handed neutrino mass state to be very heavy and allowing the physical neutrino masses to be so small. However, at present it is just a theory that must be tested. Short of observing the extremely massive neutral lepton described by the $|\nu_r\rangle$ state (which is certainly not feasible on any foreseeable time scale), the best chance for vetting this model is the observation of neutrino-less double beta decay — this is only possible if neutrinos are Majorana particles. A controversial claim of its observation has been made by a subset of the Heidelberg-Moscow experiment in 2005, however, a multitude of next-generation experiments are preparing to perform even more sensitive tests for its existence.

2.3.2 Neutrino-less Double Beta Decay

For certain unstable nuclei, normal beta decay is forbidden by energy conservation because the daughter isotope is actually heavier than the parent. For instance, the isotope ^{76}Ge has a mass of 75.9214 amu, and its would-be daughter under beta decay is ^{76}As with a mass of 75.9224 — obviously the direct decay is forbidden. However, the next isotope on the $A = 76$ mass parabola is ^{76}Se with a mass of 75.9192. Therefore, it is possible for the unstable ^{76}Ge to decay by emitting two betas and two anti-neutrinos simultaneously: $^{76}\text{Ge} \rightarrow ^{76}\text{Se} + 2e^- + 2\bar{\nu}_e$. This process is called “2-neutrino double beta decay” ($2\nu\beta\beta$) and is exceedingly slow due to it being a second-order process involving the weak interaction. The decay mentioned above has a mean life on the order of 10^{21} years.

⁵This treatment has only dealt with one particular flavor of neutrino — abstracting it out to handle all three flavors means that m^D is now a 3×3 matrix and the situation becomes more complicated. A treatment of such a case is given in, for example, [28].

There is another possible decay path for such nuclei which is only available if neutrinos have a Majorana mass: $(A, Z) \rightarrow (A, Z + 2) + 2e^-$ in which an anti-neutrino is emitted by one nuclei, and is able to be absorbed by another nuclei — this process is known as “neutrino-less double beta” ($0\nu\beta\beta$) decay. The Feynman diagram for this process is depicted below in Figure 2.10.




Figure 2.10: Feynman diagram for the hypothetical process of neutrino-less double beta decay. The process requires that a neutrino be a Majorana particle — identical to its antiparticle — and that lepton number violation occur. Observation of neutrino-less double beta decay would verify the Majorana nature of neutrinos, and can give insight into neutrino masses. The main body text deals more explicitly with these issues.

In order for neutrino-less double beta decay to occur, the neutrinos must be Majorana particles, thus $|\nu\rangle = |\bar{\nu}\rangle$. When a neutron inside of the decaying nucleus emits an anti-neutrino, it can, through the Majorana mass term, convert to a neutrino which is able to be absorbed by another nuclei and induce a beta decay — the process appears as if the two neutrinos have disappeared, leaving just the two electrons to carry the full Q value of the decay. The signature of a neutrino-less double beta decay experiment is a mono-energetic peak at the end of the two-neutrino double beta decay continuum due to the summed energy of the two electrons, this requires that the background in that region be exceedingly low and that the energy resolution be very high to be able to deconvolve the $0\nu\beta\beta$ peak from the $2\nu\beta\beta$ continuum. Integration of the $0\nu\beta\beta$ peak (minus any background components) yields

the number of events observed during the livetime, which yields a lifetime for the decay⁶.

The decay rate for the neutrino-less double beta decay process is directly related to the effective neutrino mass via:

$$\Gamma^{0\nu\beta\beta} = G^{0\nu\beta\beta} \cdot |\mathcal{M}^{0\nu\beta\beta}|^2 \cdot |m_{ee}^2| \quad (2.59)$$

$$m_{ee} = \sum_i U_{ei}^2 m_i \quad (2.60)$$

Where m_{ee} is known as the effective Majorana mass (sometimes written as $m_{\beta\beta}$), $G^{0\nu\beta\beta}$ is a phase space factor, and $\mathcal{M}^{0\nu\beta\beta}$ is the nuclear matrix element⁷. To date, only a subset of the Heidelberg Moscow experiment (lead by H.V. Klapdor-Kleingrothaus) has published a measurement of $m_{\beta\beta}$ which is not an upper limit, they report $m_{\beta\beta} = 440_{-200}^{+140} \text{meV}/c^2$ [29] (errors are 3σ). However, this result is very controversial and the scientific community is anxiously waiting for the next generation of double beta decay experiments to aid in resolving the issue. Additionally, neutrino-less double beta decay experiments can explore the mass hierarchy issue discussed earlier in the chapter. Observation of neutrino-less double beta decay is a highly desirable goal of the scientific community because of its power to define the nature of neutrino masses, their absolute mass, and the hierarchy. For a review of the current status of double beta decay experimental results, see [30].

⁶In practice, current experiments are only able to set an upper limit based on the non-observation of neutrino-less double beta decay. The limit on the number of events is just given by the Poisson probability to observe zero events.

⁷It should be noted that non-zero values for the Majorana and CP violating phases multiplying the U_{e2} and U_{e3} components will make m_{ee} smaller than the weighted sum of the mass components $m_{1,2,3}$.

2.3.3 The value of θ_{13}

Throughout this chapter, the assumption has been made that $\theta_{13} = 0$ because atmospheric experiments such as SuperKamiokande do not see an appreciable appearance signal from $\nu_\mu \rightarrow \nu_e$ oscillations⁸ but do see a $\sim 60\%$ reduction in the ν_μ flux. Therefore, one must have $\nu_\mu \rightarrow \nu_\tau$ (or $\nu_\mu \rightarrow \nu_s$ where ν_s is a sterile neutrino — however, this scenario is excluded by exploiting the fact that ν_μ, ν_τ undergo neutral current interactions in the earth matter, but ν_s do not undergo charged or neutral current interactions). Reactor oscillation experiments as well as short baseline oscillation experiments also set limits on the value of θ_{13} which give a slight indication for a non-zero value, albeit with significant errors.

If we now consider a non-vanishing value for θ_{13} , a two-neutrino mixing picture for solar neutrino oscillations is incomplete. For low-energy neutrinos such as the ${}^7\text{Be}$ neutrinos, matter oscillations do not cause significant deviations from the vacuum regime, and for the propagation distances involved the \sin^2 term dependent on L/E has been smeared to the point where we can take it to be equal to $\langle \sin^2 \rangle = \frac{1}{2}$. With these approximations in mind, the survival probability is equal to:

$$P_{ee} = P(|\nu_e\rangle \rightarrow |\nu_e\rangle) = \sin^4(\theta_{13}) + \cos^4(\theta_{13}) \left\{ 1 - \frac{1}{2} \sin^2(2\theta_{12}) \right\} \quad (2.61)$$

For the mixing parameters presented in Table 2.2, the inclusion of θ_{13} in calculating the survival probability amounts to a 2% change (reducing it from 0.569 to 0.558) in the survival probability. This difference is well below the current errors on survival probabilities for low-energy neutrino data from Borexino, however, it is an independent way to check values

⁸This of course is the dominant oscillation scheme for *solar* neutrinos. However, the atmospheric neutrinos are produced in the earth's atmosphere and so the distance L from the production point to the detector is at most $2R_E + h \ll 1A.U.$ where R_E is the radius of the earth, and h is the height above the surface of the earth where it was produced. The neutrinos are produced with energies on the order of 1GeV, thus, the factor L/E in the oscillation probability is vanishingly small and the oscillation to electron flavor is negligible.

obtained from reactor or short baseline oscillation experiments.

Still, the experimental determination of θ_{13} only affords an upper limit, and so our neutrino oscillation pictures are incomplete. If θ_{13} is equal to zero, then the CP violating phase δ cannot be measured since the only terms which contain it are products with $\sin(\theta_{13})$ (see the second matrix in equation 2.16). A new breed of experiments with the goal of measuring θ_{13} are nearing the data taking phase — most notably, the Daya Bay experiment in China and the Double-Chooz experiment in France. Both of the aforementioned experiments are reactor disappearance experiments using two (or more) detectors at two different baselines from the reactors and aim for sensitivities of ~ 0.01 to θ_{13} ([31], [32]).

Independent measurements of θ_{13} are also possible from off axis accelerator experiments such as NO ν A [33]. Placing the detectors off axis provides a neutrino flux more uniform in energy than on axis, albeit at a smaller flux. Since these experiments can produce both ν_μ and $\bar{\nu}_\mu$, they can search for CP violation by looking for an asymmetry in the rates of $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations.

2.3.4 Elemental Abundances

A table containing the solar neutrino fluxes for the reactions of the pp and CNO cycles is presented in Table 2.3. The different columns for the fluxes assume different values for the metallicity⁹ as evaluated in three separate studies in 1998, 2005, and very recently in 2009. The solar model based on the solar abundances obtained from the study performed by Grevesse and Sauval in 1998 (GS98 [34]) is commonly referred to as the “Hi Z” model and is the only one which is in agreement with the helioseismology measurements performed by Basu and collaborators[35].

⁹In solar astrophysics, the mass fraction of hydrogen is labeled as X , the mass fraction of helium is labeled as Y , and the summed mass fraction of all other elements is known as the metallicity, Z

Type	Energy (MeV)	Flux			Change from GS98 (%)		Fraction (%)
		GS98	AGS05	AGS09	AGS05	AGS09	
pp	≤ 0.420	5.97	6.04	6.03	1.17	1.01	91.3
pep	1.442	1.41	1.44	1.44	2.13	2.13	0.216
${}^7\text{Be}$	0.862^{10}	5.08	4.54	4.64	10.63	8.66	7.76
${}^8\text{B}$	≤ 15	5.88	4.66	4.85	20.75	17.52	0.009
hep	≤ 18.77	7.91	8.24	8.18	4.17	3.41	10^{-5}
${}^{13}\text{N}$	≤ 1.199	2.82	1.85	2.07	34.40	26.6	0.431
${}^{15}\text{O}$	≤ 1.732	2.09	1.29	1.47	38.28	29.67	0.319
${}^{17}\text{F}$	≤ 1.740	5.65	3.14	3.48	44.42	38.41	0.009

Table 2.3: Energy and fluxes for the various neutrino sources under the assumption of three different solar abundances: the GS98 “Hi-Z” data set, the AGS05 “Low-Z” data set, and the new AGS09 data set which has a metallicity intermediate between the two. The huge difference in the CNO neutrino fluxes predicted with the different solar abundances is now being referred to as the “solar composition problem”. The units (in $\text{cm}^{-2}\text{s}^{-1}$) for the fluxes are: pp - 10^{10} ; ${}^7\text{Be}$ - 10^9 ; (pep, ${}^{13}\text{N}$, ${}^{15}\text{O}$) - 10^8 ; (${}^8\text{B}$, ${}^{17}\text{F}$) - 10^6 ; hep - 10^3 . The last column shows the fraction of the total solar neutrino spectrum comprised by each source — the reference data was arbitrarily chosen to be the fluxes determined from the GS98 abundances. Flux data obtained from [36].

The typical process for constructing a solar model is to evolve a main sequence star to the age of the sun under the constraints that after the lifetime of the sun, it has reached the present size, luminosity, and current ratio of metallicities Z/X in the photosphere (the present day photospheric abundances are assumed to reflect the initial solar abundances). This is of course a *vastly* oversimplified description of the construction of a solar model, but it illustrates the relevance of the solar abundances.

The study performed by Asplund, Grevesse, and Sauval published in 2005 (AGS05 [37]) initiated what is now referred to as the “solar abundance problem” — and is immediately evident from looking at the “Change from GS98” columns in Table 2.3. The AGS05 model is known as the “Low-Z” model due to a change in the photospheric ratio Z/X of almost 30%, the net result of which is a lower core temperature for the AGS05 solar models. The Beryllium and especially the Boron solar neutrino fluxes are highly dependent on the core

temperature ($\Phi(^8B) \propto T^{\sim 18}$), thus the neutrino fluxes dropped considerably for these two neutrino sources for the solar models based on these results. Since less energy is released by the ^7Be solar neutrinos for the lower metallicity case, this change must be compensated by an increase in the other prominent reaction, the pp neutrinos.

A recent update to the solar abundances published by Asplund, Grevesse, Sauval, and Scott in 2009 (AGS09 [38]) determines a value for Z/X which is slightly higher than the AGS05 numbers, but is still $\sim 22\%$ lower than the GS98 abundances. The solar abundances may be determined by two complementary methods: one is via the mass spectroscopy of meteorites (assumed to have the same composition as the sun), the other is through identification of the Fraunhofer lines in the visible spectrum from the photosphere. Until recently, the meteoritic method was preferred over the photospheric one due to problems with the accuracy of the photospheric method¹¹. However, in the AGS09 paper, the authors published Z/X for both meteoritic and photospheric abundances and they were in excellent overall agreement. However, the individual elemental abundances for some of the heavier elements do show a difference from the meteoritic abundances — see reference [36] for a more complete discussion on the flux differences between the two sets of abundances.

The crux of the solar abundance problem is that the AGS09 solar abundance model would presumably reflect the current best-estimate of the metallicity of the sun, however, these data are at odds with the independent measurements of the density structure of the solar interior by way of helioseismology[35]. The solar model constructed with AGS05 metallicities has a 5.5σ discrepancy for the surface helium abundance (Y_S) relative to helioseismological results, and a 15σ discrepancy for the depth of the convective envelop (R_{CZ}/R_\odot) — the differences have decreased to 5σ and 11σ [36] for the solar model based on the AGS09 abundances. However, a disagreement at that level is a bit difficult to swallow; the differences in the

¹¹Section 1 of [38] contains a more complete discussion of the pros and cons of the meteoritic and photospheric methods of evaluating the solar abundances.

neutrino fluxes are as large as $\sim 40\%$ for components of the CNO cycle, and $\sim 20\%$ for the ^8B neutrinos in the pp chain. Solar neutrinos now play a very important role in that they provide the only method currently available to distinguish between the different abundances.

Figure 2.11 shows the predicted ^7Be and ^8B ν fluxes using the different solar abundance data sets, including model uncertainties. Also presented in the plot is the ^7Be data from Borexino and an averaged value for the ^8B flux over all of the SNO phases — both including 1σ error bars. From the figure, it is clear that we have a long way to go in being able to discriminate between the solar abundances by using ^7Be and ^8B solar neutrinos as proxies. Figure 2.12 shows the same thing except it correlates the ^8B flux with $\Phi(^{13}\text{N}) + \Phi(^{15}\text{O})$. It is clear then that the CNO neutrinos will more easily allow us to confirm or exclude the GS98 abundances and come closer to solving the solar abundance problem. By being able to refine the current estimates of the CNO contribution to solar burning, Borexino is in a unique position to aid in resolving the solar abundance problem.




Figure 2.11: Theoretical predictions of the ${}^7\text{Be}$ and ${}^8\text{B}$ solar neutrino fluxes for the four different metallicities (AGS09 uses the meteoritic abundances, AGS09ph uses the photospheric abundances) with model uncertainties. The central values (with 1σ errors) from Borexino and SNO are indicated by the black dashed line. It is easy to see that at the current uncertainty levels it is essentially impossible to consider one model better than any of the others. Figure from [36].




Figure 2.12: The same as Figure 2.11 except the ordinate axis contains the summed contribution of the neutrinos from ${}^{13}\text{N}$ and ${}^{15}\text{O}$, and the shaded grey band represents the SNO ${}^8\text{B}$ measurement with 1σ error bars. As we can see, the CNO neutrinos (measurable by Borexino) provide a slightly more powerful discrimination between the solar abundance models.

Chapter 3

The Borexino Detector

At the time of this writing, the Borexino detector is about to enter its third (out of ten) year of continuous data taking. Originally proposed by R.S. Raghavan in 1989, the BOREX[39] experiment was going to be a detector filled with several kilotons of a trimethylborate scintillator capable of measuring the Boron 8 solar neutrinos. Neutral current interactions would lead to nuclear de-excitations of the ^{11}B in the scintillator, while the charged current neutrino capture reactions would lead to production of the positron emitting ^{11}C . The Borexino experiment was to be the 100 ton fiducial mass prototype for the BOREX experiment. However, it was quickly realized that if a scintillator could be purified sufficiently, an experiment such as Borexino with a 100 ton fiducial mass could be an effective detector for the much higher rate ^7Be neutrinos.

A one ton fiducial mass prototype, the Counting Test Facility (CTF) — which will be discussed in more detail in Section 3.4.4 — was constructed in the mid 1990s with the task of demonstrating the feasibility of large scale purification of organic liquid scintillators to the levels required for ^7Be ν detection. After successful demonstration of the purification techniques used, construction of the Borexino detector began in 1998. Finally, in late 2006,

the filling process of Borexino began, first with water, and then with scintillator. After nine months of around-the-clock filling, data taking began with a full detector on May 15, 2007.

3.1 Neutrino Detection & Their Sources in Borexino

The detection method employed in the Borexino detector is that of neutrino - electron scattering which proceeds via the charged and neutral current processes depicted by the Feynman diagrams in Figure 3.1. The charged current interaction, mediated by the charged W^\pm vector bosons, is possible only for electron flavor neutrinos in the normal electronic, matter present in our detectors. The neutral current interaction, mediated by the Z^0 boson is sensitive to all three flavors of neutrinos, albeit with a smaller cross section.




Figure 3.1: First-order Feynman diagrams for the neutral (left) and charged (right) current neutrino-electron scattering interaction. Note that the charged current interaction is only possible for electron neutrinos, therefore the electron neutrinos have higher probabilities for scattering than the μ or τ flavors.

In a neutrino-electron scattering experiment such as Borexino, the neutrino interacts with an electron in a scintillator molecule (typical binding energy of a few eV, so, without loss of generality, we can take it to be free) which then produces scintillation light from excitation and ionization of other scintillator molecules which it encounters. This situation is analogous to Compton scattering, and the recoil-electron spectrum is a continuum with a sharp edge

occurring at

$$E_{max} = \frac{2E_\nu^2}{2E_\nu + m_e c^2} \Rightarrow E_{max}(E_\nu = 862\text{keV}) \approx 667\text{keV} \quad (3.1)$$

The overall shape of the neutrino-induced recoil-electron energy spectrum recorded by an ideal detector is given by the differential cross section:

$$\frac{d\sigma}{dE} = \frac{2G_F^2 m_e}{\pi \hbar^4 c^2} \left\{ g_l^2 + g_r^2 \left(1 - \frac{E}{E_\nu} \right) - g_l g_r \frac{m_e c^2 E}{E_\nu^2} \right\} \quad (3.2)$$

$$g_r = \sin^2(\theta_w) \approx 0.222 \quad g_l = \begin{cases} \sin^2(\theta_w) + \frac{1}{2} \approx 0.722 & , \nu_e \\ \sin^2(\theta_w) - \frac{1}{2} \approx -0.278 & , \nu_{\mu,\tau} \end{cases}$$

In equation 3.2, the leading constant is equal to $8.81 \times 10^{-45} \text{cm}^2$, and θ_W , the Weinberg mixing angle, is related to the ratio of the masses of the W and Z bosons: $\theta_W = \arccos(M_W/M_Z) \approx 28^\circ$. Note that the value of g_l for electron flavor neutrinos leads to an overall larger cross section relative to the heavier flavors; this is due to the fact that electron neutrinos may interact via the charged and neutral current interactions, whereas the heavier flavors can only interact through the neutral current interaction. The total cross section is then given by the integral of the differential cross section from 0 up to E_{max} :

$$\sigma = \frac{2G_F^2 m_e E_{max}}{\pi \hbar^4 c^2} \left\{ (g_l^2 + g_r^2) - \left(\frac{g_r^2}{E_\nu} + g_l g_r \frac{m_e c^2}{2E_\nu^2} \right) E_{max} + g_r^2 \frac{E_{max}^2}{3E_\nu^2} \right\} \quad (3.3)$$

3.1.1 ${}^7\text{Be}$ Neutrinos

For ${}^7\text{Be}$ solar neutrino energies¹, the total cross sections derived from equation 3.3 are:

$$\sigma(\nu_e) = 5.81 \times 10^{-45} \text{cm}^2 \quad (3.4)$$

$$\sigma(\nu_{\mu,\tau}) = 1.30 \times 10^{-45} \text{cm}^2 \quad (3.5)$$

$$\frac{\sigma(\nu_e)}{\sigma(\nu_{\mu,\tau})} = 4.48$$

The neutrino analysis performed with an experiment such as Borexino is derived from an energy spectrum, where the contents of each bin are given by (assuming 100% detection efficiency):

$$\frac{dN}{dE} = n_e V T \Phi_\nu \left(\frac{d\sigma}{dE} \otimes R(E, \sigma_0) \right) \quad (3.6)$$

The argument in parenthesis in the above equation represents the convolution of the ideal differential cross section with a resolution function (which we assume to be Gaussian — σ_0 is the width of this Gaussian), n_e is the electron density, V is the volume, T is the livetime, and Φ_ν is the neutrino flux at the detector.

The rate of neutrinos is then just the integral of equation 3.6 divided by the livetime T :

$$\frac{N}{T} = n_e V \Phi_\nu \int_0^\infty \left(\frac{d\sigma}{dE'} \frac{1}{\sqrt{2\pi\sigma_0^2}} e^{-\frac{(E-E')^2}{2\sigma_0^2}} \right) dE' \quad (3.7)$$

Obviously, for neutrino sources other than ${}^7\text{Be}$ and pep, the neutrino flux itself is energy

¹Here we deal only with the 862 keV line as the Compton edge from the 380 keV line (with a 10% branching ratio) is hidden by the ${}^{14}\text{C}$ background

dependent and must also go inside of the integral. To obtain a rough estimate of the neutrino interaction rate in Borexino, we may simply not perform the convolution² and use the result in equation 3.3. The Borexino scintillator has $2.89 \times 10^{29} e^-/m^3$, and the fiducial volume contains roughly $114m^3$ of active scintillator; the current (according to the BPS08 high-Z solution to the solar model [40] — the low Z solution gives a flux lower by about 10%) value for the ${}^7\text{Be}$ solar neutrino flux is $(5.07 \pm 0.304) \times 10^9 \nu/cm^2/s$, of which, only $\approx 89\%$ are the 862 keV neutrinos. The expected rate of ${}^7\text{Be}$ ν interactions in the Borexino inner vessel is then 75 per day if there are no neutrino oscillations, and about 17 per day if all neutrinos oscillate to μ or τ flavors. For an electron neutrino survival probability of $P_{ee} = 0.55$ (which is typical of ${}^7\text{Be}$ solar neutrinos in the current theory), Borexino will see about $P_{ee}\Phi_{\nu_e} + P_{e,\mu/\tau}\Phi_{\nu_\mu/\nu_\tau} = 75 \cdot 0.55 + 17 \cdot 0.45 = 48$ interactions per day in the fiducial volume³.

Čerenkov based detectors are fortunate in that it is relatively easy to verify the solar nature of the neutrinos simply by projecting the light cone from the recoil-electron back to its vertex, the axis of this cone, on average, points towards its source. Neutrino-electron scattering detectors however, have to rely on the fact that distance from the earth to the sun varies over the course of the year. The earth's elliptical orbit about the sun has an eccentricity, ϵ , of 0.0167 and a semi-major axis, a , of 1.496×10^8 km — see Figure 3.2.

For any elliptical orbit, the eccentricity is related to the radii at aphelion and perihelion by:

$$\epsilon = \frac{r_{aph} - r_{per}}{r_{aph} + r_{per}} \quad (3.8)$$

²This is allowable because the convolved spectrum is very similar in shape to the ideal spectrum

³Borexino has a very large background of ${}^{14}\text{C}$ which extends up to about 250 keV due to pile up and resolution smearing (see section 3.4.2), so, we are really only sensitive to the region above 250 keV




Figure 3.2: Pictorial representation of the earth-sun orbit (eccentricity exaggerated for clarity). The semi-major axis of the orbit is $(r_{aph} + r_{per})/2 = 1.496 \times 10^8$ km, and the eccentricity is 0.0167. The orbital radii (with respect to the center of the sun) at perihelion and aphelion are indicated.

$$\begin{aligned} &= \frac{\left(\frac{r_{aph}}{r_{per}} - 1 + 2\right) - 2}{1 + \frac{r_{aph}}{r_{per}}} \\ \epsilon &= 1 - \frac{2}{1 + \frac{r_{aph}}{r_{per}}} \end{aligned} \quad (3.9)$$

The flux at the earth is inversely proportional to the squared earth-sun distance, thus, the flux at perihelion will be larger by the ratio of the distances at aphelion to perihelion squared:

$$\Phi_{per} = \Phi_{aph} \frac{r_{aph}^2}{r_{per}^2} \Rightarrow \left(\frac{r_{aph}}{r_{per}}\right)^2 = \left(\frac{1 + \epsilon}{1 - \epsilon}\right)^2 \approx 1.07 \quad (3.10)$$

Thus, Borexino will observe a 7% variation in the solar neutrino flux from January to July, and back again, each year. Over the course of one period, $T = 1$ year, the earth - sun distance varies as:

$$D(t) = \frac{a(1 - \epsilon^2)}{1 + \epsilon \cos\left(\frac{2\pi t}{T}\right)} = a \left\{ 1 - \epsilon \cos\left(\frac{2\pi t}{T}\right) \right\} + O(\epsilon^2) \quad (3.11)$$

Since $\epsilon \ll 1$, we can ignore terms larger than first order. If we take the average neutrino rate to be \bar{R} , and the average distance to be $\bar{D} = a$, then we can express the time dependent solar neutrino rate as:

$$R(t) = \bar{R} \left(\frac{D(t)}{\bar{D}} \right)^2 = \bar{R} \left\{ 1 + 2\epsilon \cos\left(\frac{2\pi t}{T}\right) + O(\epsilon^2) \right\} \quad (3.12)$$

If we take $\bar{R} = 49\nu/\text{day}$, then the solar neutrino rate will oscillate over the course of three years as depicted in Figure 3.3.




Figure 3.3: Annual variation of the ${}^7\text{Be}$ ν solar neutrino rate as a function of time over a period of three years (equation 3.12, with $\bar{R} = 49$). The peak to peak difference due to the elliptical nature of the earth's orbit about the sun is 7% and is the only signature of the solar nature of the neutrino flux.

A deviation significantly different than the predicted 7% would be a sign of neutrino oscillations on a distance scale less than $r_{aph} - r_{per} = 2\epsilon a \approx 5 \times 10^6$ km, or regeneration when passing through the earth (the latter would show up as an excess of events during the night

as compared to the rate during the day).

Normalizing LENS to Borexino

The LENS experiment has the task of measuring the charged current neutrino flux from the sun via the neutrino capture reaction: $\nu_e + {}^{115}\text{In} \rightarrow e^- + 2\gamma(\tau = 4.76\mu\text{s}) + {}^{115}\text{Sn}$ which has a threshold of 114 keV. One of the uncertainties that dominates its predicted error is that of the cross section for the neutrino capture (measured to within 10%). Performing a precision flux measurement would require circumventing this error by calibrating the detector with a known source of neutrinos. Calibrating LENS can be accomplished with a MCi neutrino source, *or*, one can exploit the fact that with a threshold of just 114 keV, LENS will readily see the same ${}^7\text{Be}$ solar neutrino flux that Borexino does.

Recall from equation 3.7 that the rate of neutrino interactions (per unit volume) is given by:

$$R = \frac{N}{TVn_e} = \Phi_\nu\sigma \quad (3.13)$$

Because Borexino and LENS detect neutrinos by different interactions, the σ in equation 3.13 will be different, but the flux (Φ_ν) is the same, therefore, we have:

$$\Phi_\nu = \frac{R_{BX}}{\sigma_{BX}} = \frac{R_{LENS}}{\sigma_{LENS}} \quad (3.14)$$

Because LENS exploits a neutrino capture reaction, it is only sensitive to ν_e (neutrinos of solar energies do not have sufficient energy to produce a μ in the final state), so, σ_{LENS} in equation 3.14 is the product of the cross section for ${}^7\text{Be}$ neutrino capture on ${}^{115}\text{Sn}$ with the electron neutrino survival probability: $\sigma_{CC}P_{ee}$.

On the other hand, Borexino exploits the neutrino-electron scattering reaction which is sensitive to all active neutrino flavors. For this reason, σ_{BX} in equation 3.14 is a sum of two terms:

- The product of the charged-current scattering cross section with the electron neutrino survival probability: $\sigma_{CCS}P_{ee}$.
- The product of the neutral-current scattering cross section with the survival probability for non- ν_e flavors: $\sigma_{NCS}(1 - P_{ee})$.

Numerical values for σ_{CCS} and σ_{NCS} are given in equations 3.4 and 3.5 respectively⁴. Thus, equation 3.14 becomes:

$$\frac{R_{LENS}}{P_{ee}\sigma_{CCC}} = \frac{R_{BX}}{P_{ee}\sigma_{CCS} + (1 - P_{ee})\sigma_{NCS}} \quad (3.15)$$

$$\sigma_{CCC} = \frac{R_{LENS}}{R_{BX}} \left\{ \frac{P_{ee}\sigma_{CCS} + (1 - P_{ee})\sigma_{NCS}}{P_{ee}} \right\} \quad (3.16)$$

The rates, survival probability, and cross sections in equation 3.16 are evaluated at the ${}^7\text{Be}$ ν energy, however, if one knows how the cross section depends on energy, σ_{CCC} can be determined at pp energies. For the desired 3% measurement with LENS, the error on the capture cross section can be reduced to an acceptable level if Borexino is able to measure the ${}^7\text{Be}$ ν flux to 3%. For this reason, it is highly desirable for the solar neutrino physics program of LENS that Borexino perform a 3% measurement of the ${}^7\text{Be}$ ν flux. Doing so requires a thorough calibration of the detector to reduce and understand the systematic

⁴The subscripts CCC, CCS, and NCS mean: Charged Current Capture, Charged Current Scattering, and Neutral Current Scattering respectively.

errors. The hardware, method, and results of this calibration are presented in the latter half of this dissertation.

3.1.2 CNO and pep neutrinos

In addition to the ${}^7\text{Be}$ solar neutrinos, Borexino is also sensitive to the pep and CNO neutrinos with interaction rates of 2.75 and 5.2 neutrinos per day in the fiducial volume respectively. Table 3.1 summarizes the errors on the CNO fluxes in the BPS08 standard model with both high and low metallicities:

Source	BPS08(GS98) Hi Z	BPS08 (AGS05) Lo Z	% Diff
${}^{13}\text{N}$	$2.88 (1 \pm 0.15) \times 10^8$	$1.89 (1^{+0.14}_{-0.13}) \times 10^8$	34
${}^{15}\text{O}$	$2.15 (1^{+0.17}_{-0.16}) \times 10^8$	$1.34 (1^{+0.16}_{-0.15}) \times 10^8$	31
${}^{17}\text{F}$	$5.82 (1^{+0.19}_{-0.17}) \times 10^6$	$3.25 (1^{+0.16}_{-0.15}) \times 10^6$	44

Table 3.1: CNO neutrino fluxes and errors for two different metallicities with the BPS08 standard solar model[40].

The large errors on the CNO neutrino fluxes are largely dominated by the uncertainty in the solar composition for these elements — indeed, a measurement of the CNO neutrinos by Borexino would help to resolve the issue of this large discrepancy in solar chemical abundances (for a more in-depth discussion of the chemical abundances, see Section 2.3.4).

The CNO neutrinos are produced in inverse beta decay reactions (see Section 2.2.2) and have continuous spectra up to 1.199, 1.732, and 1.740 MeV for the neutrinos released in the decay of ${}^{13}\text{N}$, ${}^{15}\text{O}$, and ${}^{17}\text{F}$ respectively. The ${}^{17}\text{F}$ neutrinos have a flux two orders-of-magnitude lower than those for the ${}^{13}\text{N}$ and ${}^{15}\text{O}$ neutrinos and will be essentially impossible to see in Borexino.

The pep neutrinos compete with ν_{pp} as the first step of the pp chain, albeit with a flux smaller by a factor of ~ 420 . Thus, a measurement of the pep neutrinos accesses the same

information about the beginning of the pp chain as would a much more difficult measurement of the pp neutrinos. The pep reaction yields a monoenergetic, 1.44 MeV neutrino line which will produce a Compton edge at 1.22 MeV in Borexino. The energy of the pep line affords a unique opportunity to obtain a survival probability measurement in the transition region between the vacuum (${}^7\text{Be}$ energies) and matter (${}^8\text{B}$ energies) dominated regimes. By being able to measure survival probabilities in all three regions, Borexino will play an important role in testing the MSW-LMA oscillation scenario; furthermore, one or more of these measurements can be used to test the viability of various alternative neutrino oscillation models (mass varying neutrinos[41], non standard interactions[42], etc.).

The major obstacle for observation of the pep and CNO neutrinos is the decay of ${}^{11}\text{C}$ (which covers the energy region from 1.02 – 1.98 MeV). ${}^{11}\text{C}$ is produced primarily through inelastic collisions with ${}^{12}\text{C}$ nuclei when high energy secondaries (primarily γ or n) from cosmogenic muons pass through Borexino (with a flux of about $1.2/h \cdot m^2$ at the Borexino site). Fortunately, it is possible to tag and remove these events (see section 3.4.3) which cover the bulk of the pep and CNO neutrino energy window; the ability of Borexino to measure pep and CNO neutrinos is vitally linked to the ${}^{11}\text{C}$ subtraction efficiency.

3.1.3 ${}^8\text{B}$ Neutrinos

The large target mass and extreme radio-purity of Borexino also make it an ideal detector for measuring the high-energy ${}^8\text{B}$ solar neutrinos ($E_\nu \leq 16.3$ MeV). All of the previous ${}^8\text{B}$ solar neutrino spectrometers have been Čerenkov-based experiments; these detectors exhibit low light yields and thus require the imposition of energy thresholds that have typically been on the order of 5 MeV, thus, the low-energy spectral information is ultimately inaccessible. However, since Borexino is scintillator-based, the only limit for the detection of the

comparatively rare ^8B neutrinos comes from radioactive background.

The recoil electron spectrum produced by ^8B neutrinos extends from 0 to 16.3 MeV, but the naturally occurring ^{232}Th contained in the PMTs, SSS (Stainless Steel Sphere), vessels, and other materials, imposes a lower energy limit. The isotope ^{208}Tl , a product of the alpha decay of ^{212}Po in the thorium chain, decays by β emission with $Q = 5.001$ MeV. In addition to the β , in 99% of the decays, a 2.6 MeV gamma is released. It is this highly-penetrating gamma that demands a lower energy cut of 3.0 MeV (to account for the finite energy resolution) for the ^8B neutrino analysis.

Above the lower energy threshold, additional sources of background come from the high-energy tail of the decay of ^{214}Bi ($Q = 3.272$ MeV), ^{208}Tl events which occur inside of the FV (the β and γ are summed calorimetrically to give $Q = 5.001$ MeV), neutron capture gammas (usually on ^{12}C ; first excited state energy of 4.9 MeV) and cosmogenic isotopes (^{12}B , ^6He , ^{10}C , ^8B , ^8Li , etc.).

Assuming the BS07 high metallicity standard solar model (with the best fit MSW-LMA values for $\Delta m_{12}^2 = 7.69 \times 10^{-5}$ eV² and $\tan^2 \theta_{12} = 0.45$ from [43]) the expected full-spectrum count rate in Borexino is 0.49 ± 0.05 counts/day/100 tons. After removal of cosmogenics, ^{214}Bi , and statistical subtraction of ^{208}Tl , the resulting measurement of the ^8B solar neutrino flux is in good agreement with theory: $\Phi_{exp}^{ES}/\Phi_{th}^{ES} = 0.88 \pm 0.19$ for the 3.0-16.3 MeV analysis region for a 488 day data set. More detailed information on the analysis and results can be found in reference [44].

With this measurement, Borexino has provided a high-energy point on the electron neutrino survival probability curve (see figure 3.4) of $P_{ee} = 0.29 \pm 0.10$ for the median energy of 8.9 MeV, making it the first experiment to simultaneously measure survival probabilities in both the vacuum and matter-dominated oscillation regions with the same detector. In doing

so, Borexino has provided data points that are consistent with the expected suppression of survival probability in the matter enhanced regime at a level of 1.9σ . Further work is underway to reduce the systematic errors (predominantly fiducial volume determination) and to possibly expand the fiducial volume to the entire inner vessel which would increase the statistics by a factor of almost three, albeit at the expense of higher background.




Figure 3.4: Solar electron neutrino survival probability plot showing data measured by Borexino in the ${}^7\text{Be}$ analysis as well as the ${}^8\text{B}$ neutrinos (along with pp data from all solar experiments, and ${}^8\text{B}$ data from SNO). With the ${}^8\text{B}$ neutrino measurement, Borexino has become the first experiment to simultaneously measure survival probability in both the matter- and vacuum-dominated oscillation regions. Figure from [44].

3.1.4 Geo and supernova neutrinos

Modern theories of stellar evolution predict the occurrence of roughly three type II supernovae per century in our galaxy[45]; these predictions are supported by astrophysical observations over many recent decades. Poisson statistics dictates that Borexino has roughly a 26% chance of observing a supernova in its ten year lifetime (of course, the actual probability is less because the livetime is reduced due to calibrations and other DAQ outages).

Neutrinos released from supernovae

Nuclear burning in a star proceeds such that the core burns through progressively heavier elements until reaching ^{56}Fe , whereupon further fusion ceases to release any more energy. The core now ceases to undergo any further nuclear fusion⁵ and thus cannot sustain the pressure to counteract gravity, and it begins to contract and compress under its own weight.

The temperature in the core begins to rise rapidly due to its densification until it reaches a level such that photo-dissociation of the ^{56}Fe nuclei is possible — this process does create a higher number density of particles which, in principle, should increase the outward pressure, however, the dissociation of the iron nuclei is an endothermic reaction and so acts to reduce the outward pressure. The pressure loss due to dissociation of the iron is ultimately greater than that gained by the increase in number density of nuclei and hydrostatic equilibrium rapidly breaks down.

As the outer layers begin to collapse inward, electrons in the outer layers undergo electron capture on protons and nuclei, this process is known as “neutronization” and provides an initial burst of pure electron flavor neutrinos. The decrease in the number of electrons results in decreased electron degeneracy pressure which further reduces the outward force, and the collapse accelerates as the density of the core approaches the nuclear density.

At the nuclear density ($\sim 10^{14} \frac{\text{gm}}{\text{cm}^3}$, roughly the nucleon mass divide by a nucleon volume of a cubic fermi), the neutron degeneracy pressure increases and eventually becomes equal in magnitude to the inward gravitational force and provides the so-called “bounce” as the outer layers of the star are blown outwards. During this explosion, massive numbers of neutrinos are produced as ejected leptons undergo pair production processes to form neutrino - anti-neutrino pairs. The energy released in this explosion is usually assumed to be 2.5×10^{53}

⁵The outer layers are still burning progressively lighter elements.

ergs, over 99% of which is given off in the form of neutrinos [8].

During the period when the core density is above $\approx 10^{11} \frac{\text{gm}}{\text{cm}^3}$, even neutrinos become trapped through inelastic scattering inside of the “neutrino-sphere” (the spherical volume inside of which neutrinos are trapped). Due to the ability of $\nu_e/\bar{\nu}_e$ to interact via charged, as well as neutral, current interactions, the neutrino-sphere for electron flavor neutrinos is at a much smaller radius than for μ or τ flavors. The resulting effect is that the ν_e ’s that do reach the earth do so with a lower energy than ν_μ or ν_τ . Moreover, the process of neutronization creates a large imbalance of neutronic to protonic matter, thus ν_e have an abundance of capture targets with respect to $\bar{\nu}_e$, so, the neutrino-sphere for ν_e is at an even smaller radius than for $\bar{\nu}_e$. The net effect of the different reaction mechanisms is that the average energy for the supernova neutrinos reaching detectors on the earth is: $\bar{E}(\nu_e) = 11$ MeV, $\bar{E}(\bar{\nu}_e) = 16$ MeV, and $\bar{E}(\nu_\mu, \bar{\nu}_\mu, \nu_\tau, \bar{\nu}_\tau) = 25$ MeV.

Supernova Neutrino Interactions

When these neutrinos do make it to the Borexino detector, they may interact in a variety of ways, however, exact event rates are difficult to estimate and depend strongly on the energy spectrum of the arriving neutrinos. For the calculations described below, a supernova at 10 kpc (3.08×10^{17} km) with an energy release of 10^{53} ergs is assumed. Values for cross sections and expected event rates obtained from [46].

- Elastic scattering: $\nu + e^- \rightarrow \nu + e^-$. This mechanism is available to all neutrino and anti-neutrino flavors and does not have a reaction threshold. The cross section for elastic scattering is, to a good approximation, linearly dependent on energy, with the proportionality constant dependent on the leptonic flavor, but is on the order of $(10^{-45} E_\nu) \text{cm}^2$ (energy given in MeV). For the supernova characteristics assumed above,

Borexino should see approximately 5 events from elastic scattering on electrons.

- Anti-neutrino induced inverse beta decay: $p + \bar{\nu}_e \rightarrow n + e^+$. This mechanism is available only to $\bar{\nu}_e$ and has a reaction threshold of 1.8 MeV (equivalent to the $m_n + m_e - m_p$ mass difference). The cross section for this reaction can be approximately described as $9.5 \times 10^{-44} (E_\nu - 1.29)^2 \text{ cm}^2$. Once the cross section is integrated over the expected $\bar{\nu}_e$ spectrum, Borexino should expect about 80 neutrinos from this channel. This will be the dominant reaction mechanism for supernova neutrinos in Borexino.
- Neutrino capture on ^{12}C : $\nu_e + ^{12}\text{C} \rightarrow ^{12}\text{N} + e^-$. Since ^{12}C lies at the bottom of the $A = 12$ isobar mass parabola, the nuclei on either side of it are unstable to beta decay. The nitrogen created in this reaction is itself unstable to positron decay back into ^{12}C with a $\tau_{1/2}$ of 11 milliseconds. The delayed coincidence provided by the prompt electron and subsequent positron annihilation make tagging this reaction possible. The Q value for this reaction is 17.3 MeV, and has a cross section of $1.85 \times 10^{-43} \text{ cm}^2$ for a mean neutrino energy of 11 MeV. An estimated 0.6 neutrino events are anticipated for this channel.
- Anti-neutrino capture on ^{12}C : $\bar{\nu}_e + ^{12}\text{C} \rightarrow ^{12}\text{B} + e^+$. This reaction proceeds with a Q value of 13.4 MeV, and a cross section of $1.87 \times 10^{-42} \text{ cm}^2$ for a mean anti-neutrino energy of 16 MeV and is expected to produce 4.1 events in Borexino. The boron created herein beta decays back to ^{12}C with a $\tau_{1/2}$ of 20.2 milliseconds.
- Inelastic scattering of neutrinos on ^{12}C : $\nu_x + ^{12}\text{C} \rightarrow ^{12}\text{C}^* + \nu'_x$. The Q value for this reaction is 15.1 MeV, and will be characterized by the 15.1 MeV deexcitation gamma ray emitted in the M1 transition back to the ground state of ^{12}C . Note that this is a neutral current reaction and is thus possible for all neutrino, and anti-neutrino, flavors. Due to the large spread in energies (and thus mean energies) for different neutrino flavors,

the expected rates from neutral current interactions must be calculated separately for each flavor, the net result is 23 events.

In addition, since the natural isotopic abundance of ^{13}C is 1.1%, there is also the possibility for charged current neutrino capture in the $A = 13$ isobars, however, the event rate is predicted to be at the level of 0.06 events[47].

Neutrino mass limits

If we consider an ultra-relativistic neutrino (i.e. a very small mass), then $v = E/p$; applying the relativistic energy-momentum relationship, we may derive:

$$v = c\sqrt{1 + \frac{m^2c^4}{E_\nu^2}} \approx c\left(1 + \frac{m^2c^4}{2E_\nu^2}\right) \quad (3.17)$$

So, the time of arrival between two neutrinos is given as:

$$\Delta t = \frac{L}{c} \frac{m^2c^4}{2} \left(\frac{1}{E_2^2} - \frac{1}{E_1^2} \right) \quad (3.18)$$

For an experiment in which it is possible to distinguish between charged and neutral current interactions (and therefore ν_e from $\nu_{\mu,\tau}$) or ν_e from $\bar{\nu}_e$, one can obtain an upper limit on the mass of one or more neutrinos. For instance, the $\bar{\nu}_e$ are produced in very large quantities and pass more readily through neutronic matter on their way out of the star than do ν_e , thus they exhibit a narrower temporal dispersion. Additionally, $\bar{\nu}_e$ are easier to detect by virtue of their higher cross section for inducing inverse beta decay on protons. This technique was employed by the Kamiokande and IMB collaborations following supernova 1987A and an upper limit of $m_{\bar{\nu}_e} \leq 20$ eV was obtained[8].

Geo-neutrino Detection

Geo-neutrinos are comprised of the flux of the $\bar{\nu}_e$'s produced primarily by the β^- decays of the isotopes in the ^{238}U and ^{232}Th decay sequences, and the isotope ^{40}K . The detection mechanism in Borexino is induced inverse β decay via: $p + \bar{\nu}_e \rightarrow n + e^+$. The neutron will capture onto a proton, emitting a 2.2 MeV capture gamma in the process, and the positron will eventually annihilate with an electron, releasing 1.1 MeV via two annihilation gammas. These two signals allow a very efficient method of tagging geo-neutrino signals, efficient enough to allow almost the entire Borexino inner-vessel volume to be used as the target mass since external gammas cannot mimic the coincidence signal.

Due to the neutron-proton mass difference, there is a threshold on the incoming anti-neutrino energy of 1.8 MeV. This threshold makes the detection of anti-neutrinos from ^{40}K impossible due to the fact that the Q value for the β^- decay of ^{40}K is only 1.3 MeV. The expected $\bar{\nu}_e$ signal from the decay of ^{238}U , ^{232}Th , and ^{40}K in the earth can be seen in Figure 3.5.

A precise estimate of the expected geo-neutrino flux requires detailed knowledge of the geology surrounding the experimental site, as well as the flux of anti-neutrinos emitted from any nearby nuclear power plants⁶. Since the the local regions can contribute as much as 40% of the total geo-neutrino signal seen in a detector, a study of the Gran Sasso area was performed in order to measure the abundance of the radioactive isotopes ^{238}U , ^{232}Th , and ^{40}K . Combining these values for the local area with the values predicted by the leading model of the chemical abundance of the earth (the Bulk Silicate Earth, or BSE, model [49]) yields the expected geo-neutrino flux incident on Borexino.

The total number of expected geo-neutrino signals (we consider only ^{232}Th and ^{238}U due to the reaction threshold) in Borexino is given simply by the product of the cross section (\approx

⁶Borexino is very fortunate in that Italy does not currently have any operational nuclear power plants.




Figure 3.5: Theoretical spectrum of geo-neutrinos incident on terrestrial neutrino experiments. The dotted line at 1.8 MeV indicates the energy threshold for the $\bar{\nu}_e + p \rightarrow n + e^+$ reaction.[48]

10^{-43}cm^2), flux ($7.96 \times 10^6\text{cm}^{-2}\text{s}^{-1}$), exposure time ($3.15 \times 10^7\text{s}$) and number of targets ($\sim 1.7 \times 10^{31}$)⁷. Multiplying everything out, Borexino will see roughly seven geo-neutrinos per year — or 41 TNU⁸.

In 2005, the KamLAND detector in Japan released the first experimental results for geo-neutrino fluxes. Based on theoretical calculations, the expected signal was at the level of 35 ± 6 TNU — the KamLAND collaboration reported their findings to be 30 ± 11 .

The Borexino collaboration has recently (March, 2010) released its first geo-neutrino measurement in which we claim the observation of geo-neutrinos at a 99.997% confidence level⁹ [51]. The relatively small reactor anti-neutrino background near the Borexino site coupled

⁷The “targets” here are the protons in the hydrogen of the scintillator — the Q value to induce a reaction on ^{12}C is an order of magnitude larger than that of the geo-neutrinos.

⁸TNU (Terrestrial Neutrino Unit); corresponds to one event per year per 10^{32} targets

⁹Strictly speaking, this is the confidence level at which the null hypothesis is rejected.

Isotope	Local Abundance (ppm)	Predicted Flux ($\text{cm}^{-2} \text{s}^{-1}$)
^{238}U	0.95	4.22×10^6
^{232}Th	4.57	3.73×10^6
^{40}K	1.2	Not detectable ($Q_{^{40}\text{K}} \leq 1.8 \text{ MeV}$)
Total		7.96×10^6

Table 3.2: Abundance of geo-neutrino producing isotopes in a (222×165) km area around Gran Sasso and the predicted flux in Borexino [50]

with the unprecedented radio-purity of the detector afforded a geo- $\bar{\nu}_e$ to non- $\bar{\nu}_e$ ratio in the analysis sample of 23:1 — the same ratio for the KamLAND measurement was 1:4. This measurement highlights the sensitivity of Borexino to geo-neutrinos despite the small exposure time of 252.6 ton·yr.

The next decade will hopefully give rise to several much larger detectors capable of assisting KamLAND (one kiloton) and Borexino (0.3 kiloton) in the measurement of geo-neutrino fluxes:

- SNO+ — a one kiloton scintillator detector is anticipated to fill and begin their first stage of data taking in 2010[52]
- Hanohano — a 10 kiloton ocean based detector has been proposed and could begin construction in 2013[53]
- LENA — a 50 kiloton scintillator detector (part of the proposed European LAGUNA triad of detectors: LENA, GLACIER, and MEMPHYS) is undergoing site evaluations and by the end of 2011 should have decided on a site and a final detector design[54]

3.2 Detector Design

The Borexino detector is situated in Hall C of the Laboratori Nazionali del Gran Sasso in Assergi, Italy — approximately 100 miles east of Rome in the Apennine Mountains. The underground portion of the laboratory is located adjacent to the A24 highway tunnel linking Rome with the eastern shoreline of Italy; it is also covered by several kilometers of rock overburden which gives a shielding capacity of 3,800 MWE (Meters of Water Equivalent). Figures 3.6 and 3.7 depict the layout of the underground portion of the LNGS laboratory and the arrangement of the major components of Borexino in Hall C respectively.




Figure 3.6: Schematic of the underground portion of the LNGS laboratory. The direction of CERN is indicated due to the CERN to LNGS (CNGS) neutrino beam used for the OPERA experiment which sits directly in front of Borexino in Hall C.




Figure 3.7: The major components of the Borexino experiment in Hall C. From left to right: The Borexino water tank; The “Big Buildings” (East and West) which house the control room, DAQ, and portions of the purification system; the purification skids; the CTF detector; and the four scintillator storage tanks.

3.2.1 Structural components

The purity requirements of Borexino (explained further in Section 3.4) are such that a large quantity of buffer liquid is required to reduce the external gamma activity to an acceptable

level. The adopted design of Borexino is a series of concentric volumes, each with increasingly higher purity requirements — Figure 3.8 shows this onion-like structure in a cut-away view of the Borexino detector.




Figure 3.8: A partial section view of the Borexino experiment. See the text for further details about the relevant components.

Water Tank

The outermost layer of shielding is the water tank (WT, diameter of 18 meters) which holds over 2100 cubic meters of deionized water. Cosmic muons traversing the water tank will produce Čerenkov radiation which is subsequently detected by 208 photomultiplier tubes (PMTs) placed on the outer surface of the stainless steel sphere (SSS) and the floor of the water tank. In order to enhance light collection, the area between the water tank inner wall and the outside of the stainless steel sphere is covered with highly reflective Tyvek sheets.

For a detailed description of the outer muon detector see the PhD theses by D. D'Angelo and M. Wurm in references [55] and [47] respectively.

The overall height of the WT is limited by the height of Hall C, and as such, it was not possible to have the same quantity of shielding water below the SSS as there is above it. To compensate for the reduced shielding at the south pole of the SSS, there are two large steel plates underneath the WT to balance the shielding capability. Not depicted in Figure 3.8 is a small tunnel underneath the WT where one could place an intense ^{51}Cr neutrino source for neutrino magnetic moment and calibration studies[56].

Stainless Steel Sphere

Inside of the water tank is the 6.85 meter radius stainless steel sphere (SSS — see Figure 3.9) which acts as the support structure for the 2212 inner photomultiplier tubes. The SSS is constructed from 8 mm thick sheets of stainless steel in the northern hemisphere, and 10 mm thick sheets in the bottom hemisphere, and is supported by 20 massive legs. The SSS also has a three meter diameter opening (Figure 3.10) through which the vessels were inserted — the port is sealed via a large Helicoflex gasket.

Inner and Outer Vessels

Peering further into the detector we finally get to the inner and outer scintillator containment vessels. The goal we are seeking to achieve is the reduction of all internal and external backgrounds to the lowest levels possible, however, the detector components themselves (SSS, PMTs, etc.) contribute to the backgrounds, some in non-negligible ways. For instance, the glass bulbs of the PMTs contain a considerable amount of ^{40}K ; the stainless steel of the SSS contains some ^{238}U and ^{232}Th — these are examples of intrinsic radioactivity that cannot




Figure 3.9: The stainless steel sphere (radius 6.85 meters) before installation of the photomultiplier tubes.




Figure 3.10: A view of the PMTs through the three meter port on the SSS

be removed by cleaning or other surface treatments.

The aforementioned background radiation is at a level high enough that, if allowed to produce scintillation light, would overwhelm the data acquisition system. Therefore, the liquid in the region closest to the detector materials must be non-scintillating. Obviously then, we must have a physical barrier to separate the scintillating and non-scintillating volumes — this barrier is provided by the inner vessel (IV). The knowledge gained from the CTF campaigns led the collaboration to add a second concentric vessel in the outer region, the outer vessel (OV), as it was found to significantly reduce the rate at which radon was able to diffuse into the IV.

Due to the somewhat destructive nature of the scintillator cocktail used for Borexino, material compatibility requirements made a rigid acrylic or polycarbonate vessel (such as that found in SNO) impossible. Fortunately, the scintillator used in the three volumes separated by these two vessels are similar enough that they need only be able to support their own weight (under the assumption of isotropic temperature), and thus could be made very thin. Since Borexino performs its position reconstruction based on the arrival time of photons at

the PMTs, the vessel material also needs to have a nearly identical index of refraction as that of the scintillator itself.

Finally, the inner vessel in particular, must be very leak tight to avoid mixing the non-scintillating fluid with the active scintillator. The inner vessel was measured to have a leak rate of $< 5 \times 10^{-3} \frac{\text{cm}^3}{\text{s}\cdot\text{mbar}}$; the outer vessel, while not as critical since the fluids are identical on either side of it, was measured to have a leak rate $< 0.1 \frac{\text{cm}^3}{\text{s}\cdot\text{mbar}}$ ¹⁰[57].

Supporting the two vessels and maintaining their location inside of the detector was accomplished via a series of Tensylon ropes and belts. Each vessel has a series of ropes running from the top to the bottom, exactly half of which are used to hold the vessel “up” if it were trying to sink, and the other half hold the vessel “down” in the event that it tried to float. Another set of ropes encircle the vessel in the ϕ direction and are primarily used to keep the vertical ropes equally spaced. The vertical ropes are connected to load cells at either end of the vessel so the strain can be monitored at all times. These and many other components of the vessel are visible in Figure 3.11.

The vessel position can be monitored visually with the aid of the calibration cameras (see Chapter 5), however, this requires illuminating the entire detector with very high levels of light which requires turning off the high voltage to the PMTs first. This operation of cycling the PMT voltage is thought to be a catalyst for various problems with PMTs, so, this situation is avoided when possible. A less intense method is possible with the aid of 78 point light sources attached on the surface of the inner and outer vessels. These light sources are constructed from $\frac{3}{32}$ ” (2.4 mm) diameter Teflon spheres into which quartz fibers have been glued and attached to the vessels. Figure 3.12 shows one of the spheres under illumination from a green laser in air, and Figure 3.13 is a photo taken with the calibration cameras when inside of Borexino. These light sources have also been illuminated with light from a 405 nm

¹⁰The pressures here are the differential pressures between adjacent volumes




Figure 3.11: The southern end of the inner vessel assembly in the clean room at Princeton University. This photo was taken before the inner vessel was nested inside of the outer vessel[57].

laser, to simulate scintillation light, with the DAQ on and used as fixed calibration points. Unfortunately, the fibers used to channel the light did not have an opaque jacket and light leakage from the fibers made position reconstruction difficult.

The outer vessel was fabricated from Capron B737ZP¹¹ nylon (a homopolymer of nylon 6) extruded into a film 125 μm thick. The inner vessel, however, is fabricated from a mixture of 83% Sniamid¹² ADS40T and 17% Ultramid¹³ B4 pellets, also 125 μm thick[57]. The construction of the vessels for all three CTF campaigns as well as for Borexino took place in a dedicated cleanroom at Princeton University — see Figures 3.14 and 3.15 for photos of the vessels during the construction and testing stages. For an extensive review of the Borexino

¹¹Manufactured by Allied Signal / Honeywell

¹²Produced by Nyltech

¹³Supplied by BASF




Figure 3.12: One of the Teflon diffuser balls placed on the inner and outer vessels as point light sources to allow monitoring of the vessel shape without requiring the PMT high voltage to be cycled. The fiber was attached by first inserting into a hole in the Teflon sphere, then loosely secured with some nylon and formic acid paste, then several coatings with Nafion polymer, and several dippings in the nylon and formic acid paste[58].

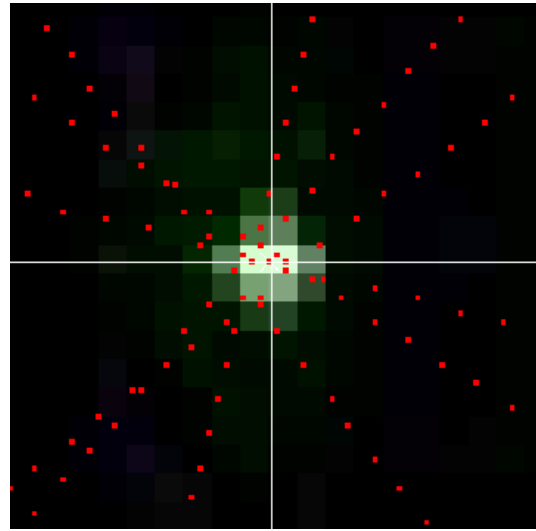


Figure 3.13: A photo of one of the inner vessel diffuser balls (seen in air in Figure 3.12) taken by one of the CCD cameras during a special calibration run in July 2007. The dotted red lines in the picture show the rays from the other six cameras to the location where each camera found the diffuser ball. A least squares minimization is done for all seven cameras to find the best fit position for the location of the source.

vessels see reference [59].

3.2.2 Scintillator

The heart of Borexino is the 315 cubic meters of *active* scintillator contained inside of the inner vessel. The base chemical for both the active and inactive scintillators is 1,2,4 trimethylbenzene (Pseudocumene, PC). The active scintillator inside of the inner vessel contains a fluor / wavelength shifter in the form of diphenyloxazole (PPO) at a concentration of 1.5 g/L. The inactive scintillator is contained in the inner and outer buffers and is comprised again of



Figure 3.14: Photo taken during the construction of the vessels for Borexino. The vessels were constructed from tear shaped panels, glued and clamped together — the IV was then nested inside of the OV prior to sealing the last OV panels. John Bahcall (middle person on the left side of the table) can be seen on a visit to the cleanroom where the construction was taking place.



Figure 3.15: Prototype inner and outer vessels undergoing a test inflation in the Jadwin Gymnasium at Princeton University.

pseudocumene, but with the addition of the scintillation quencher dimethylphthalate (DMP) at an initial concentration of 5 g/L. An important fact is that these two solutions have a very small density difference (less than one part per thousand), thus, the nylon inner vessel will not be under any considerable strain, assuming a constant and uniform temperature. Table 3.3 lists some of the important physical parameters of the chemicals comprising the Borexino scintillator. The chemical structure of these three molecules can be seen in Figures 3.16, 3.17, and 3.18.

Chemical	Formula	Specific Gravity	Boiling Point (°C)
Pseudocumene (PC)	$C_6H_3(CH_3)_3$	0.875	168
Dimethylphthalate (DMP)	$C_{10}H_{10}O_4$	1.19	283.7
Diphenyloxazole (PPO)	$C_{15}H_{11}NO$	1.06	360

Table 3.3: Some physical properties of the chemicals used in the Borexino scintillator [60], [61], [62].

The radioactivity analysis of the photomultiplier tubes require that the buffer liquid not




Figure 3.16: Pseudocumene




Figure 3.17: PPO




Figure 3.18: DMP

scintillate so as not to overwhelm the data acquisition. There are of course an abundance of non-scintillating chemicals from which one could choose, however, the design of the nylon vessels dictate that the density difference between the two fluids be kept as small as possible. Additionally, the analysis becomes much easier if the two volumes have the same indices of refraction so that the light propagation is easier to model. With these considerations, and the experience gained from the CTF campaigns, the inner and outer *buffers* were initially comprised of PC with 5 g/L of DMP. In 2008, a leak was discovered in the Borexino inner vessel and it became necessary to reduce the density difference between the inner vessel and the buffers to as small a level as possible. The purification skids were used in loop-mode to distill out some of the DMP until the average concentration was reduced to a level of 2 g/L. At this level, the DMP is still effective in quenching the scintillation light, but the upward buoyant force on the inner vessel is reduced by a factor of ~ 7 . The lower buoyant force reduces the ΔP across the film at the site of the leak and therefore the leak rate.

Scintillation light production with organic liquid scintillators

The scintillation process in organic molecules results from excitation of the electrons in the π bonds of the scintillator molecules. For an aromatic hydrocarbon such as pseudocumene, each carbon has four valence electrons which bond via trigonal sp^2 hybridizations. For

non-hybridized carbon, the electron configuration is $1s^2 2s^2 2p_x^1 2p_y^1$, but the sp^2 hybridization causes the electrons to assume the $1s^2 sp^2 sp^2 sp^2 2p_z^1$, the $2p_z^1$ electrons¹⁴ lie in a plane perpendicular to the those in the hybridized orbitals. When six carbon atoms come together in a ring, each adjacent pair is brought close enough together that the $2p_z$ orbitals “bend” towards each other and bond in a so-called π bond (the electrons in the hybridized orbitals bond together in a σ bond). The electrons in the π bonds are *delocalized* and are essentially free to move around the ring of carbon atoms like a free electron in a metal (a cartoon of this process is shown in Figure 3.19). It is these π electrons which give rise to the luminescence we know as scintillation.




Figure 3.19: Pictorial representation of σ and π bonds between two carbon atoms — the photo[63] is for ethene, but the mechanism is the same in pseudocumene. When six carbons come together in a hexagon, each adjacent pair forms a temporary π bond that can travel from carbon to carbon (not unlike the motion of a slinky) — it is these bonds that are responsible for producing scintillation light.

Incident radiation causes these molecules to be excited by either ionization or collisional excitation. In either case, luminescence occurs when a π electron which has been excited to higher energy level falls back down to the ground state with the emission of a photon.

Recall that in a π orbital, there are two electrons in a singlet state with total spin angular

¹⁴The choice of p_z is arbitrary

momentum of zero (their spins are anti-aligned). When one of these electrons is excited, it is possible for its spin to flip in this process, thus yielding a triplet state where it, and its counterpart, have their spins aligned. This sort of transition is heavily suppressed and is dubbed “spin forbidden”, and for this reason, direct excitation into the triplet states almost never occurs. However, the triplet states may be populated during recombination of ionized molecules and free electrons.

A diagram of the energy level structure in an organic scintillator is shown in figure 3.20. In addition to the principal energy levels shown, there are vibrational sublevels (S_{0x}, S_{1x} , etc), the level splittings for these are on the order of $\sim 0.16\text{eV}$ [64]. For a scintillator at room temperature $k_bT \approx \frac{1}{40}\text{eV} \ll 0.16\text{eV}$, so, almost all of the π electrons lie in the S_{00} ground state, and when excited to, for instance, S_{12} , will quickly fall back to the S_{10} state in less than a nanosecond.

After the passage of radiation through the scintillator, there are many molecules which have been excited into the various singlet energy levels (S_{10}, S_{20}, S_{30} — after already having settled from whatever vibrational sublevel into which they were initially excited). In a very short period of time, these molecules will relax into the S_{10} energy level through various non-radiative processes. The majority of the first excited state molecules will then decay back to one of the S_{0x} levels, some through radiative transitions, others non-radiatively. The radiative transition is called *fluorescence*, the intensity of which is governed by $I = I_0 e^{-t/\tau_s}$, where τ_s is the characteristic decay time of the singlet state emission. The absorption and fluorescence processes are shown by the red and blue arrows respectively in Figure 3.20. Of particular importance for the efficacy of the scintillation mechanism is the fact that the maximum photon energy emitted in the fluorescence process (from the $S_{10} \rightarrow S_{00}$ transition) is the minimum energy required for the lowest energy absorption process ($S_{00} \rightarrow S_{10}$ transition), therefore, the scintillator is largely, but not completely, transparent to its




Figure 3.20: Energy level scheme for the π electrons of an aromatic hydrocarbon scintillator like pseudocumene. Direct excitation into the triplet states requires a spin-flip of the electron and is thus negligible. The triplet states are populated through molecular recombination as well as transitions from the singlet states via a process known as intersystem crossing. See the text for further description. Figure from reference [55]

own light.

A few of the molecules in the S_{10} state can undergo a spin-flipping process called inter-system crossing into the first excited triplet state T_{10} . For the molecules in the triplet state, the decay back to the singlet ground state is suppressed due to the spin flipping requirement, thus, this state is very long lived (on the order of seconds or more). When the radiative decay from the triplet states, *phosphorescence*, does occur, it does so with its own decay lifetime τ_T . The more common method for the dissipation of the energy contained in an excited triplet state is when two molecules, each in a triplet state, collide and undergo the reaction $T_{10} + T_{10} \rightarrow S_{1x} + S_{00}$. The net result of this reaction is one molecule in its singlet

ground state, and another in one of the excited singlet states, the latter will then decay via the fluorescence process. However, the light from this process is delayed by the time required for the two triplet molecules to find one another; when emitted, the light follows the fluorescence time decay profile; this process is known as *delayed fluorescence*.

Thus far, we have only considered a unitary system comprised of pseudocumene and have made no mention of the solute PPO. Borexino relies completely on its ability to reconstruct the position of events based on the time profile of the light collected (and to a lesser extent, the amount of light collected). Important factors then are the speed at which the light is produced (related to the lifetime of the S_{10} state, and the degree of overlap of the absorption and emission curves of pure pseudocumene), the light yield of the scintillator (related to the fraction of molecules which emit light radiatively versus non-radiatively), and how well the emission spectrum matches the quantum efficiency of the photomultiplier tubes. The desired solute is one in which the primary method of excitation is through non-radiative energy migration from the solvent — a high level of radiative migration is useless as it defeats the purpose of trying to obtain a faster pulse.

It was found that the addition of PPO at a concentration of 1.5 g/L would simultaneously resolve all of the aforementioned issues. Reference [65] reports the mean lifetime of the S_{10} state in pseudocumene to be 27.2 ns, whereas in PPO the lifetime is only 1.4 ns — comparable with typical PMT jitter times. Furthermore, the same reference cites a scintillation efficiency of only 45% for pure pseudocumene, but 100% for PPO. The issue of matching the solute emission spectrum to the PMT quantum efficiency spectrum is also resolved by the addition of PPO as can be seen in Figure 3.23.

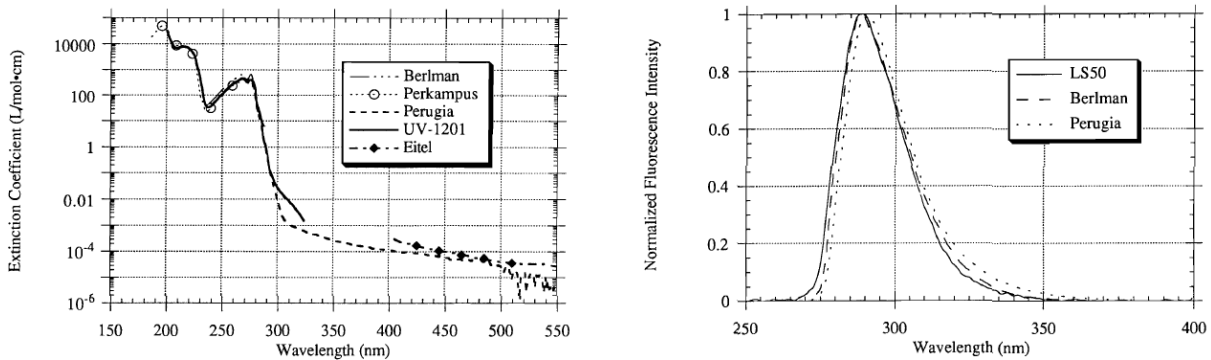


Figure 3.21: Absorption (left) and emission (right) spectra of pure pseudocumene (Figures from [65]). The non-negligible amount of overlap between the absorption and emission spectra means that light emitted by PC will undergo a considerable amount of absorption and re-emission, this results in a smearing of the time spectrum of the scintillation light. Moreover, the quantum efficiency for radiative re-emission of absorbed light is not unity, so, shifting the wavelength away from the absorption spectrum is necessary in order to avoid total extinction of the light.

Characteristics of the light produced

The light emission processes described in the previous section actually have profound consequences in terms of identifying the different types of radiation that cause excitation of the scintillator. The phosphorescence process occurs with a time constant on the order of milliseconds and is irrelevant for the short trigger gates used for Borexino; so, without loss of generality, we can neglect it. The net result of the competition between fluorescence and delayed fluorescence is a light emission spectrum that is comprised of both processes — the former referred to as the “prompt” component, and the latter the “slow” component.

Recall that the triplet states are predominantly populated by molecular-ion — free-electron recombination (triplet states result by approximately a 3:1 ratio to singlet states based purely on statistical considerations) and it is these triplet states that cause the longer lived component of the scintillation light. A particle which produces a higher density of molecules in the triplet states will yield a track along which there are more collisions between molecules




Figure 3.22: Absorption (left) and emission (right) spectra of PPO (the right spectrum is actually the emission spectrum of a PC + PPO mixture). There is markedly less overlap between the two spectra than there was for pure pseudocumene (Figure 3.21), therefore the scintillation light will undergo minimal absorption and re-emission. Close to the point of energy deposition, a few such interactions may take place, the net result being that the wavelength of the light is progressively biased towards higher values (known as a Stokes shift) where the re-absorption is less and less probable. Plots from [65].

in triplet states, and when they recombine and emit light via delayed fluorescence there will be a larger slow component of the scintillation light. Therefore, particles with a large $\frac{dE}{dx}$ such as heavy ions, α 's and to some extent protons, can be distinguished from β 's by looking at the relative abundance of light in the tail to that of the entire pulse — known as the *tail to total ratio* (TTR). Figure 3.24 shows the normalized light output of a sample of Borexino scintillator bombarded with α 's and β 's — the larger slow component of the α component is a very important result as it affords a method of tagging them.

Ultimately, the light output from an organic scintillator can be described by a sum of exponentials:

$$I(t) = \sum_i \frac{q_i}{\tau_i} e^{-\frac{t}{\tau_i}} \quad (3.19)$$

In 3.19, the relative amplitudes, the q_i , are subject to the normalization condition $\sum_i q_i = 1$. Values for the q_i, τ_i are shown in Table 3.4. Equation 3.19 does not take into account




Figure 3.23: Fluorescence emission spectrum of pure pseudocumene (blue), pseudocumene plus 1.5 g/L of PPO (red); and the quantum efficiency of the Borexino photomultiplier tubes[18].

the initial rise in output intensity from the quiescent level (which is exponential with a time constant τ_p), however, the usual approach for fast rise time scintillators is to convolve equation 3.19 with a gaussian (this is in fact what is used for the scintillator decay time probability distribution function in the reconstruction algorithms; see section 3.3.2).

Quenching

There are in fact a multitude of processes by which scintillator molecules may relax to their ground state, only some of which involve the emission of light. Like triplet excitation, non-radiative molecular recombination to the ground state is itself a function of the density of ionized molecules, and therefore occurs much more often for those particles with a large




Figure 3.24: Normalized scintillation emission spectra from α and β excitation of the Borexino scintillator[66]. The higher density of ionized molecules produced by the α 's causes a higher percentage of light emission via the delayed fluorescence process.

value for $\frac{dE}{dx}$. This process, called ionization quenching, is responsible for the reduction of the deposited energy down to the energy one actually observes. For α particles in Borexino, the observed number of photons is approximately 10% of that expected from an equal energy β . The Borexino collaboration has empirically determined the conversion factor to be:

$$Q(E_\alpha) = 20.3 - \frac{E_\alpha}{0.77} \quad E_\alpha^{observed} = \frac{E_\alpha}{Q(E_\alpha)} \quad (3.20)$$

Where the energy in equation 3.20 is given in MeV and refers only to the energy of the α particle (some α decays such as that of ^{212}Bi include a γ).

The situation is more complicated for β particles however, because their ionization density

Particle type	τ_1	τ_2	τ_3	τ_4	q_1	q_2	q_3	q_4
β - γ	6.57	17.61	59.50		0.895	0.063	0.042	
α	3.25	13.49	59.95	279.1	0.630	0.178	0.119	0.073

Table 3.4: Time constants and relative intensities for the light produced when Borexino scintillator is excited by the types of radiation listed. These values were measured by the CTF, and taken from reference [67].

is dependent on their kinetic energy. Theoretically, the light produced as fluorescence, L , by a particle along its track (from 0 to R , its range) is given by:

$$L = \int_0^R S \frac{dE}{dx} dx \quad (3.21)$$

The value S is the absolute scintillation efficiency and is defined as the fraction of the incident energy actually converted to photons. If we define a constant B such that $B \frac{dE}{dx}$ is the density of ionized and excited molecules per unit path length, and define another constant k to be the fraction of these ionized and excited molecules that participate in ionization quenching, then we can express equation 3.21 as:

$$L = \int_0^R \frac{S \frac{dE}{dx}}{1 + kB \frac{dE}{dx}} dx \quad (3.22)$$

This formulation is attributed to Birks[64] and the combined constant kB is known as the Birks parameter¹⁵ — it typically takes on a value on the order of $\sim 0.01 \frac{\text{cm}}{\text{MeV}}$. For particles with a small $\frac{dE}{dx}$ such as high energy β 's, equation 3.22 reduces to:

$$L = \int_0^R \left\{ \frac{S}{kB + \left(\frac{dE}{dx}\right)^{-1}} \right\} dx \approx \int_0^R S \frac{dE}{dx} dx \approx SE \quad (3.23)$$

¹⁵The Birks parameter is a function of the scintillator and must be empirically determined

Thus, in this regime, the result is as expected — the light produced is linearly proportional to the *energy* of the incident particle.

By contrast, for α particles with a rather large $\frac{dE}{dx}$, the denominator in equation 3.22 can be taken to be just $kB\frac{dE}{dx}$ and 3.22 reduces to:

$$L \approx \int_0^R \frac{S}{kB} dx \quad (3.24)$$

In this regime, we have saturation, and the light output varies linearly with the *range* of the particle.

3.2.3 Photomultiplier Tubes

The photomultiplier tubes (PMTs) used for the inner and outer detector of Borexino are the 8" ETL-9351 model developed specially for Borexino by Thorn. The tubes feature a special pseudocumene-resistant sealing and coupling mechanism and a low-radioactivity glass bulb. The quantum efficiency of the PMTs was shown in Figure 3.23 and has a maximum value of 26%. The PMT high voltage and the signal produced by the PMT travel on the same 55 meter long submarine cable through special feedthrus and are capacitively decoupled in the DAQ. Each PMT was completely characterized (dark rate, single photoelectron peak position and width, etc.) before its installation into Borexino[68].

Inner-Detector PMTs

In total, there are 2212 photomultiplier tubes installed on the interior surface of the stainless steel sphere; 1838 of these are equipped with aluminum light concentrators to increase the light collection efficiency (the geometric coverage provided by the PMTs is approximately

33%). The inner-detector PMTs without light concentrators comprise the inner-detector muon veto system¹⁶. The light concentrators are essentially modified Winston cones and are designed such that for any given PMT, a photon emitted anywhere in the inner vessel can be detected by that particular PMT. The acceptance for regions outside of a radius of five meters drops off rapidly[66], thus, not having the light cones on the inner muon PMTs allows them to see the entire inner vessel. Figures 3.25 and 3.26 show the Borexino PMTs from various angles. Figure 3.27 depicts the shape and angular acceptance of the PMT light concentrators.



Figure 3.25: Side-view of a Borexino PMT with light concentrator and μ -metal shielding attached




Figure 3.26: Front view of an 8" ETL-9351 PMT

The precise determination of the position of an event in Borexino is based on the arrival time of scintillation photons at the PMTs — this requires that all of the PMTs be perfectly aligned in time. To that end, each inner-detector PMT is equipped with a fiber optic cable pointing at its photocathode¹⁷. These fibers are collected in several bundles throughout the

¹⁶Strictly speaking, this system is only for detecting muons in the inactive buffer — it is quite easy to detect a muon traversing the active scintillator since it produces a vastly larger amount of light than typical α , β , γ , or ν events

¹⁷These are attached to the edge of the light cones. For PMTs without light cones, a different mount is used, but ultimately puts it at the same distance to the photocathode as the fibers on the light concentrators.




Figure 3.27: The left plot shows the side view of one of the light concentrators — the right plot shows the acceptance of the cones as a function of the incident angle relative to the photocathode normal vector. Figure from reference [66].

detector and eventually couple to a single fiber optic attached to the calibration laser. The calibration laser produces pulses 50 ps in width at a wavelength of 394 nm, and is fired at 100 Hz at the beginning of every run, as well as during much longer laser calibration runs which occur at least once a week.

Outer-Detector PMTs

While the 208 PMTs on the external surface of the SSS and inside of the WT are the same ETL-9351 PMTs that were used in the inner detector, they differ in that they are completely encapsulated as shown in Figure 3.28. The encapsulation contains the entire PMT, voltage divider, and μ -metal shielding and a mineral oil filled cap placed over the photocathode.

The outer detector has its own independent calibration system driven by laser-LEDs coupled to fiber optic cables. A more complete description of this system can be found in reference [55].




Figure 3.28: Schematic of the encapsulation design for the muon veto PMTs in the water tank[66].

3.2.4 Purification Plants

The reduction in the level of intrinsic background radiation down to that required to perform an efficient measurement of solar neutrinos necessitates a very intense purification effort. An in-depth discussion of the Borexino purification system can be found in [69], here we recap just the main points. A schematic of the overall fluid handling and filling system can be seen in Figure 3.29

Pseudocumene Purification

After being delivered in specially designed tanker trucks, the pseudocumene is pumped into one of four storage tanks (each holding approximately 100 tons)¹⁸. From here, the pseudocumene first undergoes vacuum distillation to remove compounds which adversely affect the light yield and attenuation length¹⁹.

¹⁸In actuality, only three of these tanks are used for storing scintillator, the fourth is left as a buffer volume to facilitate fluid movement operations.

¹⁹Dimethylbenzaldehyde is a compound which pseudocumene forms when exposed to oxygen and is a very strong quenching agent and is predominantly responsible for the characteristic “yellowing” of pseudocumene exposed to air.

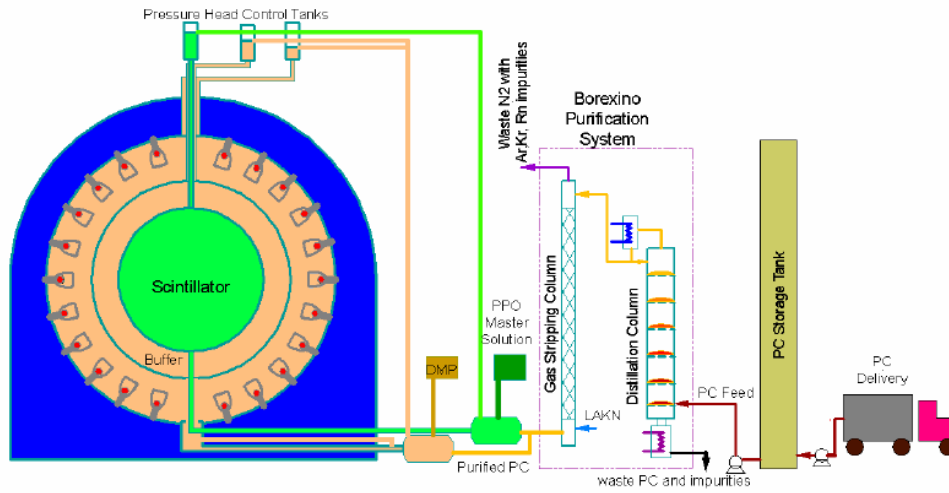


Figure 3.29: Overall view of the purification and filling procedure for Borexino[69]. After being delivered in special tanker trucks, the PC is unloaded into storage tanks and then undergoes distillation, nitrogen stripping, and then mixing with either a flash-distilled concentrated master solution with PPO, or a master solution of the quencher DMP.

Following distillation, it is possible that the pseudocumene still contains some dissolved gases such as radon and krypton, both troublesome for Borexino, therefore, the scintillator is sent through a nitrogen stripping stage to remove these contaminants. The pseudocumene is introduced at the top of a very tall packed column and allowed to flow downwards against a stream of special Low Argon and Krypton Nitrogen (LAKN) which is introduced at the bottom of the column. After leaving the nitrogen stripping column, the now purified PC is sent to a storage tank to await its mixing with either PPO or DMP. Due to its very high melting point (70°C), PPO cannot be purified directly without extensive measures being taken to avoid solidification. The method adopted is to prepare a highly concentrated “master solution” of PC with 140 g/L of PPO which can be purified by flash distillation. This solution is then metered into the pure PC when filling the inner vessel.

To guard against possible re-contamination of the scintillator during and after the purification process, all critical valves, fittings, and pumps are outfitted with sealed, nitrogen-blanketed, secondary enclosures. During the commissioning phases, all components were individually

He leak checked to provide further assurance against air leaks.

Water Plant and Cleaning Module

The purification and filling operations of Borexino require large amounts of ultrapure water (≈ 2000 tons for filling the water tank; and large volumes required for rinsing components after cleaning) and nitrogen (nitrogen blanketing of pumps, fittings, valves, and vessels; nitrogen stripping; drying). Water from the underground labs is sent through the water plant where it and passes through filters, and then undergoes nitrogen-stripping, deionization, and reverse osmosis purification stages. The resulting water has a resistance of $18 \text{ M}\Omega$ and a reduction by approximately four orders-of-magnitude in the activity of uranium and its daughters. The plant is currently used in a loop mode, drawing water from the bottom of the water tank, heating it and reintroducing it to the top of the water tank to maintain a constant temperature gradient inside of the water tank, and therefore the stainless steel sphere.

The commissioning of Borexino and its associated plants required extensive cleaning of all surfaces, pipes, and vessels. To facilitate this process, a dedicated cleaning module was constructed with the ability to circulate hot detergent, chelating agents, rinse water, and subsequent nitrogen drying. The usual cleaning procedure consists of several hours of circulating with a critical cleaning detergent (most commonly: Micro90²⁰ or Detergent 8²¹), then rinsing with ultra high purity water until resistivity measurements ceased to indicate the presence of residual detergent. In an effort to remove surface metal impurities (lead, polonium, etc.), a solution of EDTA is then circulated for several hours, followed by a water rinse until particulate counting indicates the desired cleanliness level²², and of course a final

²⁰Produced by Cole Parmer

²¹Produced by Alconox

²²For most components this was set to be Class 30 as defined by Mil-Std-1246C

drying with LAKN. This module was used extensively for the cleaning of many components of the insertion system as will be described in Chapter 6.

3.3 Data Acquisition and Software

3.3.1 DAQ

Inner-Detector DAQ

Rather miraculously, all of the physics performed with the Borexino detector can be accomplished by simply recording the precise time of arrival of photons at each PMT and the charge collected at the anode of each PMT. To facilitate this task, there are 14 identical electronics racks with dedicated electronics for extracting this information from all of the inner-detector PMTs. Each rack contains the following items (visible in Figure 3.30):

- Patch Panel — The 55m submarine cables attached to the PMTs arrive at the top of each electronics rack where they are then routed to the front-end board to which it belongs. The patch panel simply functions as a common termination point for all 158 of the submarine cables processed by one rack.
- Front-End Boards (FEB) — The fourteen front-end boards comprise the first half of the pulse processing for the inner-detector electronics. The input to a given channel on a front-end board contains both the DC high voltage for that channel, as well as the AC signal riding on top of it; the first stage of the front-end board decouples these two components. Each PMT signal is split into two copies, the first of which is amplified and shaped to be used in a subsequent stage for the timing pickoff, and the other copy is integrated in preparation for digitization. In addition to these primary tasks, the

front-end board also produces an analog sum of all twelve channels on the board — this is sent to the scaler module on each rack.

- Laben Boards — These are VME digitizer boards produced by the now defunct Laben S.p.A. electronics company in Milan, Italy. There are twenty of these boards in each VME crate, with each board capable of processing eight channels. These boards accept the timing and integrated charge outputs produced by the front-end boards and subsequently perform the time pick-off to an accuracy of 200 ps [70]. For trigger evaluation, each board contains a coincidence unit with a coincidence time of 50 ns and sends its output to the trigger processing logic. Each channel also contains its own memory registers where the data are stored for a duration of $6\mu\text{s}$. If a trigger is received, these data are copied to on-board RAM in preparation for event-building. Each VME crate possesses one PowerPC module which has the task of retrieving the data from all of the channels in all of the boards in its crate, building it into part of an event, and then sending the event over an internal network to the main event-building computer BxBuild.
- Scaler / Fast Adder — The analog sum output of all the channels in each FE board are sent to one scaler / adder module in each crate. This module performs an analog sum of all of the board-level sums, in other words, the analog sum of all PMTs in that particular crate. The crate analog sum is then used in the flash ADC based supernova monitoring system, the Princeton neutron DAQ, as well as the Genova neutron DAQ. The module also has an internal discriminator and scaler for each FE board connected to it — these are routed to a web interface which is very useful for detector monitoring and diagnostic purposes. Since the information provided to it is just the analog sum of the entire board, it is impossible to determine precisely which channel might be causing a problem, only which group of twelve channels contains a problematic PMT

or channel.

- High Voltage — Finally, each crate houses a high voltage unit for all of the PMTs in its crate. Each HV crate holds seven modules, each module capable of powering two front-end boards. The entire high voltage system is interlocked and can only be operated when a key is inserted into a special module in the counting room — in order to operate the high-powered lights used to illuminate the SSS when taking a photo this key must be removed and placed in a similar box for the camera system.



Figure 3.30: A few of the fourteen inner detector electronics racks. From top to bottom: PMT cables entering from the SSS, patch panel, FE power supply, adder/scaler module, FE crate and boards, Laben crate and boards, HV crate. See text for functional description.

Outer-Detector DAQ

The inner-detector electronics are designed primarily for the detection of rather small amounts of light ($\sim 0.2 - 2$ photoelectrons for events with energies in the neutrino window) and as such, are not well suited for the detection of the large amount of light produced by muons. Therefore, a separate electronics chain is used to cope with the much higher dynamic range required for this task. The basic process is the same since the end results are identical. In this case however, the PMTs are decoupled in special decoupling units, and then sent directly to charge-to-time converters (QTC) which are then forwarded to TDCs where they are integrated to determine the total charge seen by that channel (the width of the output from the QTC), and the time of the hit (start of the pulse). There is also a dedicated trigger system for the outer-detector which provides a muon veto signal to the main Borexino trigger. A more in-depth discussion of the outer-detector electronics can be found in [55].

Trigger Formation

During normal data taking, the Borexino Trigger Board (BTB) is constantly fed with the number of channels that have crossed preset detection thresholds in both the inner and outer detector. For the inner-detector, the coincidence counting begins at board level and is then fed to special coincidence adders attached to each crate, these adders then forward the information to another set of adders in the trigger crate; the result is then analyzed by the BTB. Inner-detector trigger formation occurs when 25 channels (corresponding to approximately 50 keV) register hits in a time window of 60 ns. This threshold is configurable at DAQ run time and is sometimes varied to reduce the trigger rate (i.e. during a calibration with a high activity source). The outer-detector produces its muon trigger when six channels fire within a time window of 150 ns. The trigger board can also produce triggers when: one

of the calibration lasers is fired, a calibration signal is fed to all of the front-end boards in lieu of a PMT signal, and at random intervals so as to measure the PMT dark noise.

Regardless of the source of the trigger, when it occurs, all of the data in both the inner and outer detector is recorded and sent to disk. The window during which data is recorded lasts a total of $16\ \mu\text{s}$, appropriately delayed such that the recording of the data starts before the hits that actually caused the trigger. There are several sets of radioactive isotopes which decay in rapid succession, quickly enough that the two events will often occur inside of the same $16\ \mu\text{s}$ trigger window (^{212}Bi - ^{212}Po , part of the ^{232}Th decay chain $\tau = 431\ \text{ns}$ and ^{85}Kr - $^{85\text{m}}\text{Rb}$ $\tau = 1.46\ \mu\text{s}$). In order to tag these isotopes, they will have to be identified as two separate events — thus there is a software reconstruction stage focused on deconvolving the entire trigger window into as many physical events as appropriate.

After Borexino began data taking in 2007, it was discovered that the inner-detector electronics can become saturated (and their charge information unreliable) after a muon traverses the inner-vessel. The reason for this is that immediately following a muon, there are a large number of channels that have seen so much light that their ADCs have been saturated; the result being that these channels report a value of zero for their measured charge, thus biasing the event energy towards lower values. To fix this problem, the firmware was modified such that if both the inner- and outer-detector triggers occur at the same time, the normal trigger is issued and immediately after it expires, a new $1.6\ \text{ms}$ trigger gate is opened — this has the effect of clearing the ADCs and allowing the multitude of neutrons and cosmogenically produced isotopes in the wake of the muon to be detected.

3.3.2 Software Reconstruction

The task of transforming the time and charge information obtained from each channel into the variables used for analysis is performed by the software reconstruction programs. Prior to 2005, the analysis code was in a state of disrepair, with many components written in different languages and otherwise largely inefficient. The collaboration then decided to revamp the code completely — the resulting program, *Echidna*, is the code written and used by the European collaborators. In 2007, fueled by some short-comings of the *Echidna* code and the desire to have a second independent analysis code, the American collaborators from Princeton and Virginia Tech took on the task of writing a separate analysis code. After a few weeks, the code, named *Mach4*, was in a sufficiently mature state to perform an analysis on just six days of Borexino data and extracted a β spectrum showing features very reminiscent of the ${}^7\text{Be}$ Compton edge.

The *Mach4* code (documentation at [71]) was designed from the beginning to have an emphasis on speed and modularity. Written predominantly by a group of graduate students at Princeton University, *Mach4* is written entirely in C++ and produces as its output a flat ROOT[72] ntuple for each run. The code is open to all collaboration members via CVS access, and the data are stored on RAID disk arrays in Italy, Princeton, Virginia Tech, and UMass Amherst. The code follows a typical development cycle of modifications, testing, freezing, and culminates with a batch reprocessing of all of the data with the new code cycle. During all but the batch reprocessing stages, a `cron` script runs every night and starts jobs that download the raw data and reconstruct it with the current release version of *Mach4* and then copies those data to all of the disk arrays in remote locations. When a new release of the code is made, the mass reprocessing is performed on two computer clusters at Princeton and Virginia Tech — at the time of this writing, the most recent mass reprocessing involved over 6000 thousand runs and took less than a week to process.

Precalibration

The first task performed by either code is that of *precalibration*²³ — a special set of routines runs over the raw data file and calculates:

- Timing offset values are calculated to apply on a channel-by-channel basis to equalize the timing of all channels. This information is obtained from the pulser triggers at the beginning of each run (these are the “fake” PMT signals fed to the calibration inputs of the FE boards), and the calibration laser triggers at the start of each run. The pulser events are used to align the timing of the FE and Laben boards, and the laser events are used to account for any cable length differences between the FE boards and the PMTs themselves.
- Charge calibration is performed with the pulses from the calibration laser fibers pointed at each PMT. In particular, the charge histogram for each channel is fit to a function and the position of the single photoelectron peak is recorded and used later in calculating the number of photoelectrons seen by that channel in an event.
- Calculation of the dark rate for each channel is performed with the dedicated random triggers. These triggers occur both at the beginning of each run and also throughout the entire run. This dark rate is used during calculation of the number of photoelectrons or hits occurring during a cluster: the number of expected dark hits for each channel is computed and subtracted to yield the “true” number of hits.

During this process of precalibration, individual channels may be marked as dead for not possessing any events in laser and/or random triggers, having too high of a dark rate, having an abnormally wide or offset single photoelectron response, or requiring a timing correction

²³The majority of this portion of the code was written and maintained by B. Loer of Princeton University.

too far out of specifications. Despite careful installation procedures, there are some channels whose calibration fiber is broken or damaged — in the former case, little can be done to extract energy calibration data for this channel and it is simply not used for calculating the energy of an event. A distinction is made between broken and damaged fibers; the latter typically do channel some light, but the laser statistics in a normal run are insufficient to allow the determination of its energy calibration variables. Therefore, at least once per week, a special laser calibration run is performed where the laser is fired at 100 Hz for a total of 250,000 triggers — this is usually sufficient to extract the energy calibration data with high accuracy from all channels, including those with damaged fibers.

Particularly during the first few months of data taking, it was not uncommon for the DAQ to disable, at any point during a run, all data coming from a particular Laben crate due to communication or synchronization issues. Additionally, there are cases where a PMT starts to fail or exhibit problematic behavior and can cause a trip of a particular HV board. Since the energy variables used in the analyses are calibrated to 2000 PMTs by scaling to the number of live PMTs, significant departures in the number of live channels during a run will bias the energy reported after the failure to a lower value. To guard against this, both analysis codes monitor the number of channels responding to pulser triggers throughout the run²⁴. The pulser signal is injected in place of the PMT signal, thus, any failures in the electronics chain will cause that channel not to report the fake hit issued by the pulser. If at any point during a run, the number of channels responding to the pulser triggers falls below 80% of the mean number that responded to those issued prior to the start of the run, that event number is saved as the variable `Last_Good_Event` which is used later in the analysis scripts to skip over any portion of a run after a serious electronics failure. The *Echidna* code actually adjusts their scaling of the energy variables to 2000 PMTs on an

²⁴Recall that during a normal run, service triggers are issued on the order of once every second. These triggers are either calibration laser triggers, random triggers, or pulser triggers.

event-by-event basis; the *Mach4* group may implement this feature in a future release.

Clustering

The $16\mu\text{s}$ trigger gate is long enough that it can often contain two or more events which occur in rapid succession (^{212}Bi - ^{212}Po , ^{85}Kr - $^{85\text{m}}\text{Rb}$, neutrons after muons, etc.) and being able to look at these events individually requires that the hits be split, or *clustered*, into physical events. This task, performed in the *Mach4* clustering module²⁵, follows the following algorithm[73]:

- The hits in the trigger window are loaded into memory and are checked for the validity of their hit times, invalid times are discarded. As another quality check, the trigger is discarded if there were not at least three crates with four or more hits.
- The times of the hits are then put into a histogram with a bin-width of 16 ns. The algorithm then searches for the first three bins that each have at least nine hits in them; the first hit in this set of three bins is defined to be the start of the cluster.
- The clustering module then defines the *temporary* end of the cluster by searching for the first six bins that have less than three standard deviations above the dark noise level; the last hit in these bins is set to be the temporary end of the cluster.
- If the total number of hits between the start of the cluster and the temporary end of the cluster is less than twenty, the cluster is discarded and not written to disk. If there are more than twenty hits, the cluster is shelved temporarily and the next part of the trigger window is searched for a cluster immediately following the one just processed. If there is another cluster, the temporary end of the cluster is defined to be the real end of the cluster and it is then stored.

²⁵This module was written and is maintained by R. Saldanha of Princeton University.

- If another cluster was not found immediately following this one, then an additional tail is added to the end of the cluster. The length of this tail is a function of the number of individual hits inside of the cluster. If the cluster contains 100 or less hits, the tail is extended by another 512 ns (32 bins); if the number of hits is greater than 100, then the tail is extended by $32 \times \left(1 + \log \left\lceil \frac{\text{hits}}{100} \right\rceil\right)$ bins (each bin is 16 ns long). The last hit in this tail is set to be the end of the cluster.

A modification of this routine is used for the 1.6 ms triggers that occur immediately following a muon, for details of this method see the *Mach4* documentation, or reference [73].

Once the clusters have been defined, it is possible to compute the energy of the event described by the cluster. There are two energy variables used in Borexino, one is simply the number of PMTs that registered valid hits during the entire cluster, and the other is the sum of the ADC values from all of the PMTs that were hit. We have already established that from run-to-run, the number of live channels varies and therefore the energy reported for a given event will vary as well. To make the analysis effort easier, we normalize these values to 2000 live channels. In the ROOT ntuples produced by *Mach4*, the raw values (unnormalized) are stored as `npe` and `nhits` respectively. The former is calculated simply as $\sum_0^{2218} \text{npe}_i$, where npe_i is the value of the `npe` in each channel i .

The normalized version of `npe` (often referred to as “charge”) is called `npe_corrected` and is calculated as $\text{npe_corrected} = \text{npe} \frac{2000}{N_{PMT}}$.

The problem with `nhits` is that it cannot account for the fact that higher energy events are more likely to have PMTs that are hit more than once²⁶, thus the energy reported by `nhits` is biased towards lower values as the deposited energy increases. The probability to obtain two or more hits in a photomultiplier tube is given by summing over the Poisson distribution

²⁶In other words, you can only tell if a PMT has been hit, not how many times — the latter must be determined from the charge

for $n \geq 2$, and reaches the level of 10% for an energy of 2.3 MeV (assuming a 500 pe/MeV light yield).

It is possible to use the information obtained from the ADCs to calculate what the number of hits *should* have been if it were possible to correct for multiple occupancy — this is in fact done (at least to first order) in Mach4 and is referred to as `laben.cluster.nhits_corrected`²⁷.

If we make the simplifying assumption that an event occurs at the origin (therefore, all PMTs have the same solid angle coverage of the event), and that all of the PMTs have the same quantum efficiency (we also assume that they all have a detection threshold of zero photoelectrons), then, the probability for any given PMT to be hit by a photon is just $1/N_{PMTs}$. Since we are just looking for the number of PMTs that are hit *at least once* this is given by:

$$N_{hits} = \left(1 - \left\{1 - \frac{1}{N_{PMTs}}\right\}^{npe}\right) N_{PMTs} \quad \Rightarrow \quad 1 - \frac{N_{hits}}{N_{PMTs}} = \left(1 - \frac{1}{N_{PMTs}}\right)^{npe} \quad (3.25)$$

Taking the natural logarithm of both sides, and then normalizing to 2000 PMTs, gives us the `nhits_corrected` variable, or, “npe from nhits”:

$$\text{nhits_corrected} = \frac{\ln\left(1 - \frac{N_{hits}}{N_{PMTs}}\right)}{\ln\left(1 - \frac{1}{N_{PMTs}}\right)} \cdot \frac{2000}{N_{PMTs}} \quad (3.26)$$

The behavior of equation 3.26 as a function of the number of hit PMTs can be seen in Figure 3.31.

²⁷In addition to the multiple occupancy correction, it is also normalized to 2000 PMTs.




Figure 3.31: The dependence of nhits corrected — whose functional form is shown in equation 3.26 — on the number of hit PMTs. To illustrate the departure from a linear relationship, a straight line with unit slope is drawn.

Position Reconstruction

The previous sections have dealt with how to turn the charge information from the individual channels into physical variables; now the timing information can be used to derive an estimate of the spatial coordinates of the event. The sequence of events that ultimately lead to detection of a photoelectron are the following:

- A particle deposits energy in the scintillator thus exciting it, this occurs at global time T_{event}
- The scintillator takes a finite amount of time for the de-excitation of the pseudocumene, energy transfer to PPO, and emission by the PPO. A given photon is emitted by the scintillator after an interval T_{PDF} — which is distributed according to the probability distribution function (PDF) for the scintillator — modeled in *Mach4* as a Gaussian

convolved with two exponentials.

- The photons then take a time T_{TOF} to travel from their emission point to the individual phototubes. Ideally, this is just given by $n \frac{d_i}{c}$ where d_i is the distance to the i -th PMT. However, in a detector as large as Borexino, effects such as absorption and re-emission, Rayleigh scattering, etc., all act to delay the light. These effects can be partially accounted for by increasing the index of refraction beyond its true value as a way to “slow down” the light.
- There is a small additional amount of time, T_{jitter} , required for the photoelectron emitted from the cathode to reach the first dynode. The path that the first photoelectron must take from photocathode to the first dynode depends upon where on the photocathode it was produced — this uncertainty is known as the PMT transit time jitter, and is on the order of 1 ns for the Borexino PMTs[74]. The jitter time simply acts to broaden the distribution of hit times slightly.

Assuming that all of the PMTs and electronics are equalized in time, there are no further corrections to the arrival time required, and the hit times are given by:

$$T_{hit_i} = T_{event_i} + T_{PDF_i} + T_{TOF_{i,j}} + T_{jitter_{i,j}} \quad (3.27)$$

where i runs over all the hits in a cluster, and j corresponds to a given PMT since they all will have different distances from the event position, and in principle can have different jitter characteristics. We can make a simplification by absorbing the jitter time into the emission time — it is thus accounted for in the gaussian component of the PDF. If we let \vec{X}_0 be the (x, y, z) coordinates of the event (for notational consistency, we will also switch to $T_{event} = T_0$), and \vec{X}_{PMT_i} be the coordinates of the i -th PMT, then we can express equation

3.27 as:

$$T_{PDF} = T_{hit_i}(\vec{X}_{PMT_j}) - T_{TOF}(\vec{X}_0, \vec{X}_{PMT_j}) - T_0 \quad (3.28)$$

It is easy to see that for the actual event position, the time-of-flight-subtracted hit times ($T_{hit} - T_{TOF}$) equal the scintillator PDF. Since they are not known a priori, the goal then is to find the position and time of the event (\vec{X}_0, T_0) that minimizes the difference between the time-of-flight-subtracted hit times and the scintillator PDF. This is the task of the position reconstruction module in *Mach4* and it does so by performing a likelihood fit of the hit times to the scintillator PDF.

For a likelihood-based fit, the fitter is given the list of PMTs that were hit, their locations, and the times at which they were hit — the fitter then guesses a location and time for the event, and computes the right half of equation 3.28. The fitter then iterates through various guess positions trying to minimize the value of the reduced times with that of the log of the scintillator PDF. The formula for the latter in *Mach4* cycle 9.5 is:

$$\log(T_{PDF}) = \log \left(\sum_{i=0}^1 \frac{q_i}{2\tau_i} e^{\left(\frac{\sigma^2}{2\tau_i^2} - \frac{t}{\tau_i}\right)} \text{Erfc} \left\{ \frac{\sigma}{\sqrt{2}\tau_i} - \frac{t}{\sqrt{2}\sigma} \right\} \right) \quad (3.29)$$

Where in equation 3.29 the q_i are the relative strengths of the two exponentials with mean lives τ_i , σ is the width of the Gaussian component, and Erfc is the complementary error function given by:

$$\text{Erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-t^2} dt \quad (3.30)$$

However, this is only true when PMTs only see one photoelectron. The arrival time of a photon in a multi-hit event is biased towards earlier times, therefore, the PDF must be modified by the following formula[57]:

$$T_{PDF}^{n_\gamma} = n_\gamma \cdot T_{PDF}^{n_1} \cdot \left(1 - \int_{-\infty}^t T_{PDF}^{n_1}(t') dt'\right)^{n_\gamma-1} \quad (3.31)$$

The second term in the brackets is just the cumulative distribution function of the PDF. As it turns out, this modification of the PDF can be accommodated by modifying the parameters q_i, τ_i , and σ for the different numbers of photoelectrons seen by a PMT. The *Mach4* code treats the three distinct cases of $n_\gamma = (1, 2, > 3)$ separately using the parameters in Table 3.5.

n_γ	q_0	q_1	τ_0	τ_1	σ
1	0.5509	0.4491	0.7415	13.8752	3.6679
2	0.2066	0.7934	6.8285	0.6126	3.635
≥ 3	0.9006	0.0994	0.4979	4.3609	3.631

Table 3.5: Values for the scintillator PDF as a function of the number of photoelectrons seen by a PMT.

This reconstruction is currently performed twice in Mach4 — once with an index of refraction of 1.70, and another with an index of refraction that varies with the energy of the event. The former was found to be the index of refraction that reconstructed the edge of the inner vessel the best. The latter was developed by J. Xu of Princeton University based on calibration data obtained with a ^{222}Rn source and is given by:

$$n = \begin{cases} 1.696 + 6.775 \times 10^{-6} \text{npe} & \text{npe} \leq 1000 \\ 0.083 * \log(\text{npe} + 655.721) + 1.146 & \text{npe} > 1000 \end{cases}$$

Pulse Shape Discrimination

The efficiency for tagging various types of backgrounds that are present in Borexino is greatly improved when one can use particle identification techniques. Such methods are facilitated by exploiting the different time profiles of the scintillation light produced by variations in the ionization density of different types of radiation — see Section 3.2.2, and Figure 3.24. There are two such methods employed by Borexino: the tail-to-total ratio (TTR), and the Gatti parameter [75]. In both cases, these statistical tests are based on the differences in the shapes of the light output, they are performed on the *emission* time spectrum — not the hit time spectrum. Once the position reconstruction module has determined the location and time of the event, it produces the emission time spectrum by subtracting the time of flight of each photon from the event position to each PMT that was hit.

The separation power of any α/β separation technique can be quantified by the “figure of merit” defined as [18]:

$$D_{\alpha\beta} = \frac{\Delta\mu}{\sqrt{\sigma_\alpha^2 + \sigma_\beta^2}} \quad (3.32)$$

This is computed by producing histograms of the alpha-beta discrimination variable for alpha and beta events — each of these histograms will, to a good approximation, be Gaussian. Then, in equation 3.32, $\Delta\mu$ is the distance between the means of these two Gaussians, and σ_α and σ_β are the respective widths; clearly, a “good” alpha beta separation variable produces two very narrow and well separated peaks.

The tail-to-total ratio is given simply as the relative fraction of hits which occur in the tail of the pulse ($\sum_{i=tail}^n f_i / \sum_i^n f_i$), where the tail start time is a user-defined variable. In *Mach4*, there exist quite a number of these variables with tail start times in the range 30–110 ns, in

5 ns intervals. It has been found experimentally that the tail offset time of 70 ns (TTR70) gives the best α/β discrimination ability (as measured by the aforementioned figure of merit) in the *Mach4* software.

The second method used is the Gatti linear discriminant. This statistic is aimed at determining how similar the pulse under consideration is to one of two reference shapes. The reference shapes used by *Mach4* are currently obtained from ^{214}Bi - ^{214}Po events identified in a dedicated processing of the data before a mass processing occurs. For a given normalized scintillation pulse, the algorithm loops through the bins and computes the Gatti parameter defined by:

$$G = \sum_i P_i f_i = \sum_i \left(\frac{\alpha_i - \beta_i}{\alpha_i + \beta_i} \right) f_i \quad (3.33)$$

In equation 3.33, the α_i and β_i are the bin contents of the normalized α and β reference shapes respectively, and f_i is the content of the normalized emission spectrum for the event under consideration. It is clear then that, α 's are defined by $G > 0$ and β 's as $G < 0$. Figure 3.32 shows the α and β Gatti and TTR70 distributions for events obtained during a calibration with a ^{222}Rn source vial at the center of Borexino in June of 2009.

In addition to the tail-to-total ratios and Gatti variables, the α/β module in *Mach4* computes several other variables that can also be used for pulse shape discrimination purposes:

- Skewness: The third standard moment of the distribution. It is a measure of the asymmetry of a probability distribution, or, in this case, a scintillation emission spectrum. In the CTF data analysis, the skewness of pulses was often used to perform α/β discrimination.
- Kurtosis : The fourth standard moment of the distribution. The kurtosis is a measure




Figure 3.32: $\alpha - \beta$ discrimination variable distribution for events tagged as either ^{214}Bi or ^{214}Po during a calibration with a ^{222}Rn source at the center of Borexino. The plot on the left shows the better separation (figure of merit $D_{gatti} = 4.64$, $D_{TTR70} = 4.0$) obtained between α and β events with the Gatti parameter versus the tail to total ratio plot on the right.

of how “peaked” the distribution is — alpha events have larger values for kurtosis as they are characterized by a longer, more pronounced, tail than betas.

- Gatti parameters using reference pulse shapes from the skewness, kurtosis, and tail-to-total ratio methods of discriminating alpha and beta events.

Other Modules

In addition to the core modules described above, the *Mach4* code also contains specialized modules for handling the inner and outer detector muon veto systems, creating the reference shapes for the Gatti parameter determination, reconstructing the tracks of muons as they pass through the inner and outer detectors, calculating parameters which describe the distribution of hits in the detector, and for producing scintillation PDFs from calibration sources. The muon tracking module uses the output of the inner and outer detectors and, based on the hit patterns of the inner and outer PMTs, produces an estimate of the muon’s path through the detector. The ability to reconstruct the track of a muon is of vital importance

for removing the cosmogenically-produced ^{11}C which obscures a very important region for pep and CNO neutrino analyses. The PMT hit distribution information turns out to be quite useful in filtering out hits due to electronic noise (which comprise less than 1% of all events, but the fraction varies considerably from run-to-run) — which appears as an isolated distribution of hits as opposed to the more uniform distribution indicative of an isotropic scintillation event.

Finally, the “position calibration” module is used in conjunction with the deployment of calibration sources and ultimately produces histograms of the emission time spectra seen by PMTs which lie at a specified distance from the event position. If one makes the assumption that all the light propagation effects like absorptions and re-emission, Rayleigh scattering, etc., can be parameterized solely by the distance from an event to a PMT, then, one can generate a PDF whose parameters are a function of that distance. To that end, the module takes as its input the position of the calibration source (be it radioactive or laser) and groups all of the PMTs by the distance d from the source to the i -th PMT in 0.5m bins in d . The resulting output is a series of emission time spectra which can then be fit with the PDF — in the case of the current release of *Mach4*, this is two exponentials convolved with a Gaussian. Each of these histograms is fit with this PDF and the parameters of the fit $q_1, q_2, \tau_1, \tau_2, \sigma, t_0$ are recorded — at the end of this procedure, one has these values for each distance bin. Then a fit of each parameter (as a function of distance d) is made with a second order polynomial — the net effect is that we now have a PDF whose parameters (time constants, relative magnitudes of the exponentials, and width and mean of the Gaussian) depend on the distance. At the time of this writing, the data from this module are still being evaluated and the PDFs so generated have not been put into final production use. However, the work is ongoing and it is ultimately the goal of the *Mach4* group to use a PDF derived from these calibration data. Some results from a test of a PDF generated from the calibration data are

presented in Section 8.3.3.

A set of `bash` shell scripts distributed with the *Mach4* code are used for performing mass re-processing of the data (on individual computers as well as clusters with the PBS scheduler), automatically analyzing each run to determine if it is a valid run to be used for analysis, and several other tools for automating daily downloads and processing of new data.

3.4 Backgrounds

The scientific campaign of Borexino is centered around measuring very rare event rates, often with signals in areas normally dominated by natural background radiation. Borexino is aiming for an ultimate accuracy at the level of 3% for its ${}^7\text{Be}$ ν measurement. The most recent reported measurement for Borexino indicates a signal to background ratio of 1 in the neutrino window (largely due to ${}^{85}\text{Kr}$ and ${}^{210}\text{Bi}$)[20]. From basic Poisson statistics, the statistical error on the neutrino signal ($N_\nu = N_{tot} - N_{bkg}$) is given by

$$\begin{aligned} \frac{\delta N_\nu}{N_\nu} &= \frac{1}{N_\nu} \sqrt{N_{tot} + N_{bkg}} \\ &= \frac{1}{N_\nu} \sqrt{N_\nu + 2N_{bkg}} \\ &= \frac{1}{N_\nu} \sqrt{N_\nu} \sqrt{1 + 2\frac{N_{bkg}}{N_\nu}} \end{aligned} \tag{3.34}$$

$$\frac{\delta N_\nu}{N_\nu} = \frac{1}{\sqrt{N_\nu}} \sqrt{1 + \frac{2}{S/B}} \tag{3.35}$$

The first factor is just the standard Poisson statistical error, and the second factor takes into account the presence of the background (where S/B is the signal to background ratio). For the signal to background ratio of 1 obtained in the last neutrino analysis, the statistical

error on the measured neutrino signal is worse than the ideal (i.e. no background) case by 73% due to the large number of background counts. Obviously, the presence of systematic errors (at the level of 8% in [20]) added in quadrature with the statistical error makes the reduction of backgrounds of paramount importance for achieving the 3% error measurement.

Table 3.6 shows the major sources of background present in Borexino, the approximate natural activity level encountered in Hall C of the underground labs, the contaminant level which is tolerable by Borexino, the method by which the contaminant is removed and the level actually obtained by Borexino. One can see that in several categories, the Borexino purification system has performed remarkably and exceeded the requirement by an order of magnitude or more. The upper limit on the level of ^{85}Kr is higher, by a factor of about three, than what was originally planned; this is likely due to an air leak that occurred during the filling (further discussed in Section 3.4.2). The reduction of the ^{85}Kr is the goal of the first purification campaign since the initial filling, which began in the Spring of 2010; being a noble gas, krypton, can be removed by nitrogen stripping with special Low Argon and Krypton Nitrogen.

3.4.1 Decay Chains

Many of the radioactive isotopes which contribute to the background in Borexino are parts of the decay chains of either ^{238}U or ^{232}Th . These heavy-element decay chains each produce ten or more isotopes with different mobilities, energies, and particle types. The source of these isotopes is the stainless steel of the sphere, the glass and inner components of the PMTs, the nylon of the inner vessels, etc.

Contaminant	Source	Normal Conc.	Borex Req.	Reduction Method	Borex Achieved
μ	Cosmic	$200/(s \cdot m^2)$	$10^{-10}/(s \cdot m^2)$	Location, OD veto, pulse shape	$< 10^{-10}$
^{14}C	PC/PPO	$10^{-12}g/g$	$10^{-18}g/g$	Old oil, check in CTF, threshold cut	$10^{-18}g/g$
^{238}U	Dust	$10^{-4}g/g$	$< 10^{-16}g/g$	Purification	$< 10^{-17}$
^{232}Th	Dust	$10^{-4}g/g$	$< 10^{-16}g/g$	Purification	$< 10^{-17}$
^{210}Po	PPO, Rn			Purification, α/β	$0.01c/d/t$
^{222}Rn	Air, emanation	$10 - 100Bq/L$	$< 1c/d/100t$	Stripping, delayed coinc	$< 0.02c/d/t$
^{85}Kr	N_2	$1Bq/m^3$	$< 0.01ppt$	Leak tightness	$< 0.035ppt$
^{208}Tl	^{232}Th	$10^{-4}g/g$		Purification, stat. sub.	$\sim 28c/y/100t$

Table 3.6: Sources of background, typical level of the background, maximum tolerable signal level for Borexino, methods by which they are reduced, and what was ultimately achieved in Borexino[76].

The ^{238}U Decay Chain

The ^{238}U decay chain is a set of 15 decays which transform ^{238}U to ^{206}Pb — several different branches are possible along the way as can be seen in Figure 3.33.

The inner and outer buffers shield the inner-vessel from the vast majority of the radiation originating from uranium which may be embedded in the construction materials of the detector, and thus, to a good approximation, the signals seen from this chain originate in the inner-vessel. The purification skids, most notably the particulate filters, reduce the number of suspended heavy elements in the inner-vessel to a practically negligible level. The bulk of the inner-vessel contamination attributable to ^{238}U is due to the first gaseous and somewhat long-lived ($\tau_{1/2} = 3.8$ days) isotope of ^{222}Rn . Radon introduced during filling has undergone many mean lives since the original operations, therefore, any radon still present in the inner-




Figure 3.33: The decay series of ^{238}U . The bulk of the contamination present in the inner-vessel is in the form of the radioactive gas ^{222}Rn and its daughters. It is important to note that one atom of ^{238}U will cause a total of fourteen α and β particles deposited in the scintillator. See the text for details on the tagging of a portion of this chain. Figure from [77].

vessel is a product of the emanation of radium in or on the inner-vessel surface²⁸. Since the radon has alternate entry points into the detector, and not just the decay of internal radium, it is usually only safe to quote a number for the ^{238}U entire chain after many mean lives have elapsed since any detector operations have occurred — otherwise secular equilibrium does not hold.

Among this very long list of decays, there are two which occur in rapid succession and afford a very efficient means of tagging these two isotopes, and therefore the chain (at least as far

²⁸There is also some contribution from the diffusion of radon present in the inner-buffer through the inner-vessel.

back as radon). The first isotope to decay is ^{214}Bi , releasing a β — and often one or more γ 's — in a decay with a Q value of 3.272 MeV. The result of this decay is an atom of ^{214}Po with a half life of $164.3\mu\text{s}$, which then decays by α emission to ^{210}Pb with a high Q value of 7.833 MeV. The half life of this decay is long enough that less than 1% of the decays occur within the same trigger window²⁹, but short enough that the probability of chance coincidences is exceedingly low. An additional bonus is provided by the fact that the polonium alpha is sufficiently higher in energy than that from the much more frequent ^{210}Po decay, therefore, false ^{214}Bi - ^{210}Po coincidences are also very rare. These events can be tagged with relative ease even without using any alpha-beta discrimination; but, once tagged, these events are very effective in creating the reference shapes for α and β events — as noted in Section 3.3.2.

An even closer inspection of the decay chain (relevant parameters reproduced in Table 3.7) shows that it is possible to exploit the ^{214}Bi - ^{214}Po fast coincidence to remove another chunk of the ^{238}U decay chain. Upon examining Table 3.7, one can see that the half lives of the isotopes between ^{222}Rn and ^{214}Po are all less than 30 minutes. If one assumes a reasonably accurate position reconstruction code, and negligible diffusion on the time scale of a few hours, all of these decays can be removed by vetoing around the ^{214}Bi - ^{214}Po coincidence position for a time period of several hours preceding the fast coincidence. This method has been used in the *Mach4* analyses by vetoing a spherical region with a radius of 1 meter around the mean position of the ^{214}Bi - ^{214}Po coincidence for a time period of three hours prior to the coincidence. The amount of livetime excluded by this cut is obviously a function of the number of coincidences and is less than 1% for the ^{222}Rn rate reported in Table 3.6, so, it provides a very efficient means of removing a sizeable portion of the ^{238}U chain. In general, these events can be tagged with an efficiency of roughly 90%.

One final note is that of the ^{210}Bi beta decay: the Q value of this reaction is 1.162 MeV,

²⁹Not that it is impossible to resolve them via the clustering algorithm, but doing so does reduce the overall efficiency of the cuts used

Isotope	Particle Type	Energy	Half Life
^{222}Rn	α	5.590 MeV	3.82 d
^{218}Po	α	6.114 MeV	3.10 m
^{214}Pb	$\beta + \gamma$	1.024 MeV	26.8 m
^{214}Bi	$\beta + \gamma$	3.272 MeV	19.9 m
^{214}Po	α	7.833 MeV	164.3 μs

Table 3.7: The portion of the ^{238}U decay chain bounded by ^{222}Rn and ^{214}Po . The short half lives of all of these isotopes allows them to be removed by blinding a region around the ^{214}Bi - ^{214}Po coincidence.

which is very close to that of the neutrinos emitted in the ^{13}N portion of the CNO cycle — these two spectra are very difficult to distinguish and in most spectral fits the sum of the two contributions is left as a free parameter. In principle, it would be possible to determine the rate of ^{210}Bi by measuring the ^{210}Po rate³⁰, however, we know a large amount of ^{210}Po was introduced with the PPO, so, we can not deconvolve the rate attributed to the ^{210}Bi in the scintillator and that introduced with PPO — the only solution for now is to wait until the ^{210}Po rate decays away. Section 8.2.1 presents results of an analysis of the ^{238}U contamination in Borexino.

The ^{232}Th Decay Chain

Also present in the detector materials is radioactive ^{232}Th , which is the parent of a slightly shorter decay chain than ^{238}U , but one that still manages to release ten α and β particles into the scintillator per atom of ^{232}Th . Thorium is present in the earth’s crust at an abundance of three to four times that of uranium, and is often used in welding filler rods. Therefore, special welding procedures and material selection guidelines were established to reduce the contamination introduced by thorium. A pictorial representation of the decay chain can be

³⁰The 22.3 year half life of ^{210}Pb and the non-equilibrium rate of ^{222}Rn makes it impossible to extract the rate from the ^{214}Bi - ^{214}Po coincidence rate.

seen in Figure 3.34.




Figure 3.34: The decay series of ^{232}Th . Not unlike the uranium chain, one atom of ^{232}Th ultimately gives rise to ten particles deposited in the scintillator. See the text for details on the tagging of a portion of this chain. Figure from [78].

Not unlike the ^{238}U series, the ^{232}Th series also contains a fast coincidence which is also a Bismuth followed by a Polonium, however, in this case these are the mass-212 isotopes. The ^{212}Bi decays by emitting a β with a Q value of 2.254 MeV in a 64.06% branch (the other 35.94% of the time, it decays by emitting a 6.207 MeV α to ^{208}Tl). The resulting ^{212}Po lives for a very short time and decays to stable ^{208}Pb with the emission of a 8.954 MeV α . This coincidence is slightly more difficult to tag due to the short meanlife of the ^{212}Po isotope of 432.8 ns because over 97% of the coincidences will occur inside of the same trigger gate,

thus the efficiency of the clustering algorithm must be included in calculating the overall efficiency of these cuts.

Among the many isotopes in the ^{232}Th decay chain there is the particularly troublesome ^{208}Tl , discussed in Section 3.1.3 on the ^8B solar neutrinos. ^{208}Tl , produced in the 36% branch of the decay of ^{212}Bi as described in the previous paragraph, has a mean life of 4.4 minutes³¹, and decays by β emission with a Q value of 5.001 MeV. In the decay process, several gammas are released, most notably one with an energy of 2.614 MeV — the very highly-penetrating gamma ray of concern in the ^8B analysis. The rate of these expected from internal contamination can be estimated by counting the ^{212}Bi - ^{212}Po coincidences — however, even ^{232}Th isotopes outside of the inner vessel will produce ^{208}Tl γ 's that will produce a signal in the inner vessel. The degree to which the buffer shields these events is the subject of a dedicated campaign with a ^{228}Th (a parent of ^{208}Tl) source inserted at the radius of the PMTs (6.35m) using the external source insertion system (described in Chapter 7 of [18]) currently planned for May of 2010.

The ^{232}Th chain also provides a tertiary method of identifying its presence through a series of three consecutive alpha decays — those of ^{224}Rn , ^{220}Rn (often referred to as “Thoron”), and ^{216}Po . Table 3.8 presents the energies and lifetimes of these three isotopes. The efficacy of this method is compromised by the fact that the energy of the ^{224}Rn is only ≈ 380 keV higher than that of ^{210}Po , therefore, one has to either accept a large number of false coincidences, or discard a large portion of the ^{224}Rn events by setting a large cut value for the energy. Indeed, it is still possible to perform the “double-alpha” coincidence of the ^{220}Rn - ^{216}Po events. The ability to tag two different portions of the chain is a valuable diagnostic tool in searching for unsupported (i.e. out of equilibrium with the rest of the chain) contamination introduced by various operations. Following the first deployment of calibration sources in Borexino

³¹This is not an exceedingly long time and one can also attempt to tag this coincidence with ^{212}Bi .

with a “bare-bones” on-axis system, a spike in the number of ^{212}Bi - ^{212}Po coincidences was noticed without a corresponding increase in the number of ^{220}Rn - ^{216}Po coincidences — the likely contaminant was trace ^{212}Pb (parent of ^{212}Bi) remaining on the surface of the inserted components. An analysis of the ^{232}Th contamination performed with both the ^{212}Bi - ^{212}Po coincidence, and with the double- α coincidence, is presented in Section 8.2.2.

Isotope	Particle Type	Energy	Half Life
^{224}Ra	α	5.788 MeV	3.66 d
^{220}Rn	α	6.404 MeV	55.6 s
^{216}Po	α	6.906 MeV	0.145 s

Table 3.8: The “triple-alpha” coincidence present in the ^{232}Th decay chain. The proximity of the energy of the α from ^{224}Ra to that from ^{210}Po makes tagging all three troublesome, but the Rn and Po can be tagged with relative ease.

A further possibility for this set of decays is to perform a likelihood-based analysis in which the triad of events with the maximum probability is selected.

3.4.2 Other isotopes

While *any* isotope present in Borexino has the ability to mimic or mask neutrino signals of interest, the ones of most concern, aside from cosmogenics and those mentioned in the previous sections, are: ^{14}C , ^{85}Kr , and ^{40}K .

^{14}C

The absolutely inescapable background for Borexino is the ^{14}C present in the carbon of the pseudocumene itself. The decay of ^{14}C results in emission of a β with a Q value of 156 keV to produce ^{14}N ; this process occurs with a mean life of 7833 years. The natural abundance of ^{14}C is approximately one part per trillion. However, the Borexino scintillator

is derived from very old petroleum, thus, its ^{14}C concentration has been reduced to a level of $\sim (5 \pm 2) \times 10^{-18}$ grams of ^{14}C per gram of ^{12}C [57] — this corresponds to a ^{14}C rate of $\sim 90\text{Bq}$ for the entire inner vessel³². This very high rate makes the measurement of any other signals below the endpoint of 156 keV impossible (such as the pp neutrinos, tritium in the scintillator, etc). The finite energy resolution of Borexino and the event pile-up does smear the endpoint of the ^{14}C to an effective endpoint of approximately 250 keV.

The high ^{14}C rate can cause problems when two ^{14}C events occur very close in time, but not necessarily in space — this is referred to as “pile-up”. The result is a pulse that will not reconstruct well, or at all, due to the fact that the resulting arrival time spectrum does not approximate that of a point like event. Depending on how closely in time these two events occur, it is possible to deconvolve them using a Fourier analysis like that done in the *Echidna* code, however, this is not done in *Mach4*. For those pairs of events where the fitter can find a minimum, the energy recorded will be that of two ^{14}C events, so, biased towards higher energies (the pile-up spectrum is just the theoretical ^{14}C spectrum convolved with itself and a Gaussian to describe the detector resolution). These events can be tagged to some degree by looking at the pattern of hit PMTs: for pile-up events, one will see two clusters of PMTs (once sorted by time) instead of just one cluster of hit PMTs.

There is a silver lining to the cloud of this very high ^{14}C rate and that is its use for energy calibration — which was necessary during the two years prior to the first source calibration. The energy spectrum of ^{14}C is complicated due to the presence of a poorly constrained spectral shape factor; the details of the fit can be found in reference [79] where the fit was done to CTF data, however, the method in Borexino remains the same.

³²The actual rate is somewhat smaller, both due to the detector threshold, and the uncertainty in the precise $^{14}\text{C}/^{12}\text{C}$ ratio.

^{40}K

Radioactive ^{40}K is abundant in the earth's crust and is often found in large amounts in glass (the PMT bulbs and camera lenses in particular), fortunately the ^{40}K present in the outer buffer does not have the ability to penetrate into the inner vessel. One major area of concern for the introduction of ^{40}K is on the surface of the calibration sources and insertion system components which may have been inadvertently touched by a person and not properly cleaned afterwards. The decay of ^{40}K proceeds in one of two ways:

- β decay to ^{40}Ca with a 89.28% branching ratio. The reaction proceeds with a Q value of 1.311 MeV which means that any potassium present in the inner vessel will create β 's with energies in the ^7Be , pep, and CNO neutrino windows.
- Electron capture decay to ^{40}Ar with a 10.72% branching ratio (see Figure 3.35 for the decay scheme). This decay is identifiable because the decay is to an excited state of ^{40}Ar , which subsequently de-excites with the emission of a 1.46 MeV γ . Unfortunately, the monoenergetic line that this γ provides is initially hidden under the β spectrum of the cosmogenically-produced ^{11}C (the β^+ from this decay is calorimetrically summed with the annihilation γ 's which pushes the total energy deposit into the energy range above 1 MeV). After ^{11}C subtraction has been performed, it will in principle be possible to obtain a value for this contaminant.

 ^{85}Kr

Perhaps the most dangerous background for Borexino is that of ^{85}Kr because it decays by β emission with a Q value of 687 keV — extraordinarily close to the 662 keV Compton edge produced by the ^7Be ν 's. As can be seen in Figure 3.36, in 99.566% of the decays,




Figure 3.35: Level scheme of ^{40}K electron capture decay to ^{40}Ar [80]. The 10.72% branch to the excited state of argon produces a 1.46 MeV gamma ray which should be visible in Borexino after subtraction of the ^{11}C background which covers this area.

the ^{85}Kr goes directly to the ground state of ^{85}Rb by emitting a β with the aforementioned “unlucky” endpoint energy. In the other 0.434% of the cases — in other words, one out of every 231 ^{85}Kr decays (only 1.59/yr in the fiducial volume assuming the acceptable activity listed in Table 3.6) — the ^{85}Kr decays to a metastable state of ^{85m}Rb which has a mean life of $1.464\mu\text{s}$ and then decays to the ground state of ^{85}Rb with the emission of a 514 keV gamma. The β emitted in the transition to the metastable state has a Q value of 173 keV, unfortunately in the energy region of the tail of ^{14}C , greatly decreasing the efficiency of this delayed coincidence tag. Furthermore, the short mean life means that the majority of these decays occur inside of one trigger window, many of which occur in a time scale too short to even be clustered individually. It is unfortunate that the efficacy of this method is so low as it would greatly improve the error on the ^7Be ν measurement by being able to fix the level of the ^{85}Kr in the spectral fit. At present, the statistics are still too low to implement this, but the two methods (delayed coincidence and spectral fit) give results consistent to within the errors of each method.

This isotope of krypton is particularly a nuisance because we humans have created this




Figure 3.36: Level scheme of ^{85}Kr β decay to ^{85}Rb [80]. The 0.434% decay to the metastable ^{85m}Rb allows a means, albeit a rare one, for a delayed coincidence tagging of the ^{85}Kr content of Borexino.

problem ourselves: ^{85}Kr is not a naturally occurring isotope, but is produced in copious amounts in the fission of uranium and has been dispersed through atmospheric weapons testing and venting of radioactive gases produced in spent nuclear fuel rod assemblies. At present, ^{85}Kr exists in air at a concentration of $1 \text{ Bq}/\text{m}^3$; this is particularly troublesome because commercial nitrogen is produced by fractional distillation of liquified air, therefore, strict limitations must be placed on the nitrogen used in Borexino. Assuming a limit of no more than 1 count per day in the entire fiducial volume, and a partitioning coefficient of 1.4 of krypton to nitrogen versus pseudocumene, the nitrogen used for Borexino must have a ^{85}Kr specific activity no higher than $0.14 \mu\text{Bq}/\text{m}^3$ [81]. An analysis of the ^{85}Kr contamination in Borexino is presented in 8.2.3.

3.4.3 Muons and Cosmogenics

The overburden of the mountain on top of the Borexino experimental site in Hall C is large enough to shield all components of the cosmic ray flux except for muons (and atmospheric neutrinos of course). During their flight from their production in hadronic showers in the atmosphere to the detector underground. The muon flux is attenuated by a factor of $>$

5×10^5 from its sea level flux of $\sim 175 \text{ m}^{-2} \text{ s}^{-1}$ [82] to an underground flux of approximately $1.1 \text{ m}^{-2} \text{ h}^{-1}$ [83]. Once inside of the detector however, muons have the ability to produce a very large amount of light depending on their path length — a typical muon traversing the scintillator has a linear energy density of 2 MeV/cm. Therefore, even a muon which grazes the inner vessel can effectively blind the detector for a time period of up to a few ms while the electronics and DAQ recover; Figure 3.37 shows the hit pattern from a muon which passed through the scintillator. The more troublesome aspect of the still-persistent underground muon flux is the production of spallation neutrons and cosmogenics, predominantly through gamma-induced neutron emission from carbon.




Figure 3.37: PMT spatial and temporal hit pattern for a muon event in Borexino (Run 8060, Trigger 1136). The muon entered on the left side of the picture as can be seen by the larger proportion of red pixels on the left (pixel size represents the amount of charge seen, and the color represents the hit time according to the palette in the upper right of the window). It is also possible to see the outer detector PMTs on the floor and outer surface of the detector — these also are consistent with a large, early, energy deposit on the left side of the photo. The track was reconstructed with coordinates: entry(x,y,z) = (−5.7, −3.8, −0.2)m, and exit (x,y,z) = (−4.8, −4.2, −2.5)m.

During the late 1990s, a subset of the Borexino collaboration performed the NA54 experiment with the SPS muon beam at CERN, the goal of which was to study the production of cosmogenics in the Borexino scintillator. Muon beams of two energies, 100 GeV and 190 GeV, were used to study the overall production rates of the cosmogenic isotopes and also to determine the energy dependence of the cross sections³³ — the average muon energy for those that pass through Borexino is 320 GeV[84].

The cosmogenically-produced isotopes most relevant for Borexino are (cross sections are for a muon energy of 190 GeV and are obtained from [84]):

- ^{11}C — Produced by numerous channels in Borexino with a cross section of $905 \pm 58 \mu\text{b}$, with the most common channel being $^{12}\text{C}(\gamma, n)^{11}\text{C}$ [85]. The resulting ^{11}C atom has a half life of 20.38 minutes and decays by positron emission to ^{11}B ; the positron released has a maximum energy of 0.96 MeV, which gets calorimetrically summed with the two 511 keV annihilation gammas that are released when the positron finally annihilates with an electron. The total visible energy resulting from this decay is then 1.02 – 1.98 MeV. The long half life of the isotope requires that the beginning of non-consecutive runs be discarded since any ^{11}C produced by a muon which occurs in a gap between runs is not taggable. In almost all cases, ^{11}C is produced along with a neutron, so, this affords a three-fold coincidence with the muon, neutron, and β^+ from the ^{11}C . Since this is such a predominant isotope in Borexino, it is exploited to determine the γ light yield of the detector by examining the starting point of the ^{11}C spectrum in the spectral fit — in the most recent *Mach4* analysis, this value is found to be about 420 pe/keV. The ^{11}C rate measured by Borexino was expected to be 15–17 counts per day in the fiducial volume ([84], [85]), however, Borexino measures a rate about 50% higher

³³As one would expect, the deeper one places their detector, the higher is the average energy of the muons which reach it, so, to a good approximation, the energy dependence of the production rate is indicative of its dependence on overburden.

at 25 ± 1 counts per day in the 100 ton fiducial volume [20] — it is postulated that this discrepancy is due to the underestimation of high multiplicity neutron events[86] and ^{11}C production by $(n,2n)$ reactions. The presence of this very high rate for ^{11}C makes any pep or CNO neutrino analysis impossible if the ^{11}C cannot be tagged, as its rate is roughly an order of magnitude larger than the expected pep flux, and about five times larger than the expected CNO flux. It is worth noting that the KamLAND collaboration sees the same discrepancy in the ^{11}C production rate that is visible in Borexino (albeit at a different depth)[87].

- ^{10}C — Produced predominantly via $(\gamma, 2n)$ and $(n, 3n)$ reactions with a cross section of $115.4 \pm 14.6 \mu\text{b}$, with a half life of 19.3 seconds. The decay of ^{10}C proceeds via positron emission to ^{10}B . The positron has an endpoint energy of 1.9 MeV, followed by the two annihilation gammas, and also a 718 keV γ (98.5% branch) due to the decay into an excited state of ^{10}B — the total energy visible in the decay lies in the range 1.74 – 3.65 MeV. These events can also be determined via a coincidence with one or two neutrons and the β^+ . The theoretical rate for ^{10}C calculated in [84] is 2.0 ± 0.2 counts per day in the 100 ton fiducial volume; however, unpublished spectral fits indicate a rate of roughly double this number — this can also be attributed to an underestimation of high neutron multiplicity events described above. The high energy tail of this portion of the spectrum does provide a source of background for the ^8B analysis, so, its reduction is a very important part of that analysis.
- ^7Be — Produced by muon-induced disintegration of ^{12}C with a cross section of $230 \pm 23 \mu\text{b}$. Once produced, the ^7Be decays by electron capture to ^7Li with a half life of 53.12 days and a Q value of 861 keV. In almost 90% of the cases, the decay goes directly to the ground state of ^7Li and is effectively invisible since the decay energy partitions almost completely to the neutrino, however, in 10.5% of the decays, the

transition is to an excited state of ${}^7\text{Li}$ which de-excites by emitting a 478 keV γ . This decay energy is worrisome due to its proximity to the ${}^7\text{Be}$ edge, and the effect is exacerbated by the limited energy resolution of the detector which scales approximately as $5\%/\sqrt{E(\text{MeV})}$ and is thus about 7.5% at the energy of this gamma. The long half life and large cross section also dictated that the scintillator spend as little time as possible above ground on its way from the chemical plant where it was produced in Sardinia to the underground labs. In the end, this isotope has not yet been directly observed in Borexino.

- A handful of other short lived ($t_{1/2} < 15\text{s}$) isotopes as shown in Table 3.9. Saturation of the electronics after a muon precludes the ability to specifically detect some of the shorter lived isotopes like ${}^{12}\text{B}$ and ${}^{11}\text{Li}$. The very high beta decay endpoints of these isotopes do make them the only source of background for the high energy ${}^8\text{B}$ neutrino analysis; their removal from the data set is possible by simply vetoing the detector for a time period of five seconds after each muon that crosses the stainless steel sphere [44].

3.4.4 CTF

Before the construction of the full-scale Borexino experiment could begin, it was necessary to prove that it was possible to reach the requisite purity levels and light propagation properties of an organic liquid scintillator. To that end, the Counting Test Facility (CTF)[88] was constructed in the mid 1990s. The CTF contained the same liquid scintillator cocktail ultimately used in Borexino³⁴ inside of a one meter radius spherical nylon vessel (3.7 tons of scintillator) suspended inside of a cylindrical water tank. The CTF vessel is viewed by

³⁴This is only true for CTF campaigns 1 and 3 — CTF 2 was a campaign dedicated to testing the feasibility of an alternative scintillator (PXE). Further information on the CTF 2 campaign can be found in [89].

Isotope	Half Life	E_{max}	σ (μb)	Th. Rate (cpd/100ton)
^{11}Be	13.8s	11.5 MeV β^-	< 2.34	< 0.034
^8Li	0.84s	16.0 MeV β^-	4.02 ± 1.46	
^6He	0.81s	3.5 MeV β^-	16.02 ± 1.60	0.26 ± 0.03
^8B	0.77s	13.7 MeV β^+	7.13 ± 1.46	0.11 ± 0.02
^9Li	0.18s	13.6 MeV β^-	$2.12 \pm 0.34^\dagger$	$0.034 \pm 0.007^\dagger$
^9C	0.13s	16.0 MeV β^+	4.83 ± 1.51	0.077 ± 0.025
^8He	0.12s	10.6 MeV β^-	$2.12 \pm 0.34^\dagger$	$0.034 \pm 0.007^\dagger$
^{11}Li	0.09s	13.6 MeV β^-		
^{12}B	0.02s	13.4 MeV β^-		

Table 3.9: Half lives, beta decay endpoint energies, cross sections (for a muon energy of 190 GeV), and theoretical event rates in Borexino for the short-lived cosmogenics produced by muon interactions with ^{12}C . No cross section data was available for the last two isotopes, and the cross section and theoretical rate for those marked with † are for the combined $^9\text{Li} + ^8\text{He}$ decay since the half lives are very similar and difficult to distinguish.

one hundred photomultiplier tubes (identical to those used in Borexino) mounted on a series of circular pipes — see Figure 3.38 for a photo of the CTF PMT support structure and inner-vessel, and Figure 3.39 for a dimensioned drawing of the CTF.

To a large extent, the contamination introduced to the scintillator after purification is due to the inner vessel surface; since Borexino has a much smaller surface to volume ratio than CTF (by a factor of 4.25), the purity levels measured by the CTF will, in general, be even lower in Borexino with all other things being equal. The fact that the CTF has less shielding than Borexino, meant that the singles rate was too high to check at Borexino levels, so, the heavy-element contamination analysis was performed by exploiting the Bi-Po coincidences. After the first CTF campaign, the collaboration published its results on the $^{14}\text{C}/^{12}\text{C}$ ratio, and the ^{238}U and ^{232}Th levels — these are repeated in Table 3.10. After the CTF1 data-taking campaign, the concentration of PPO was reduced in order to verify that the fluor was not a significant source of ^{14}C .

In 1996, the CTF was refitted with a new inner vessel and the outer radon shroud was also




Figure 3.38: The PMT support structure of the CTF and its inner vessel (one meter in radius, viewed by 100 PMTs) mounted inside of the cylindrical water tank. This photo was taken prior to the CTF2 campaign where an additional outer vessel, or “shroud” was added in an effort to reduce the amount of radon diffusing into the inner vessel. Unlike Borexino, the region outside of the inner vessel is filled with high purity water, thus, the CTF inner vessel is thicker than that for Borexino. The inside of the water tank is coated with a black Permatex resin to minimize reflections.

added; additionally the experience of CTF1 dictated the need for a dedicated muon veto system which was also installed at that time. The CTF2 campaign was centered around evaluating an alternative scintillator as described in the footnote on page 120; however, the results were not as promising as was hoped, and the collaboration decided to adopt pseudocumene as the scintillator.

In 2001, the CTF was re-fitted with new vessels (a few tears and cracks developed in the CTF2 vessels) and the CTF was fitted with three digital cameras³⁵ in order to monitor the

³⁵These are the same model as those used in Borexino — see section 5.1.1 for further information




Figure 3.39: Dimensioned section-view of the CTF during the CTF2 campaign. Note the addition of the new outer vessel which was added after the experience of CTF1. The central yellow region is filled with active scintillator, while the rest is filled with ultrapure water. The entire vessel assembly can be adjusted up or down by a rope and counterweight system in the cleanroom atop the CTF — after filling of the CTF, the counterweight was adjusted to place the center of the spherical inner vessel at the point PNT0. Figure from [57].

filling and condition of the vessels. The CTF3 campaign was dedicated to testing various purification methods. However, during one of these tests in August of 2002, approximately 50 liters of scintillator was inadvertently released into the laboratory drain, and thus a nearby lake. The resulting investigation revealed problems with the laboratory infrastructure, and all fluid operations were halted for a period of two years. Since re-starting operations, the CTF has been used to screen the purified scintillator before being loaded into Borexino, and also for performing tests on gamma quenching, component contamination and various other measurements. There are currently proposals to use the CTF for dark matter detection with

Isotope	Value (CTF)	Method	Reference	Value (Borex)
$^{14}\text{C}/^{12}\text{C}$	$1.94 \pm 0.09 \times 10^{-18} g/g$	Spectral Fit	[79]	$10^{-18} g/g$
^{232}Th	$4.4^{+1.5}_{-1.2} \times 10^{-16} g/g$	^{212}Bi - ^{212}Po Coincidence	[90]	$< 10^{-17} g/g$
^{238}U	$3.5 \pm 1.3 \times 10^{-16} g/g$	^{214}Bi - ^{214}Po Coincidence	[90]	$< 10^{-17} g/g$

Table 3.10: Radiopurity measurements of the CTF scintillator after the first CTF campaign; for comparison, the values measured in Borexino (previously reported in Table 3.6) are repeated in the last column. The measurement of the ^{14}C was the most sensitive ever performed of the isotopic abundance relative to ^{12}C .

internal NaI crystals, as well as a double beta decay facility with Xenon loaded scintillator. Since its construction in 1995, the CTF has been used to publish papers on: solar axion searches[91], pulse shape discrimination[92], deep underground production of ^{11}C [93], electron anti-neutrino interactions[94], experimental limits on the Pauli exclusion principle[95], experimental limits on heavy neutrino mixing[96], neutrino electromagnetic properties[97], limits on invisible nucleon decay channels[98], electron decay via $e \rightarrow \nu + \gamma$ [99], and studies of light propagation in large scintillator detectors[100]; even this is not a completely exhaustive list! It is safe to say that the CTF has been, and will continue to be, a very valuable tool for the Borexino collaboration.

Chapter 4

Motivation For A Calibration

In Chapter 3, the goal of a sub-5% measurement of the ${}^7\text{Be}$ solar neutrino flux was set; and in Section 3.4, the statistical error on a signal comprised of neutrino events and background was established. In addition to the statistical error, one must also add in quadrature the systematic errors on the measurement; these include, but are not limited to, errors on the: energy resolution, light yield, fiducial mass, trigger efficiency, live time, etc. In the most recent ${}^7\text{Be}$ analysis [20], the largest components of the total systematic error were the fiducial mass and the detector response — each contributing 6%.

A rigorous calibration with internal radioactive and optical (i.e. laser) sources can be used to reduce the errors associated with:

- The detector response function

Energy response as a function of position, particle type, and energy (see Section 4.2)

Empirical measurement of the Birks quenching coefficient (presented in Section 3.2.2) with gamma sources

- Fiducial mass determination

Position reconstruction bias as a function of energy, position and particle type

- Alpha beta discrimination efficiency as a function of energy and position
- Efficiency of a distance separation cut for events assumed to occur at approximately the same location (see Section 8.2)
- Efficiency of the Borexino trigger and its threshold (see Section 4.3.1)
- Temporal deviations in all of the aforementioned parameters due to PMTs being connected or disconnected, board changes, modifications to the software reconstruction, etc.

A plan for how one might use calibration sources to quantify and reduce the errors associated with many of these effects is presented in the sections that follow.

4.1 Fiducial Volume Determination

The Borexino detector utilizes a virtual fiducial volume centered on the origin in an effort to obtain the cleanest possible data sample. A complicating factor in the experiment is that this volume does not have any physical structure associated with it, it is just a cut made in the software and as such, its precise definition relies on an accurate determination of the reconstructed position. The measurements released by Borexino (and other such detectors) consists of two parts — the numerator: the rate of neutrino events — and the denominator: the target mass (or, the fiducial volume multiplied by the density of the scintillator). Therefore, the accuracy of the position reconstruction is intimately linked with the ultimate accuracy with which Borexino can report a measurement.

For a spherical fiducial volume like that used in Borexino, applying standard propagation of errors yields the relationship between the fractional error on the fiducial mass due to the error on the reconstructed position:

$$\frac{\sigma_V}{V} = \frac{1}{V} \sqrt{\left(\frac{\partial V}{\partial r}\right)^2} \sigma_r = \frac{3}{r} \sigma_r \quad (4.1)$$

Thus, for a fiducial radius of three meters: $\frac{\sigma_V}{V} = \sigma_r$, so, every 1 cm error on the fiducial radius adds another 1% error to the fiducial mass.

It was explained in section 3.3.2 that the PDF used by the likelihood fitter must account for the case when more than one photoelectron is collected by a photomultiplier tube — the distribution of arrival times is biased towards earlier times for more than one photoelectron. Not applying this correction at all, or applying it incorrectly, will cause the fitter to report a location closer to the PMTs (larger radii). In the *Mach4* analysis framework, this correction is applied for the three cases $n_{pe} = 1, 2, > 3$ ¹, however, this correction was coded before any calibrations had taken place, so, it was not possible to study any position or energy dependence of this effect, it is simply implemented as a blanket correction for all energies and positions.

To show the outward radial bias as a function of energy, we may plot the radial event distribution of various classes of events, and extract the radius with a fit to the expected event distribution. The fitting function must take into account the fact that for most event classes, there is a component due to the uniform distribution in the inner vessel, and also a component due to contamination from the surface (the latter particularly for radon, thoron and their daughters)². A detailed derivation of the distribution functions can be found in

¹Here n_{pe} is the number of photoelectrons detected by that PMT

²Technically speaking, the fit should also include a component due to the external background, however, this is best evaluated with a Monte Carlo simulation as the depth to which the background can reach is dependent on the energy.

reference [57]. The fitting function used is the sum of the bulk and surface components, the former is given by:

$$f_{bulk}(r; R, \sigma) = r^2 C_{bulk} \frac{\pi}{2} \left\{ \text{Erf} \left(\frac{R-r}{\sqrt{2}\sigma} \right) + \text{Erf} \left(\frac{R+r}{\sqrt{2}\sigma} \right) + \frac{2\sigma}{r\pi} \left(e^{-\frac{(r+R)^2}{2\sigma^2}} - e^{-\frac{(r-R)^2}{2\sigma^2}} \right) \right\} \quad (4.2)$$

In equation 4.2, R is the radius of the inner vessel, σ is the position resolution, and C_{bulk} is an overall normalization constant for the bulk portion of the fit. In addition to the bulk component, the component due to the surface contamination may be represented by convolving a Gaussian with a delta function centered on the vessel position; the result is:

$$f_{surf}(r; R, \sigma) = C_{surf} \frac{r}{\sigma R} \frac{1}{\sqrt{2\pi}} \left\{ e^{-\frac{(r-R)^2}{2\sigma^2}} - e^{-\frac{(r+R)^2}{2\sigma^2}} \right\} \quad (4.3)$$

Figure 4.1 shows an example of the fit to the expected radial distribution whose bulk and surface components are described by equations 4.2 and 4.3 respectively.

This fit has been performed to several classes of events during normal, non-calibration, runs and the values are reported in Table 4.1. The results of the fit show an outward bias in the reconstruction of high energy events, but an inward bias as a function of energy for the $\beta + \gamma$ emitting ^{214}Bi . Further, this method of examination for a bias does not account for any position dependent bias (i.e. perhaps the bias is stronger for large radii, or for larger values of x , y , or z); it also assumes a spherical vessel, which we know is not the case from the photos taken with the calibration cameras (see Chapter 5). It is also assumed that the reconstruction is identical in all three coordinates, however, this is not a safe assumption since the distribution of live PMTs is very anisotropic (it is weighted heavily in the bottom of the detector), so, it is reasonable to assume that the z coordinate will, at the very least,




Figure 4.1: The reconstructed radial distribution of a subset of the ^{214}Po contamination events identified by the delayed coincidence technique. The blue line shows the bulk component of the fit described by equation 4.2, and the light green line depicts the surface component of the contamination as described by equation 4.3. The actual fitting function is the sum of these two components, represented by the black line in the figure.

exhibit a larger resolution, if not also a systematic bias.

The only way to truly discover the bias due to both energy and position is to deploy sources with different energies in different locations. The resulting information can be used to either create a correction function for the reconstructed position which is itself a function of position and energy, or, use the information obtained to tweak the position reconstruction module such that the reconstructed positions are determined correctly.

4.2 Energy Response

Even before Borexino began operation, the spatial distribution of photomultiplier tubes was slightly anisotropic as items such as the calibration camera ports, external source insertion

Event Class	Energy Range (npe)	Vessel Radius (m)
^{14}C	60 - 80	4.14
^{214}Po	350 - 500	3.85
^{214}Bi	100 - 600	4.01
^{214}Bi	600 - 800	3.97
^{214}Bi	800 - 1100	3.84
^{214}Bi	1100 - 1600	3.76
neutrons	1000-2000	4.20

Table 4.1: Reconstructed radius of the inner vessel based on fits to various event classes. The fit to the radial distribution of these event classes was done to a sum of components to describe the bulk and surface components described by equation 4.2 and equation 4.3 respectively.

tubes, scaffolding supports, and SSS legs required deviations from a symmetric distribution. After the detector was sealed, some PMTs developed problems for one reason or another, the distribution of which is stronger in the southern hemisphere — see Figure 4.2.

Since there are fewer inactive photo-elements in the south, an event occurring closer to the south pole will exhibit a degradation in energy due to the fact that a larger fraction of those PMTs which “see” a larger solid angle are inactive. This effect is visible to some extent³ by looking at the location of the ^{210}Po alpha peak as a function of the z coordinate. Figure 4.3 shows the result of such a study — approximately 70 days of data taken during the spring of 2008 were concatenated and the ^{210}Po events were selected via an energy and α/β discrimination cut in one meter slices in the z coordinate; the histogram was fit to a Gaussian, and the mean is plotted in the figure. It is quite easy to see the energy degradation which occurs along the z axis.

³Since there really isn’t a tag for ^{210}Po any sample of events selected by energy and α/β cuts will include some component of neutrino events and internal background, but, their event rate is lower by approximately two orders of magnitude, and the underlying spectrum is approximately flat in that region, so, it is not such a terrible approximation.




Figure 4.2: This figure is a graphical display of the status of the PMTs and their electronics chain as of October 21, 2009. The red dots are channels that have a zero rate, and the black channels are those with a very low rate, the latter is indicative of trouble and may just be electronic noise riding on top of an otherwise dead channel. The important point is that there are more troublesome channels in the southern half of the detector than in the northern half.

4.2.1 Spectral Impacts

Figure 4.4 shows a spectral fit to the Borexino energy spectrum obtained with a portion of the available statistics just for illustrative purposes. One can see a very short dip in the spectrum just beyond the ${}^7\text{Be } \nu$ component (blue line) and just before the start of the ${}^{11}\text{C}$ β decay spectrum (dashed green line). One detrimental effect of an energy degradation such as that which has been shown in this section is that the events in this dip, as well as the beginning of the ${}^{11}\text{C}$ spectrum which occur in the southern hemisphere will bleed into the ${}^7\text{Be } \nu$ region. The effect of the energy degradation can be thought of as having a series of ${}^7\text{Be } \nu$ and ${}^{85}\text{Kr}$ spectra, each shifted slightly to the left and summing them all up, the result is a larger smearing of the edge and a larger systematic uncertainty in the spectral fit. In Borexino, the γ light yield is determined by the starting point of the ${}^{11}\text{C}$ component




Figure 4.3: The energy of the ^{210}Po α peak as a function of the z coordinate. The data from the spring of 2008 were cut into one meter thick slices in z , and a Gaussian fit was performed in the appropriate energy range ([140, 240] npe). The non-uniformity of the energy response causes a systematic error on the reported value for the ^7Be ν 's since small energy shifts have a large impact on the balance of counts attributed to ^7Be ν 's and ^{85}Kr .

of the spectral fit⁴; thus, a smeared energy spectrum will give a larger error on the gamma light yield.

4.2.2 Electronics Saturation

Another effect which is important to understand for higher energy analyses (^8B neutrinos for instance) is that of the electronics saturation. The ADCs present in the Laben boards have a dynamic range of [0 – 8] npe, and thus events of sufficiently high energy to produce 8 or more npe in a phototube will have a reconstructed energy biased lower in energy. To

⁴Recall that the ^{11}C signal seen in Borexino is a calorimetric sum of the β^+ decay spectrum and the $2 \cdot 511\text{keV}$ annihilation gammas. When the decay energy is partitioned completely to the neutrino, the only energy visible is that of the annihilation gammas, and this is used to determine the conversion factor from npe to MeV for gammas.




Figure 4.4: The result of a spectral fit to the region of the Borexino energy spectrum relevant for the ${}^7\text{Be}$ ν analysis[101]. The energy degradation present in the southern hemisphere of Borexino due to the larger fraction of inactive electronics channels causes the data obtained in the southern hemisphere to be shifted to the left (lower energy) compared to that in the northern hemisphere. A particular area of concern is the valley between the ${}^7\text{Be}$ ν signal and the ${}^{11}\text{C}$ signal — the degradation of energy in the southern hemisphere means that the latter could bleed into the former and falsely extend the neutrino edge.

illustrate this effect, a simple simulation was performed, the result of which is presented in Figure 4.5.

The simulation of the saturation effect was performed as follows:

- The coordinates of all of the Borexino PMTs were cached locally — all PMTs were assumed to be live, and the presence of a light cone on each PMT was noted.
- An event location (x, y, z) was generated by sampling a uniform spherical distribution with a three meter radius, and an event energy E was also generated from a uniform distribution in the region $[0, 20]$ MeV — the number of photons emitted in the event is obtained by multiplying by the photon light yield.




Figure 4.5: The results of a simplistic simulation to show the effect of the electronics saturation. The Laben boards have a dynamic range of up to 8 npe, therefore, for very high energy events, or events very close to a given phototube, the recorded energy is capped at 8 npe and the resulting energy spectrum for this event is biased towards lower values. Many effects were not included, so, the scale on the x axis could be shifted appreciably, however, the purpose was simply to show the threshold effect — see the text for details on the simulation.

- The distance from the event position to each PMT was used to calculate the solid angle coverage of the event by the i -th PMT (g_i) — the radius of the light concentrator was used for those PMTs which had one, and the radius of the photocathode was used for those that did not.
- The expected number of photoelectrons recorded by the i -th PMT is calculated from $N_{photons} \cdot g_i \cdot QE$, where QE is the quantum efficiency of the PMTs, which is uniformly set at 30%.
- Two running sums of the number of detected photoelectrons were kept — one with the threshold effect (the effect caused by the fact that the ADCs cannot record any more than eight photoelectrons), and one without. For the “no-threshold” sum, the npe value was recorded as is, however, for the “threshold” sum, the npe value was capped

at eight. The sums were then binned into the relevant histograms, and the process began again for a total of 10^6 events.

The simulation does not account for any light attenuation, dead PMT channels, losses due to the light concentrator, reflections, PMT thresholds, different quantum efficiencies, statistical fluctuations, etc., so, it is not to be used for any analysis, but it does show the biasing effect. For analyses performed with the entire inner vessel (i.e. those which can exploit a tag such as anti-neutrino detection, backgrounds, etc.), the effect is even stronger as the threshold is more easily reached by a lower energy event if a PMT covers a larger solid angle of the event, so, the effect is both energy and position dependent. The correction to be applied due to the saturation effect can be determined by deploying high energy sources at radii throughout the IV — a radon source works well as the ^{214}Bi β decay produces events throughout the range $[0 - 3.2]$ MeV, as does an AmBe neutron source since fast neutrons and capture gammas will produce events up to ~ 10 MeV.

4.3 Other Effects

4.3.1 Trigger Threshold

During normal data taking, the detector issues a “neutrino” trigger whenever 20 or more channels fire within a time interval of 60 ns — this value is configurable at the start of the run. It was set to 25 hits for the first several months of Borexino while electronics problems were found and diagnosed. Unfortunately, this does not always correspond to a fixed energy — a noisy PMT or channel for instance will increase the probability that noise / random hits contribute to the threshold crossing, in such a case, the threshold is effectively lowered since only 19 channels need to fire from scintillation light, and the addition of one noisy PMT

allows it to cross the threshold. Obviously, the trigger efficiency is a function of energy, as the fractional contribution of a few noisy channels to the threshold crossing is much smaller when more light is emitted in an event, thus it is only relevant for low energy events.

Of particular interest for Borexino is the decay of ^{85}Kr (discussed in detail in Section 3.4.2) which has a taggable coincidence in a 0.043% branching ratio with the emission of a β with a Q value of 173 keV. A simulated spectrum of this β decay is showing in Figure 4.6. Nominally, the detector threshold is placed at 40-50 keV (BTB = 20-25, the vertical red line in the figure) which includes only $\sim 65\%$ of the events (the cumulative distribution function of the spectrum is presented by the green line in the figure, the fraction of accepted events is just one minus the value on the y axis). The desired location for the ^{85}Kr analysis threshold is just above the apex of the trigger threshold variations due to event position, energy, and live solid angle coverage of the event. As one can see in Figure 4.6, the cumulative distribution function of the energy spectrum has a rather steep slope in the region around the nominal hardware threshold, thus indicating that small variations in the true threshold equate to large enhancements or reductions in the accessible region of the ^{85}Kr β spectrum. It is ultimately the goal of the collaboration to be able to use the coincidence method to obtain an accurate enough value for the ^{85}Kr contamination to constrain it in the spectral fits, therefore, it is vitally important to place the lowest possible cut on the β spectrum to be used in the analysis.

Part of the goal of a calibration should thus be to measure the efficiency of the trigger with sources of different energies, at different positions so as to examine the effect of variations in the live solid angle coverage of the source position. Measuring the trigger efficiency however does require that one knows the precise activity of the source to begin with — since the goal is to see how effectively Borexino can record this value. To that end, the activity of the sources must be measured to a high precision, in a Germanium gamma ray detector for




Figure 4.6: Simulated β decay spectrum of the $Q = 137$ keV decay of ^{85}Kr . The simulation is just the result of using the acceptance-rejection method of generating an energy spectrum according to the Fermi β decay equation, so, it does not include any detector or electronics effects. If it is possible to determine the ^{85}Kr concentration with sufficient accuracy, we may use this to constrain its contribution in the spectral fit and thus reduce our error on the ^7Be ν flux. Also drawn is the cumulative distribution function, or just the running sum of the normalized bin entries; the rather steep slope present in the region around the hardware threshold (approximately marked by the thick red line) illustrates the sensitivity of the actual threshold value to the efficiency of the ^{85}Kr detection efficiency (which is tied directly to how much of the energy spectrum one can consider in the analysis). Precise knowledge of the range of threshold variation as a function of energy, position, and detector status allows one to set the ^{85}Kr analysis threshold in the lowest stable region.

instance, prior to being inserted into the Borexino detector (indeed this was done and the results are presented in Chapter 8).

4.3.2 Photomultiplier Tube Timing

The electronics for Borexino are calibrated in time and energy via a UV laser and fiber optic system[102] which distributes the light to each photomultiplier tube while simultaneously feeding reference pulses to the DAQ to sync the detector trigger with the emission of the

laser light. The central assumption built into the system is that the light will reach all PMTs at the same time and the reconstruction software later computes any offset that is necessary to apply to the photon arrival time to equalize all channels in time (an example of the alignment time of pulser events before and after the time alignment is shown in Figure 4.7).




Figure 4.7: A calibration laser distributes light simultaneously to every photomultiplier tube in the inner-detector. A precalibration module in the reconstruction code uses the laser data to determine the offsets to apply to the photon arrival time. The figure shows, in red, the reconstructed time of the pulser signal (also simultaneously delivered to each *electronics* channel) before the electronics correction is applied; and the result after the correction is shown in black. Figure from [103].

It is thus desirable to test if the time equalization is done correctly, and specifically to test the underlying assumption that the light does indeed reach all of the photomultipliers at the same time. Perhaps the most direct way to test this is by placing a 4π light source at the origin of the PMT coordinate system — since the distance to each PMT is the same, they should all be hit at the same time⁵. Any channels which are not calibrated correctly will be

⁵The PMTs installed on the three meter access port are set back by ~ 40 cm, so, the light arriving at these tubes are delayed by a further ~ 2 ns.

evident as their hit time will depart from that of the mean. This very calibration can be used as a way to determine the time equalization parameters as it is guaranteed to provide simultaneous distribution of light to all channels.

A directional light source can also be used to ensure that all PMTs are cabled correctly. During the early CTF campaigns, a partially opaque vial was filled with radon loaded scintillator and inserted into the detector, it was quickly realized that miscabling had occurred when a set of PMTs facing the opaque portion of the vial fired at the same time as those on the opposite side of the detector.

4.4 Design Considerations for a Calibration System

In the preceding pages, a few of the leading reasons to perform a calibration with radioactive and optical sources in known locations inside of Borexino was presented. Periodically repeating some of these so as to examine any time drift present in the detector is also very important for understanding the temporal drift the detector (if any). The next logical question is to ask during the early design stages is: what would constrain such a system?

4.4.1 Maintaining Purity

It goes without saying that the merit of a calibration campaign which introduces a significant amount of contamination is, at the very least questionable, and at worst can be fatal to the experiment. The entire system must be designed around the goal of minimizing the chance for introducing any contamination. The contamination can be introduced in several fashions:

- Surface emanation — The emanation in this case would be an isotope of radon from radium in the surface or shallow layers of the inserted components. This sort of con-

tamination can be minimized by using materials that are known to be lower in concentration of ^{238}U and ^{232}Th . Another possible source is from the gas inside of any voids in the inserted components — a plastic tube for instance would allow radon gas to diffuse through it, therefore, filling such a tube with LAKN would reduce this. Once the radon has been emanated from the surface, it will produce a series of subsequent α and β decays (for the ^{222}Rn in the ^{238}U chain, it will produce eight; for the thoron in the ^{232}Th chain it will produce five). It is well known that the elastomers which are used to produce O-rings (Viton, Buna-N, etc.) have rather high radon emanation rates and short diffusion lengths ([104], [105]), therefore, O-ring material and size should be considered when designing components for the system.

- Gaseous — The gaseous radioactive isotopes of concern to Borexino are primarily radon and ^{85}Kr . The ^{222}Rn is more troublesome than thoron (^{220}Rn) because it eventually produces ^{210}Bi which mimics the CNO spectrum and also produces counts throughout the neutrino window. The contaminants can be introduced through direct contact with the scintillator just below the gate valve at the top of the inner vessel fill pipe, therefore, any voids must be purged thoroughly with LAKN. There may be pockets of gas trapped in the threads of a bolt or between mating surfaces which could represent very slow “leaks” of this gas into the detector volume, therefore, storing all such components in a LAKN atmosphere will ensure that any gas that does leak out in this fashion is not harmful.
- Particulation — The surface of inserted components may, if not cleaned and handled properly, contain some dust which would wash off into the scintillator during its time inside the detector. In the realm of low-background detectors, dust is considered to be radioactive at it contains 30 mBq/kg of ^{238}U , 41 mBq/kg of ^{232}Th , and 810 mBq/kg of ^{40}K , thus, the assembly of the system and many of the operations should be per-

formed in a clean room to avoid the accumulation of particulate material on component surfaces[106]. The condition of the surfaces themselves also plays a major part in the degree to which particulates will adhere to it; very rough or grainy surfaces will trap and hold particulate more readily than a smooth surface, therefore, any stainless steel entering the detector should be electropolished to reduce the surface roughness. Particulate matter can also be generated by abrasion between surfaces (although these particulates will have the elemental composition of the material from which they were released). Fortunately, pseudocumene has moderate lubrication properties. However, care must still be taken with any closely mating surfaces especially those which have to move relative to each other.

The other component of “purity” which has heretofore remained unmentioned is the *optical* purity of the scintillator which can be degraded through various quenching agents. It is a well known fact that the introduction of elemental oxygen to a scintillator provides a non-radiative path for the de-excitation of the molecular energy. In an aromatic compound like pseudocumene, long term exposure to oxygen results in the formation of, among other things, dimethylbenzaldehyde which is a strong quencher of scintillation light — the typical yellowing of a scintillator which has been exposed to oxygen is a result of this process as dimethylbenzaldehyde is, like most -benzaldehydes, yellow in color. Thus, in addition to minimizing exposure to radioactive gases, care must also be taken to minimize exposure to oxygen due to the quenching effects it causes.

The optical properties of the scintillator may also be altered if components which are partially dissolved or broken down by pseudocumene are inserted into the detector. While detailed chemical compatibility information is not easily available for pseudocumene, one can exploit its chemical similarity to Toluene (methylbenzene) and Xylene (dimethylben-

zene) to anticipate the result of immersion in PC⁶. Of the common construction materials, some of those which are compatible with pseudocumene are: stainless steel, brass, Teflon (PFA, FEP, PTFE), nylon, Viton, nickel, copper, and glass. One should also consider the compatibility of these materials with the common cleaning detergents utilized by the collaboration — brass for instance is not recommended for use with certain alkaline cleaning detergents.

4.4.2 Accuracy of Position Determination

In Section 4.1, the importance of having an accurate reconstruction of an event was highlighted. In particular, the relationship between the error on the reconstructed radius of an event and the error on the fiducial volume was calculated: each centimeter of error on the radius yields one percent error on the fiducial volume. To obtain an approximate figure for the required position reconstruction error, we may examine a simplistic model of the systematic errors. The errors which affect Borexino can fall into three categories: position reconstruction error, detector response function, and a pool of smaller corrections which we will group into “higher order corrections”. If we assume that the higher order corrections contribute a 2% error, and the detector response function contributes another 4%, then, in order to maintain a 5% measurement of the ${}^7\text{Be}$ ν flux, the position reconstruction error must be less than 2% — therefore an accuracy of 2cm on the reconstructed radius. If we wish to go to a 3% error on ${}^7\text{Be}$, the errors must be reduced further — one possible partitioning is 1% for the higher order corrections, and 2% each for the response function and position reconstruction accuracy. Another way to express it, is that six 2% errors combine to give a

⁶Examination of a typical chemical compatibility chart for most materials shows a trend of increasing compatibility with the aromatic hydrocarbon family as the number of methyl groups increases, so, usually one can assume that if a material is compatible with Toluene it will be so with pseudocumene. This is, in the end, just an assumption which must be tested before accepting a material for the system.

5% total systematic.

While these simplistic models do not assume any correlations between the sources of the errors, and neglect the contribution of statistical error (which can become negligible by waiting a sufficient period of time), they are nevertheless useful in creating a goal for the accuracy of the calibration system. A 2% fiducial volume error is not an unachievable goal and it still allows the observation of the annual variation of the ${}^7\text{Be}$ ν flux with room for other errors. Therefore, we adopt the guideline that the calibration system to be implemented must have an uncertainty in the source position of 2 cm *or less*.

4.4.3 Mechanical Considerations

The constraints that have been discussed so far have been motivated by physics and we have not yet considered any restrictions that are driven by mechanical requirements.

Light Blocking

In order for a calibration source to be useful for an energy calibration, the ideal situation is one in which the same amount of light is detected for a “synthetic” source event as for a natural background event. Any differences are due either to quenching in the scintillator in the source vial, or blocking⁷ of the light by the insertion system components themselves. If a significant fraction of the light is reflected off of the components, it could also smear the arrival time spectrum enough that the position reconstruction is affected. Given the relative scale of the detector, source, and insertion system, the only likely source of light blocking will be the device that holds the source. With these considerations in mind, a design goal of less than 1% light blocking has been established for the system.

⁷In this context, blocking refers to reflection of the light off of materials or absorption which leads to the “death” of the photon.

Pressure Control

The filling station of Borexino located in a clean room atop the detector is engineered to maintain a positive pressure in each vessel. A set of three head tanks, one each for the outer buffer, inner buffer, and inner vessel, contain columns of scintillator (PC + DMP for the two buffer tanks, and PC + PPO for the IV tank) and a small nitrogen pressure on top of each. The total pressure in each is approximately 100 mbar relative to atmospheric pressure; and the relative pressures between the vessels are only a few millimeters of water. The positive pressure relative to atmosphere means that a source deployment system must maintain this pressure head on the inner vessel.

One must also consider the implications of a pressure loss during the calibration operation — if all three head tanks are left open to their respective vessels, a pressure loss in the inner vessel region will cause the scintillator to rise in the inner vessel fill pipe and into whatever regions are directly connected to the fill pipe. In order to minimize the amount of liquid lost, the procedure during the calibration operation is to close the valves from the IV and IB to their head tanks — therefore, all volume fluctuations (whether they are caused by the introduction of the calibration source and hardware, or temperature/pressure changes) will be accommodated by flexing the inner and outer vessel films and result in level changes in the outer buffer head tank. In this configuration, if we lose pressure in the inner vessel, the IV will collapse inwards as the liquid level rises in the fill pipe, this will cause the outer vessel to collapse slightly as the scintillator in its fill pipe drops to take up the new volume supplied by the reduction in the size of the IV, then the same thing happens in the outer buffer where the liquid comes from the outer buffer head tank. The total amount of scintillator that can be lost is just the volume of the OB head tank — a few hundred liters.

Access To The Inner Vessel

There are three concentric pipes that lead from the top of the stainless steel sphere to the filling stations, the innermost pipe goes to the IV, the middle to the IB, and the outermost to the OB — it is the innermost one through which we must pass a source and the rest of its components. The Borexino design is rather limited in this respect since many of the other neutrino experiments (KamLAND, SNO, etc.) have very large openings, with diameters well in excess of 30 cm — the IV fill pipe for Borexino is only *four inches* (10 cm). The diameter of the items to pass through this pipe should therefore be kept relatively small to avoid collisions with the inside of the pipe since it was unintentionally installed at a very slight angle off of the z axis. Also, the opening in the IV end region where the fill pipe and the IV meet has only a 1/4" radius fillet on its inner edge, which, on the scale of the experiment, is practically a 90° corner and provides a very minimal “guiding” transition into the fill pipe, therefore, the system components must be designed with this abrupt transition in mind.

Safety

It behooves one to keep safety in mind in the early design decisions as it is much easier to design a system to be safe, than to retrofit an otherwise dangerous system. When dealing with very large quantities of a flammable substance like pseudocumene, care must be taken to prevent ignition sources from coming into contact with the liquid itself or its vapors (the latter are not lethal to humans but cause headaches and respiratory irritation; olfactory fatigue is also a concern with PC and can lead to over-exposure). The prevention of ignition can be facilitated by a combination of 1) attempting to remove the vapors themselves either through containment or a high flow ventilation system, and 2) removing as many ignition sources as possible by selecting explosion proof and non-sparking instrumentation.

In the previous section on pressure control (see page 144), the configuration of the filling station during a calibration was discussed with the emphasis on the minimization of the quantity of pseudocumene that may be released if the pressure over the inner vessel were to be lost. In the event that this were to happen, common sense dictates that one would like it all to be contained in a location where it can be easily cleaned up and not damage other equipment or people. With that in mind, the system must be designed such that any scintillator which has lost containment is directed into a catch basin; further, this catch basin should be equipped with a means to extract the lost scintillator into a sealed tank. It is also highly advisable to have a way to alert the shifters if the scintillator in the inner vessel fill pipe does begin to rise out of its nominal operating range.

As recent events have shown, the location of the Gran Sasso laboratory is very near to *several* fault lines, and *all* Borexino installations must be designed in accordance with the earthquake requirements set forth by the collaboration's engineering group. For the experiment to survive an earthquake it is necessary that the sphere be allowed to move independently from the water tank, and the latter from the laboratory itself. To that end, the concentric pipes which lead from the sphere to the filling stations pass through a flexible rubber coupling on their way out of the water tank; it is to the top of the innermost pipe that we must access the detector for calibrations. It is therefore necessary that the insertion system be attached only to this pipe and not to the top of the water tank, or any walls or the ceiling. Fortunately, the system which will be discussed in later chapters *was* designed with these considerations in mind and everything did survive the L'Aquila earthquake of April 2009.

4.4.4 Adopted Design Philosophy

Having established the requirements and constraints for a calibration system, one can see that it requires a very precise piece of equipment to deploy a source with a 2 cm accuracy from a distance of at least six meters away (it is approximately six meters from the top of the gate valve on the IV fill pipe to bottom of the IV north end cap; add another eight and a half meters to reach the bottom of the detector). A constraint which has not yet been mentioned is that of cost, and it would be nigh impossible to design a system with this degree of position feedback without requiring an enormous machining cost for the extremely high precisions required. The design of the system is therefore centered around the decoupling of the deployment of the source and determining the location of the source — one then uses the deployment system to place the source in the detector and the location system to determine its coordinates.

Early attempts at performing source location by detecting sound emitted from a source proved to be difficult and the idea was abandoned. It was eventually decided to install a set of digital cameras on the sphere and to illuminate the source when it is deployed. The source location system was the subject of the PhD thesis of H. Back [18], however, some important modifications were made to the system during the time of this thesis work and will be presented in Chapter 5.

After considering various possibilities for articulated and flexible arms to be inserted to position the source inside of the detector, the final design adopted is modeled after the human arm: two straight sections with a hinge somewhere in between the two. In this way, the entire system can be inserted straight down on axis through the narrow 4" inner vessel fill pipe and then the hinge can be bent upwards to position the source at the required location. The detailed description of the source deployment system and its operation are described in

full detail in Chapter 6.

Chapter 5

A Source Imaging System for Borexino

In the previous chapter, the arguments were presented for developing a system which can deploy and locate a calibration source with an accuracy of 2 cm, what remains is to explain the system that was adopted by the Borexino collaboration, and in particular, the Virginia Tech group. If a set of cameras (at least two) were to be placed on the detector where they can see the entire inner vessel within their field of view, then the position of a light source inside of the detector can be determined. This process of reconstructing three-dimensional position information from two, or more, two-dimensional photographs is known as *photogrammetry*. Of course, one must ensure that all of the cameras that are used are calibrated such that they all agree with each other. In the sections that follow, the hardware, software reconstruction, calibration, and performance of the system will be presented.

5.1 Source Location System Hardware

To resolve the spatial coordinates of a point using two-dimensional photos requires only two cameras in principle; however, the non-ideal nature of real-world cameras and lenses introduces uncertainties which can be best resolved with the addition of more cameras. Given the anticipated ten-year lifetime of the Borexino solar neutrino phase, and factoring in that the system should be installed a year, or more, prior to the beginning of data taking, the likelihood of a camera failure after initial installation is high enough to warrant contemplation of “extra” cameras beyond the two that are absolutely required. The initial design was one in which six cameras would be installed in a tilted coordinate system on the stainless steel sphere. After some time, the collaborators at Princeton University responsible for the containment vessels requested another camera be installed in a spare port near the north pole to facilitate monitoring of the north end region during the initial filling to prevent gas bubbles from forming. The seven cameras were eventually installed in the locations listed in Table 5.1 — the coordinates given are those of the effective pinhole of the camera. A graphical depiction of the location of the cameras can be seen in Figure 5.1.

5.1.1 Digital Cameras

The cameras used for the source location system in Borexino must meet a few requirements, chief among them is remote controllability because once the water tank has been filled, access to the cameras is impossible. Moreover, the cameras should not contribute appreciably to the radioactive background already provided by the sphere and the PMTs, and should also have the resolution necessary for obtaining the 2cm accuracy requirement. Many types of scientific grade cameras were available that met these requirements, however, their extremely high cost per unit was prohibitive, therefore, a thorough search of consumer grade cameras

Camera	Radius (cm)	Theta ($^{\circ}$)	Phi ($^{\circ}$)
1	660	55.1	65
2	660	55.1	185
3	660	55.1	305
4	660	124.9	125
5	660	124.9	245
6	660	124.9	5
7	660	34.35	240

Table 5.1: The seven calibration cameras are deployed in the spherical coordinate locations listed above. The coordinates given are for the effective pinhole of the camera which is designed to be coincident with the origin of the viewport / dome. The values for theta are measured with respect to the z axis, and phi is measured with respect to north (in Hall C, this is opposite from the entrance doors). Cameras 1-6 form an orthogonal set of coordinate axes, tilted down in theta by 55.1 degrees and off by 5 degrees in phi; camera 7 is installed at a location closer to the north pole in order to monitor for trapped gas pockets during filling. A graphical depiction of the camera locations is presented in Figure 5.1.

capable of meeting the requirements with minimal modification was performed — at the conclusion of which, the winner was chosen to be the Kodak DC290 (whose specifications can be found in Table 5.2).

In order to be useful as a location system, the cameras must obviously be able to see the entire inner vessel, however, the on-board lens of the DC290 did not have a field of view large enough. To solve this dilemma, a fish eye lens (Nikon FC-E8) was added to each camera — this lens has a field of view of 183° and can easily see the entire inner vessel. A complicating factor is that the thread on the Nikon lens does not match that on the Kodak; the differing threads are accommodated by the camera mount described on page 152.

The chief advantage of the DC290 is that it lends itself quite nicely to remote controllability with its DigitaTM scripting language. When turned on, the camera can load its settings from a properly formatted configuration file (`startup.csm`) copied to the on-board memory card, thus, configuring the camera is a simple matter of writing such a script, copying it to

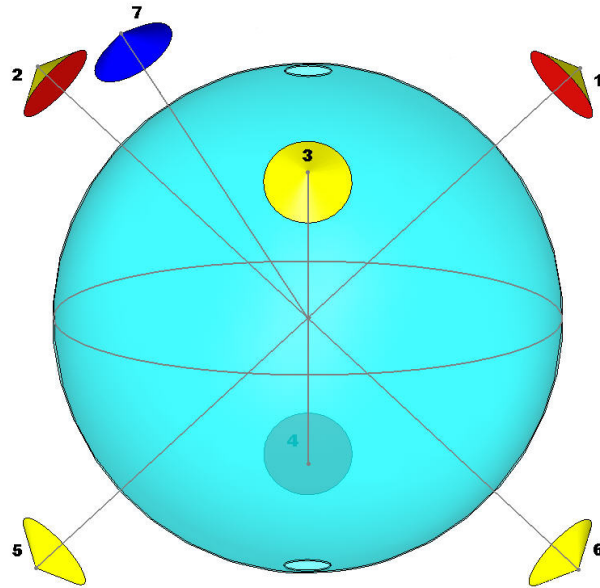


Figure 5.1: The calibration cameras used for the source location system are deployed at the locations shown. The coordinates of the effective pin holes of the cameras are given in Table 5.1.

the camera and cycling the camera power. One important “side effect” is that if one were to copy a malformed startup script, the camera will stall during boot up and require manual intervention. Since this is not possible once the detector is full, extreme care is taken that the software writes correctly formatted scripts. Figure 5.2 is a screen shot of the startup script configuration utility in the camera control software.

5.1.2 Camera Housing

The seven digital cameras, their fish eye lenses and USB repeaters (see Section 5.1.3) are all mounted inside of a stainless steel housing with a thick glass viewing dome on the front (the glass dome is actually manufactured for use in under-water photography). The housing and dome assembly are welded to a DN250 flange which mounts directly onto the stainless steel sphere via a Helicoflex gasket. To make the installation and any necessary repair options

CCD Resolution	1901 x 1212 (2.3 million pixels)
Color Depth	24-bit — > 16 million colors
Picture Format	Exif 2.1 or TIFF
Memory Card	CompactFlash (ATA compatible)
Lens Focal Length	38.4 to 115.2 mm
Zoom	6X
Width	118 mm
Length	63 mm
Height	106 mm
Weight	525 gm (without batteries)
Communication	USB 1.0
²³⁸ U contamination	4.53 Bq/camera — (3 kBq for PMTs)
²³² Th contamination	2.77 Bq/camera — (450 kBq for PMTs)
⁴⁰ K contamination	1116 Bq/camera (4.2 kBq for PMTs)

Table 5.2: Important parameters for the Kodak DC290 digital cameras used in the Borexino source location system. The values for the contamination are taken from reference [18]. Only the ⁴⁰K contamination from the cameras exceeds that from the PMTs. Fortunately, the gamma produced by the decay to the short-lived state of ⁴⁰Ar is only 1.46 MeV and is easily shielded by the several meters of buffer fluid. Only the 2.6 MeV gamma from the ²³²Th daughter ²⁰⁸Tl is worrisome and the cameras contribute a negligible amount to that already produced by the PMTs.

easier, the camera and other components all mount to another flange which seals to the DN250 flange — in this way, the cameras can be removed without also removing the dome and can assembly. Figure 5.3 shows a section view of these components.

To prevent any light from the status LEDs and LCD display from acting as a light leak into the detector, the camera lens collar is attached, via a pressure fit, into a delrin mount. The top side of this delrin mount has the required thread to mate to the Nikon fish eye lens. The mount has a chamfered leading outer edge which mates with the inner trailing edge of the glass dome mount — this ensures that these two components are concentric with each other and fixes the global coordinates of the effective pinhole of the cameras, a vital piece of information in the software reconstruction of the images (see Section 5.2). To allow electrical connections to the lights and LEDs on the detector side of the mount, there are

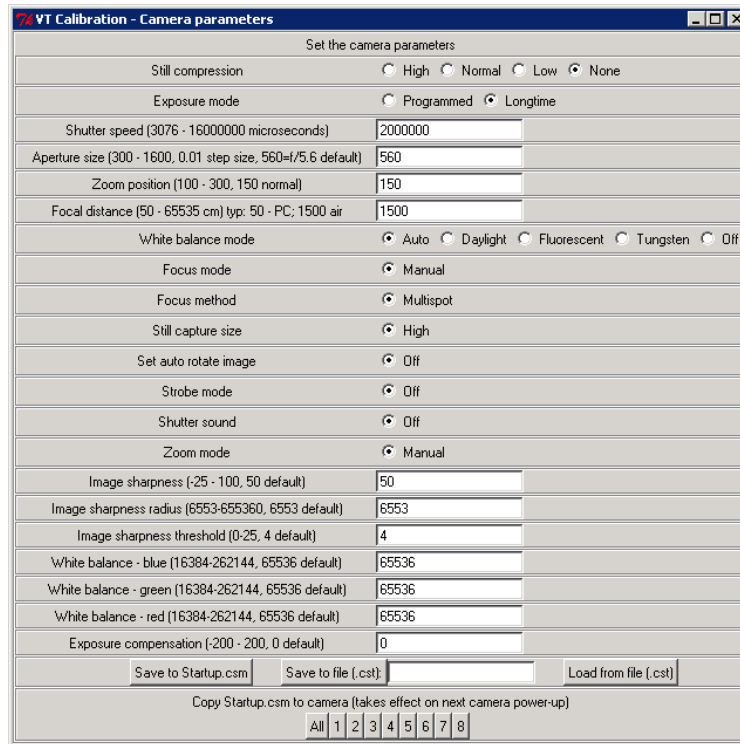


Figure 5.2: Screen shot of the portion of the camera software used to create the configuration script for the cameras. The cameras will load a file on its memory card when it is turned on, thus, preparing a properly formatted script and placing it on the camera and cycling the power allows one to change the camera settings remotely. An improperly formatted script will result in a camera error requiring physical access to the camera to fix it, thus, the program is very careful about only allowing valid values to be placed into the script.

several tapped holes through which pieces of brass threaded rod are inserted and function as mounting posts for the electrical connections.

The USB extender is cable-tied to a thin delrin disc which has the additional function of providing the upward pressure on the camera to keep it pressed securely into the mount. This entire assembly is supported by three pieces of threaded rod — one end of each rod screws into a blind hole in the delrin mount, the other end is attached to a thick aluminum ring. This ring is then attached to the eleven inch flange with shoulder bolts — the water tank side of the shoulder bolts has extension springs which act to provide an inward (towards the detector origin) force on the assembly which keeps the mount pressed securely on the

inside of the dome. The two 1/2" conduits containing the control wiring (see Section 5.1.3) transition to the run ends of a one inch stainless steel tee, the branch end of the tee connects, via a short piece of one inch tubing, to a Swagelok compression fitting welded into the flange.




Figure 5.3: CAD representation of the major components of the camera housing. The dome and housing are mounted to a flange which mates with the SSS; the camera, USB extender, and fish eye lens attach to a separate flange which is bolted to the housing flange, this facilitates the removal of the camera system without disturbing the dome and housing. Picture from reference [18].

Illumination

In order to perform the calibration of the cameras (described in Section 5.2.4) it will be necessary to take photos where the photomultiplier tubes of Borexino are visible, and therefore, the inner-detector must be illuminated by an external light source. Additionally, as a way to normalize images to one another, having one or more fixed light sources visible by

each camera can provide reference points to perform this normalization to. To that end, the camera mount is outfitted with a set of halogen bulbs for the illumination of the entire detector, as well as a set of LEDs for image normalization.

A total of eight quartz halogen lamps, 50 Watts each, are installed in a ring just outside the fish eye lens assembly. The bulbs are wired in two banks to provide redundancy should one set fail. A test set of bulbs were repeatedly turned on and off with an oscillator to verify that they would last throughout the lifetime of Borexino. When an illuminated photo is taken, the lights in all of the cameras except the one taking the picture are used, therefore, the detector is illuminated with 1200 or 2400 watts of light (depending on whether one, or both, banks of bulbs were used in the six cameras not taking the photo¹). This amount of light would be fatal if the photomultiplier tubes were to be on during this operation, therefore, power to the quartz bulbs is only enabled when a key-operated interlock has been activated — this same key is used in an interlock for the PMT high voltage, therefore, enabling the quartz bulbs guarantees that one has already shut down the high voltage.

With the capability for each camera housing to produce 400 Watts of light, the system has been designed to ensure that the lights do not stay on for an excessive period of time and cause any melting or damage to the components inside of the housing. There is a software limit imposed on the maximum amount of time that the lights may be illuminated (capped at 2 seconds) and there is also a hardware limit imposed by a timer circuit which will only enable current to flow to the lights for a maximum of ten seconds.

When the cameras were initially installed, a piece of Tyvek fabric was placed on top of the delrin camera mount to aid in reflecting the light out into the detector. However, when the cameras were removed during the cabling change in 2005, it was noted that some of these

¹Having the lights on in the camera that is taking the photo would cause the image to be saturated, thus, the software does not permit a photo to be taken if the quartz bulbs are enabled for that camera.

pieces had started to distort in the area below the quartz bulbs. When the modifications were finished and the cameras were reinserted, these Tyvek pieces were replaced with a thin sheet of PTFE Teflon, which is much more suited for the high temperatures encountered when the lights are on, and is equally as reflective.

One disadvantage of using consumer grade cameras is that each time a photo is taken or the power is cycled, the lens assembly moves slightly; this introduces a non-negligible shift and/or offset in subsequent images that must be corrected for. These corrections can be accomplished if there are one or more fixed sources of light present in each image. To that end, the initial camera system installation used two standard red LEDs mounted diametrically opposite to each other and in the field of view of the camera (but not obstructing the image of the detector in any way). When the photos were analyzed, the software searched for each of these LEDs in the images, found the centroids and used their positions to apply a relative scaling and (x, y) shift of the image. It was eventually realized that one could do even better if smaller light sources were used.

During the camera system upgrade in 2005, a set of four “pin-hole” LEDs were installed in each camera. The light sources themselves were made from rectangular red LEDs which were wrapped in thick aluminum foil and then painted black with India ink to minimize reflections. Finally, a hypodermic needle was used to carve a hole through the foil², producing the “pin-hole” light source. The LEDs were wrapped with stainless steel wire and soldered into a copper conflat gasket; the gasket fit snugly around the fish eye lens and two extension springs were used to hold them on to the lens. A photo of the pinhole assembly during a test at Virginia Tech in 2005 can be seen in Figure 5.4.

²This process was done in a special fixture to ensure the holes were in the same place on each LED.




Figure 5.4: A test of the pinhole LEDs performed at Virginia Tech in 2005. Also visible in the upper half of the picture are two instances of another pinhole LED system which was being tested at the time but ultimately proved difficult to align.

5.1.3 Cabling

The Digi-Script scripting language allowed the user to completely control the cameras remotely, this included changing parameters without cycling the power, and taking the pictures remotely. Unfortunately, Kodak stopped supporting the scripting language, and we were forced to abandon it as well. The cameras can still be controlled by writing an appropriately-formatted configuration file to the on-board memory card and cycling the power. However, two highly-important functions are not remote controllable in this way: the shutter release button and the power button. These two buttons simply provide a circuit closure, thus, the function of the button can easily be replicated with a relay. After receiving the cameras, they were given to the VT electronics shop where the buttons were removed and 22 gauge wires were hardwired to the terminals.

Communication with the cameras is facilitated with a USB 1.0 connection through a standard digital camera cable. A complicating factor is that the USB protocol has a maximum transmission length of only three meters, whereas the cameras are located at a far greater distance. A solution was provided by commercially available USB repeaters; these devices accept a standard USB cable, amplify the signal and transmit it over twisted-pair CAT-5 cable to a decoder box at the other end where the signal is sent out to the device on a standard USB cable.

There was some concern regarding noise pickup since the CAT-5 cable would be in a cable bundle with wires carrying 220V³. Therefore, a test setup was created at Virginia Tech with the assistance of the University's Communication and Network Services where a network signal analyzer was used to search for the presence of any noise induced by the 220V line. Fortunately, the test provided a null result for any noise pickup.

When the camera system was originally installed in the early part of the decade, all of the cabling for a given camera was installed in a one inch polyethylene conduit which entered the camera housings through a Swagelok connectors without compression supports inside of the conduit. It was eventually decided that over the course of Borexino's lifetime, the probability of a leak developing because of plastic creep was non negligible and in 2005, the plastic conduits were replaced with two $\frac{1}{2}$ " annealed stainless steel conduits for each camera. The wiring is as shown in Table 5.3. Inside of the camera housing, the wires listed in the table connect to either a nine pin (the first eight in the table) connector or a four pin connector (the last four in the table) — the USB power wires go directly to the repeater box, as does the CAT-5 cable.

Also contained in the conduits are the CAT-5 cable for the communication with the USB

³When carrying network signals, the potential for noise pickup is much less because the networks are designed with this in mind, however, we were unsure of the method in which the USB signals were transmitted over the CAT-5 cable, so, the test was in order.

Terminal Number	Function	Color (except Cam. 4)	Cam. 4 Color (below CR4 splice only)
1	220 VAC Neutral	Green	
2	Bank 1 Lights 220 VAC	White	Gray
3	Bank 2 Lights 220 VAC	White	Black
4	LED Power return	Blue	
5	Bank 1 LED +5 VDC	Yellow	
6	Bank 2 LED +5 VDC	Yellow	
7	Camera Power +7.5 VDC	Red	Blue
8	Camera Power return	Black	Blue
9	USB Power +15 VDC	Brown/White	
10	USB return	Brown	
11	Camera on/off	Red	
12	Camera on/off return	Black	
13	Shutter release	White	
14	Shutter release return	Green	

Table 5.3: Wiring table for the camera control wiring after replacement in 2005. Camera 4 was the last one to be rewired, and during the process, it was discovered that there was insufficient wire of some colors to complete the job. Fortunately, there was sufficient wire of other colors that were used for some of the connections, these are noted in the rightmost column of the table. Once reaching the cleanroom, these misfit colors were spliced to the correct colors so the wires arriving at the control box are as listed in the third column.

repeater module, and a piece of $\frac{1}{4}$ " Teflon tubing which is used to flush the camera housing with nitrogen to remove any moisture from inside of the dome⁴. The Teflon tube is potted into its conduit with RTV silicone sealant, this ensures that the exhaust gas exits through the other conduit. All of the conduits enter a nitrogen-flushed break-out box in the cleanroom where all the cables transfer to a piece of large flexible conduit to the enclosure which houses the camera control box.

In addition to the wiring for the seven digital cameras, there is also one spare set of wiring which is routed to a table in the clean room above the detector. This eighth set of wiring is used to connect either an additional camera for testing and debugging, and also for operating

⁴The lack of oxygen also ensures that there cannot be any fire inside of the dome if something were to spark or the lights be left on too long.

the cameras when removed from their housings. When the wiring was upgraded, this spare set of connections was used extensively to align the new pinhole LEDs (see Section 5.1.2) in an easily accessible location.

5.1.4 Camera System Control

Control over the camera system is provided by a program written at Virginia Tech which interfaces with the control box built at Virginia Tech (shown in Figure 5.6) that actually contains the power supplies, relays, and interlocks. Each camera housing contains a series of devices which must be able to be operated independently from each other, as well as from the other cameras: bank 1 quartz bulbs, bank 2 quartz bulbs, bank 1 LEDs (pinhole style), bank 2 LEDs (“old” style), the shutter button, and the camera power switch. In addition to the aforementioned functions, we also need to cycle the power to the AC adapter input to the cameras, and cycle the power to the USB repeater boxes — Figure 5.5 shows the software interface for all of these operations. All of these functions are operated by individual relays for each of the cameras.

When a source calibration is in progress, it is necessary to illuminate a diffuser attached to the source in order to determine the source location and this process must be done with the PMT high voltage on, therefore, the amount of time that the diffuser is illuminated should be kept as small as possible. The minimum amount of time can be achieved only if the cameras are properly synchronized so that they all take their photos at the same time while the laser is on. Since the pinhole LEDs also have to be illuminated at the same time as the photos are taken, they need to be synchronized as well. However, the control box is operated via a serial port, the serial nature of which precludes synchronous photo taking. The solution is provided by connecting the output of all of the shutter release relays to one

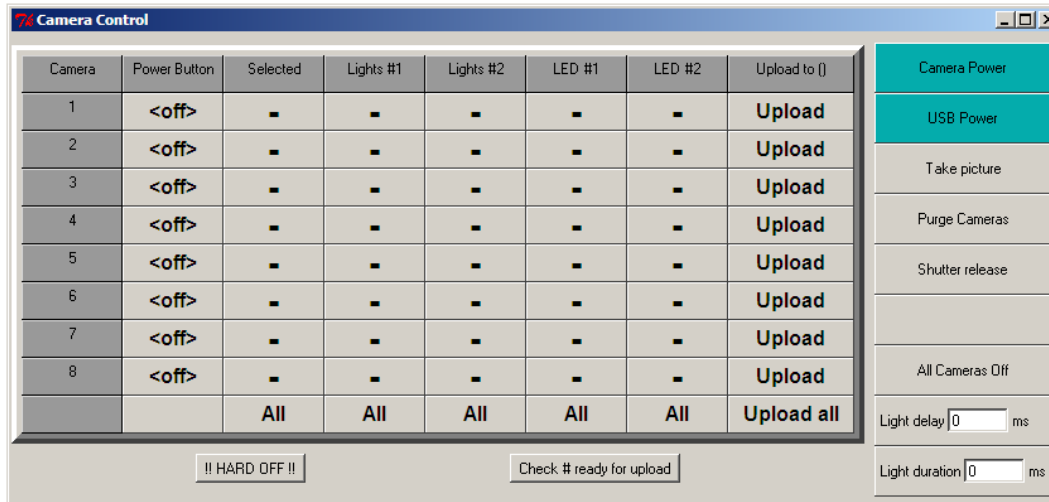


Figure 5.5: The software interface for controlling the various aspects of the camera system. Note that there are only seven cameras installed, but there is hardware for controlling an eighth camera which can be either simulated or a real camera attached for debugging purposes. The cameras to be used to take the photo will appear with a “+” in the third column; bank one and two lights can be selected in the next two columns, and the two banks of LEDs are operated from the sixth and seventh columns.

master “execute” relay for the shutter release function. To take a photo with all cameras (or any combination of cameras), first the software closes the shutter release “enable” relays for all of the cameras — this provides the shutter release signal to the upstream side of the execute relay, and when the photo is ready to be taken, the execute relay is closed and the shutter release signals leave the execute relay at the same time for all cameras. The same configuration of eight enable/mask relays that feed one execute relay is used for the bank one and two lights as well as the bank one and two LEDs.

For reasons not explained here, when power is applied to the control box, all of the relays are closed. To prevent accidental photo taking, LED flashing, etc., an interrupt relay is placed in the system which prevents any signals from being sent to the camera assemblies until a manual “Output Enable” switch is pressed. The box was initially constructed with this switch and the quartz light interlock installed on the face of the box, however, safety



Figure 5.6: The camera system interface box in the computer cabinet in CR4. The box houses power supplies, relays, and interlocks for controlling the lights, LEDs, power, and picture taking ability for all of the cameras. A series of status LEDs along the top of the box are used to give visual feedback of which lines are drawing current. The original design of the system used an LED inserted with the source, the intensity of this LED would be controlled by the potentiometer knob on the front of the box, in its current implementation this is unused. The interlock for the quartz bulbs and the output enable switch have been moved to an enclosure accessible by shifters in the cleanroom.

considerations dictated that the box be placed outside of the cleanroom to guard against sparking and ignition of pseudocumene fumes. Thus, these two items were re-routed to a break-out box mounted on a table in the cleanroom.

5.2 Software Reconstruction of Images

The source location system for Borexino is rather highly over-constrained in that the seven cameras each give a set of (x, y) pixel coordinates — thus the inputs are the 14 parameters

from the cameras and the output is simply the three source coordinates, so, there is significant room for loss of information due to either non-functioning cameras or a camera which is unable to find the diffuser attached to the source. The process is highly non-trivial however, and relies on the cameras having been calibrated to a set of fixed points.

5.2.1 Camera System Geometry

In a digital camera like those utilized in Borexino, the photosensitive element is a CCD which is located where film would normally reside. The CCD is organized in a two-dimensional grid of individual photo-sensors which simply integrate the number of photons which enter it during exposure — in essence, they are miniature PMTs without the multiplication stage. If this were all there is to the setup, the images taken would all be in gray scale since this method does not afford color discrimination. To produce a color image, a series of red, blue, and green filters — known as a Bayer array — are laid atop the photo-sensors, thus, each sensor is capable of seeing only one of the primary colors; they are grouped into 2×2 arrays consisting of two green filters⁵, one red and one blue. These 2×2 arrays, or mosaics, are then grouped together to form a color sensitive cell. A fortunate side effect of this is that the individual color pixels can be read out when the image is taken in TIFF or RAW format; thus, for the red diffuser light used during a calibration, the software can search just the red pixels for the bright spot. It should be noted that since these cameras were created in the early 2000's, newer methods of obtaining color images from CCDs are now being used by most camera manufacturers.

A problem with consumer grade cameras is that they are not made to the exacting standards required for precise, repeatable, location determination — the lenses are not necessarily

⁵The human eye is much more sensitive to green light than red or blue, thus the duplication of the green pixels.

concentric with the optical axis nor are they necessarily coplanar with the other lenses or with the image plane — the same goes for the CCD itself, it may be off center or tilted. For a perfect system, the image analysis is a simple exercise in ray tracing; however, the aforementioned departures from the ideal situation require a more involved approach.

The method adopted to account for these errors and analyze the images is to first *transform* our non-ideal system to an ideal one, and then, once we have an ideal system, we perform the ray tracing. The geometry of our simulated ideal camera is the one depicted in Figure 5.7 — a ray starts from a light source, passes through the simulated CCD and finally reaches the camera pin-hole situated at the origin. The task then is to develop a method to transform between the two systems.

5.2.2 Transforming Between Simulated and Real Camera Geometries

In order to perform this transformation between the real and ideal systems, we must consider the possible departures from ideality for each lens and CCD assembly⁶:

- Pixel scale — Each pixel is considered to be a square element of the otherwise rectangular CCD array but the pixels may not have unit size, thus, we leave the overall pixel scale as a free parameter.
- CCD offset — It may be that the CCD is not centered on the optical axis and could be off in the x or y direction (recall z is in the direction of the optical axis), therefore, an offset in both the x and y directions are left as parameters for each camera. The offset is depicted in the top right box in Figure 5.8.

⁶In the following discussions, the coordinate system adopted is left handed: x is to the right, y is up, and z is through the camera.




Figure 5.7: Graphical representation of the ideal geometry of the *simulated* camera. The pin hole of the camera resides at the origin, and the CCD array sits in front of it at a fixed offset — this offset is the real-world distance from the focal point of the camera assembly to the actual CCD. One of the rays from the light source is shown as it passes from the center of the light source, through the CCD and terminates at the camera pin hole.

- Pitch and Yaw — The CCD and camera may not lie perfectly in the x-y plane. If the CCD is rotated by an angle about the y axis, this is known as “yaw” (top left in Figure 5.8); a rotation about the x axis is known as “pitch” (bottom left in Figure 5.8). These angles are also parameters that must be determined.
- Roll — The entire CCD might be tilted slightly about the z or optical axis (so the edges are not parallel with the x and y axes); this “roll” angle (bottom right in Figure 5.8) is yet another parameter to be determined for each camera.

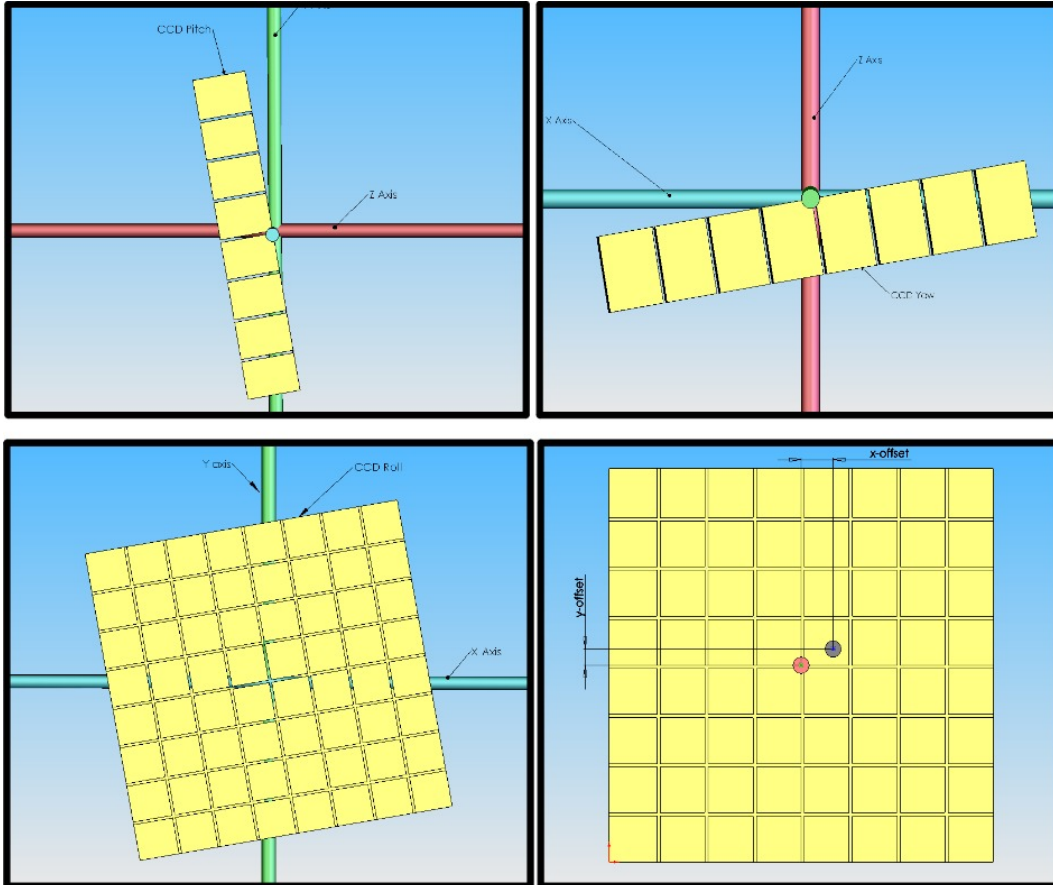


Figure 5.8: Pictorial representation of the three rotations and one offset for the CCD relative to the origin. Counterclockwise from top left: CCD Yaw — rotation of the CCD about the Y axis; CCD Pitch — rotation of the CCD about the X axis; CCD Roll — rotation of the CCD about the Z axis; CCD offset — a shift of the center of the CCD from the z axis (the red circle is the center of the CCD, the black circle is the z axis). In addition to these, there is also the pixel scaling which is not shown.

- Radial distortions — The lenses in each camera may (and indeed all real lenses do) induce a point-symmetric radial bias in the reconstructed positions. The radial distortion is handled with an expansion in odd powers of r up to ninth order: $r' = \sum_i^{1,3,5,7,9} c_i r^i$, where the c_i are parameters to be determined for each camera[107].
- De-centering distortions — It is possible that the lenses cause an angular distortion in the θ direction (using polar coordinates for the CCD) due to components of the compound lenses not all being coaxial. The correction for such de-centering distortions is accomplished via equations presented in Section 5.2.3; for now, we just mention that it depends on two parameters p_1, p_2 for each camera.

Image Tweaking

Whenever the camera power is cycled or a photo is taken, the lens assembly moves slightly and does not return to its previous location, therefore, subsequent images of a static object may yield different results. To ameliorate the problem, a set of fixed point light sources were installed in the field of view of each camera (the hardware was described on page 157). When the cameras are calibrated, the location of the centroid of each of these LEDs is saved as the “tweak master” position (four points for each of the seven cameras). When a subsequent image is to be processed, the software searches the image for these same four points which have of course shifted because of lens movements. The difference in the master positions and those in the current photo can be resolved with an overall image scaling and offset in the x and y directions. Once these three parameters are determined for each camera, they are saved and used in reconstructing the position of a light source. A set of photos from the software which shows the difference between a randomly selected image and its calibration reference image is shown in Figure 5.9.

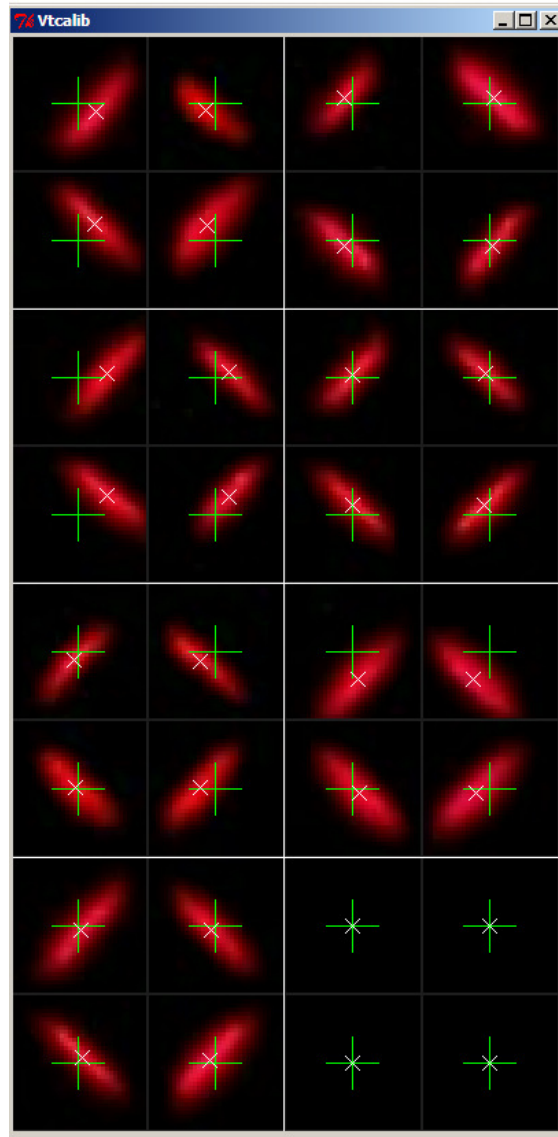


Figure 5.9: Tweak correction of an image. Since the camera lens does not return to exactly the same location after cycling power or simply taking a photo, a method for normalizing images to each other was required and is facilitated with the four pinhole LEDs shown (the light source is circular, but the fish eye lens distorts it to the ellipsoidal shape shown). The four pinhole LEDs for each camera are shown, and their centroid is depicted with a white 'x' — the location of this centroid in the calibration image is shown with the green '+'. The software minimizes the distance between the 'x' and '+' by applying a scaling and (x, y) shift to the image — these are the three tweak parameters.

For the moment, we assume that we know all of the camera parameters (yaw, pitch, roll, offsets, etc.) and then ask how to determine the spatial location of an identifiable object in the detector — this is the subject of the following section. For the discussion of how these parameters are determined, see Section 5.2.4.

5.2.3 Position Reconstruction From Images

The entirety of the previous section has explained the various corrections that are required to transform an image from the real world system to our ideal camera system; now we turn our attention to how these are used to determine the location of an object in the detector. First, we must obviously have taken a photo with two or more cameras — more correctly, we require that the object to be identified is *visible* in at least two cameras — then, for each camera, we have the x and y pixel coordinates that correspond to the point on the object we wish to know the location of, we call these sets of points $\{p_x\}$ and $\{p_y\}$. The general procedure is to construct a ray from the origin, through the fired pixel and onto the dome; then refract the ray through the dome and transform to the detector coordinate system. With rays from two or more cameras, we find the location where the distance between these rays is a minimum, this is our reconstructed point. The details of this procedure are discussed in what follows.

For a given camera, we are given the x and y coordinates of the hit pixel on the virtual CCD (recall the geometry from Figure 5.7) and we of course know the position of the pinhole (defined to be the origin). However, we cannot directly use the fired pixels because these are pixels on the *real* CCD — whereas, our method relies on working with the *ideal* CCD pixels, so, we must transform them to “undo” the effects discussed in section 5.2.2.

Pixel Corrections

We are tasked with performing the transformation from $(p_x, p_y) \rightarrow (p'_x, p'_y)$ — the primes signify *ideal* coordinates, unprimed denotes the *real* coordinates. The first thing to do is to utilize the tweak information to normalize this image to the calibration image that was used to derive the parameters used in the rest of the corrections. The tweak corrections (see page 168) are an offset in the x and y directions — x_{tweak}, y_{tweak} — as well as an overall scaling s_{tweak} . The first stage of the correction is then:

$$p'_x = \frac{p_x - x_{tweak}}{s_{tweak}} \quad p'_y = \frac{p_y - y_{tweak}}{s_{tweak}} \quad (5.1)$$

At this point, the pixel coordinates in the ideal system (p'_x, p'_y) have been corrected for the lens shift between subsequent images, now it is at least consistent with the calibration image. From Figure 5.8, we see that we also have to shift the pixels in the x and y directions by (x_{offset}, y_{offset}) so that the origin of the real and ideal CCDs coincide:

$$p'_x = p'_x - x_{offset} \quad p'_y = p'_y - y_{offset} \quad r = \sqrt{p'^2_x + p'^2_y} \quad \theta = \tan^{-1} \left(\frac{p'_y}{p'_x} \right) \quad (5.2)$$

The next step in the process is to correct for the de-centering and radial bias introduced by the lens system. In performing these corrections, we will find it easier to work with the radius normalized to six meters, so, we define $r_n = r/600$ with r as defined in equation 5.2. The radial bias correction was defined in the previous section to be an odd-power expansion in r_n to ninth order — to that end, we define the corrected radius as:

$$r' = c_1 r_n + c_3 r_n^3 + c_5 r_n^5 + c_7 r_n^7 + c_9 r_n^9 \quad (5.3)$$

In the penultimate step of the pixel correction process, we apply the radial and de-centering corrections according to:

$$p'_x = \cos(\theta)r' + r_n^2 \{p_1 (1 + 2 \cos^2(\theta)) + 2p_2 \sin(\theta) \cos(\theta)\} \quad (5.4)$$

$$p'_y = \sin(\theta)r' + r_n^2 \{p_2 (1 + 2 \sin^2(\theta)) + 2p_1 \sin(\theta) \cos(\theta)\} \quad (5.5)$$

Finally, we take the pixel coordinates in equation 5.4 and de-normalize it by multiplying by 600, and then scale the pixel up by the overall pixel scale p_{scale} :

$$p_x = p'_x \cdot 600 \cdot p_{scale} \quad p_y = p'_y \cdot 600 \cdot p_{scale} \quad (5.6)$$

We now have the coordinates of two pixels on our ideal, virtual, CCD which lies in front of the camera origin. This process is done for each camera that we are using in the analysis, each of which has its own values for x_{tweak} , y_{tweak} , s_{tweak} , x_{offset} , y_{offset} , c_1 , c_3 , c_5 , c_7 , c_9 , p_1 , p_2 , and p_{scale} . The next step is to use the corrected pixels from each of the cameras to define rays in the detector and find the point where the distance between the rays is a minimum.

Ray Propagation and Minimization

After performing the previous pixel corrections, we now have two points inside of the camera housing — one is the origin / pinhole, the other is the coordinates of the pixel which was fired. We now have a ray that shoots from the origin, through the fired pixel and into the glass dome — we can calculate the point of intersection of a line with the spherical dome surface using standard formulas from geometry, this point is (x_d, y_d, z_d) . The dome obviously has different indices of refraction on each side, so, we refract the ray through the glass into

the scintillator and finally we have a ray *in the detector* which passed through the fired pixel, however, the coordinate system is still that of the camera that took the picture.

In order to perform any sort of minimization, we need to know the points defining the ray in the detector frame of reference — recall the camera uses a left-handed coordinate system, and the detector uses a right-handed system, so, there is obviously a transformation required there; in addition, the CCDs are slightly tilted about the principal axes, so, those have to be undone as well. The roll, pitch, and yaw angles that have been determined for each camera are used, in that order, to re-orient the CCD with the camera's principle axes with standard rotation matrices. Then, the transformation to a right-handed coordinate system is done simply by swapping the coordinates: $x \rightarrow -z$, $y \rightarrow x$, $z \rightarrow y$. At this stage, we have a right-handed coordinate system whose axes are still oriented with the camera — in other words, our new y axis still points through the CCD — so, we need to rotate our coordinate system in θ and then by ϕ according to the camera mounting angles reported in Table 5.1.

In an ideal situation, at the conclusion of this process we have one ray for each camera which points from its dome to the point in question and all of these rays intersect. In the real world, various imperfections will cause the rays to come very close to intersecting, but never actually intersect, the goal is then to find the location whose distance to each ray is at a minimum. The procedure for doing this is to loop through every pair of two rays, and for each pair we calculate the shortest line possible which connects the two rays and record the coordinates of the midpoint of this line, the weighted average of these points is the starting point for the minimization. Then, a simple grid search is performed around this point; at each candidate point, the distance to each ray is calculated and the process stops when the minimum has been located — this is our reconstructed position. The software reports the minimized location as well as the residuals, i.e., the distance from the reconstructed position to the ray from each camera — in general, we reject any camera whose residual is greater

than 2cm, although sometimes they were permitted if only a few of the cameras found the LED. The graphical result from this process can be seen in Figure 5.10 for a calibration source located at (13.2, -4.4, 200.4) cm — note, this is the position of the *source*, not the diffuser (the object in the photo).

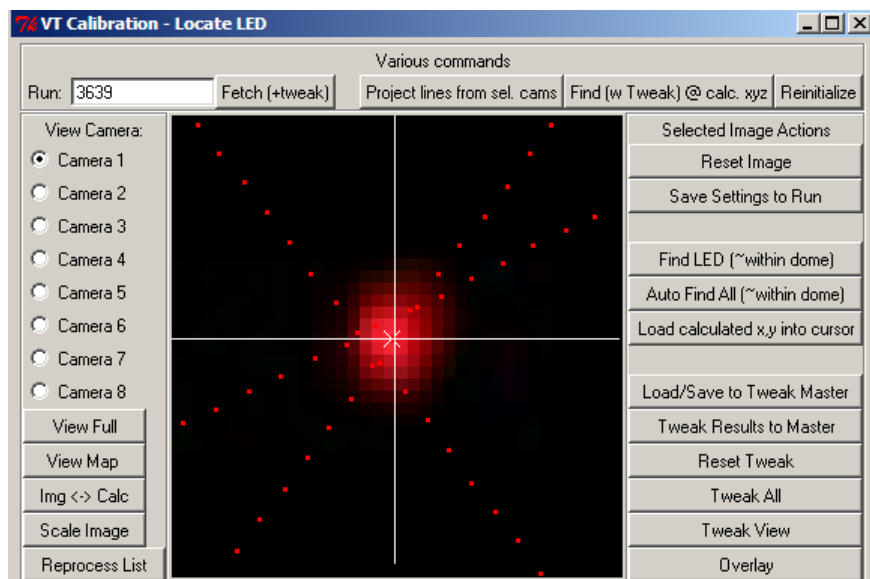


Figure 5.10: The result of the image location process of a diffuser illuminated with red laser light. The centroid of the illuminated pixels in each camera were used to propagate a ray from the origin, through the CCD, refracted through the dome and into the detector. Thus, each of the four cameras used for the reconstruction produced a ray from the dome to the diffuser; the software then performed a least squares minimization to find the three-dimensional position where the distance to each ray was at a minimum. The dotted red lines in the picture show the rays from each of the other three cameras as they pass near the source — the residuals (distance from the reconstructed position to the ray from a given camera) were: (0.78, 1.27, 1.25, 1.20) cm for cameras (1, 2, 4, 6).

5.2.4 Camera Calibration

In the treatment of the camera system software up until now we have taken the camera parameters as values we are given, but of course they must come from somewhere and this is the subject of this section. In normal usage, we utilize the camera system, with its fixed

list of parameters⁷, to determine the *unknown* position of an object in the detector; this of course relies on us knowing *a priori* the camera parameters. However, if we were to, in effect, run this system in reverse, we can use objects in the detector with *known* positions to work backwards and determine the camera parameters. Quite fortunately, there are 2214 objects in the detector whose positions are known very accurately: the PMTs!

The method of determining the camera parameters is referred to as the camera “calibration” and is one of the very few non-automated tasks of the calibration system. The general procedure is to load the known positions of the PMTs from a text file⁸ and then, for each camera, we can create a picture of what the PMT layout should look like from that camera. This calculation of the simulated picture is done with all of the parameters set to zero (no tweaking is done either — all subsequent images are tweaked relative to this one, thus, the pinhole LED positions from this photo are considered to be the “Tweak Master” data). This simulated image is then overlaid onto an actual photo of the detector with the lights on, and the user can drag the overlay until he or she is satisfied that the simulated photo matches the real one *for one PMT* and then saves the point. This process is repeated for a hundred or more points, and finally a minimization is performed to determine the camera parameters required to match up the simulated and real images. The details of the process are explained further in what follows.

The Simulated Image

Recall that in order to determine the parameters that define each camera, our chosen method is to exploit the fact that we know, with a very high accuracy, the position of the 2214

⁷Neglecting the tweak corrections which change with each picture — here we are referring to the CCD rotation angles, radial and de-centering distortion coefficients, etc.

⁸This same file is used by both the *Echidna* and *Mach4* analysis codes for performing position reconstruction

photomultipliers, therefore, we can create a simulated picture of what an ideal camera would see. Then, if we can make the ideal picture match the simulated one, we can use this information to determine the parameters which transform the ideal camera to a real one. The features which we wish to match are the outer edges of the PMT light concentrators and the outer edge of the μ metal shielding, although, the former is the most commonly used (recall that the PMTs used for the inner detector muon veto do not have a light concentrator). Since we know the dimensions of the PMT assembly as well as the location of the mounting hole for each PMT, we can determine the shape of a light concentrator located at angles $\Delta\theta$ and $\Delta\phi$ from the camera Z axis.

The simulated image is produced by a Fortran routine (`raytrace.f`) which steps through every single pixel of the CCD and projects a line from the origin through that pixel, however, it does not perform the tweak or rotation corrections because we are producing the simulated image from an *ideal* camera. This ray is propagated up to the glass dome (if the ray hits the inside of the housing, or is not in the camera field of view, it is terminated at this point), refracted through it, transformed to the detector coordinates, and then propagated through the detector until it reaches one of the following surfaces: inner vessel, outer vessel, interface between two liquids, a light concentrator edge, a mu metal edge, or the stainless steel sphere. If the ray intersects the edge of the light concentrator or the mu metal shield, this pixel in the simulated image is colored blue or red respectively; otherwise, the pixel is left uncolored⁹. The simulated image is then overlaid on top of the photo which we wish to calibrate the cameras to — this is shown in Figure 5.11.

As the `raytrace` program traverses the entire CCD array, it saves the coordinates (pixel and spatial) of each PMT that it finds to a file for later use. Once this picture has been created and is overlaid on the image to be calibrated, it is usually very far away from agreeing

⁹This is not entirely true: the software also paints the edges of the inner and outer vessels yellow, and the three meter port and camera housings are painted green in the images.

with the photo. The user then begins a process of dragging the overlay until a PMT on the overlay image is well aligned with the real image — the user can then save this point (what is actually saved is the center of the light concentrator or μ metal shield). The user then repeats this process for many points — usually more than 100 points *per camera*.

At the end of this process, each camera has a series of calibration points which are: the (x_{OL}, y_{OL}) pixel coordinates of the center of the PMT in the *overlay* (hence the “OL” subscript) as well as the (x_{obs}, y_{obs}) pixel coordinates of the center of the *observed* PMT position — the “overlay” coordinates refer to the result of the image corrections described previously in Section 5.2.2, but at this stage no corrections have been applied yet. It is then possible to calculate the squared deviation between the overlaid PMT and the observed PMT — we refer to this as a quasi- χ^2 .

At this point, the software performs a least squares minimization on all of the camera parameters. After each iteration step, this “ χ^2 ” is calculated and the process iterates until a minimum has been located — the 91 parameters (13 parameters for each of the seven cameras) for which the χ^2 is a minimum are then saved in the file `camparam.dat` where they are later parsed for use in other image processing operations.

In theory, at the conclusion of this step, one would be finished with the calibration of each camera, however, it is often the case that further improvements can be made by repeating the process beginning from generating the simulated with the `raytrace` program (of course using the camera parameters found during the first iteration). The user then goes through and fixes more positions by dragging the overlay image until it matches a PMT in the photo, and then performs the minimization again. As of the time of this writing, the current camera parameters are listed in Table 5.4.




Figure 5.11: A simulated image overlaid on top of the image to be calibrated to. The simulated image is produced by propagating a ray from the origin through every single pixel, refracted through the dome (if it hits the dome) and into the detector. If the ray encounters either the edge of a light concentrator or a mu metal shield, the pixels are colored blue or red respectively. The user then drags the overlaid image until it lines up with a PMT and the point is saved (and marked with a white 'x' in the photo — like the one located two PMTs above the cross-hairs), these saved points are then used in a least squares fit to determine the camera parameters that would make the two images match up.

Camera	1	2	3	4	5	6	7
Yaw	0.041054	0.26040	-0.093115	-0.026698	-0.067292	0.18676	-0.07366
Pitch	-0.014955	-0.086475	0.12501	0.15042	0.059848	0.011957	-0.17960
Roll	-0.10646	-0.086475	0.86653	-3.0165	1.0044	-0.3718	-1.2675
Scale	1.6809	1.6807	1.6831	1.7815	1.7651	1.7854	1.6771
x_{offset}	15.169	-0.33155	19.622	-9.3367	14.370	-2.9745	-5.7562
y_{offset}	20.466	20.856	8.1595	17.759	8.8103	2.7927	47.270
c_1	1.0032	1.0134	0.99291	0.97317	0.97209	0.96368	1.0044
c_3	0.39752	0.41123	0.39480	0.41812	0.39488	0.40862	0.40353
c_5	0.054112	0.062903	0.048767	0.081309	0.065534	0.068253	0.066770
c_7	0.20049	0.20573	0.19313	0.20548	0.20738	0.21061	0.20069
c_9	3.8072E-3	6.0607E-3	-0.013232	-3.5001E-4	4.3533E-3	6.3846E-3	-3.1759E-3
p_1	9.6614E-4	1.7438E-4	7.0919E-4	3.2818E-4	1.1894E-3	8.6198E-4	-4.5426E-4
p_2	-2.8259E-3	-8.7652E-4	-2.3246E-3	9.4348E-5	-2.6158E-3	-4.0795E-3	-2.2238E-3

Table 5.4: The calibration parameters in use at the time of this writing for each of the seven cameras. The units for yaw, pitch and roll are in degrees, and the units for the x and y offset are in pixels. From the magnitudes of the p_1 and p_2 de-centering correction coefficients, it is clear that these introduce very small corrections to the images — the largest effect seems to be the misalignment of the CCD origin with the optical axis.

When the calibration is finished, a good quality check to perform is to visually inspect a vector plot of the residuals ($\vec{r}_{overlay} - \vec{r}_{image}$) for the calibration points. If the camera calibration were not taking an effect into account, or was doing so improperly, this could appear in the vector plot as a systematic trend in a particular direction; thus, a random distribution of residual directions is desirable. Moreover, the fact all of the cameras typically converge on a location to within 2 cm is further evidence that our parameterization is correct.

5.3 Analysis of the Vessel Volumes

The ideal description of the Borexino detector would be one in which the vessels were perfectly spherical, thus, providing a uniform quantity of non-scintillating buffer liquid for all (θ, ϕ) for a given radius inside of the inner vessel. Moreover, the vessel film is the source of the vast majority of the internal ^{232}Th and ^{238}U contamination in Borexino, thus, the radial

profile of these events would ideally be expected to follow the sum of a volume and surface term (see equations 4.2 and 4.3) — however, this is only true for a spherical vessel since non-sphericity causes the surface component to be smeared appreciably. A good quality check for the position reconstruction codes is to perform this fit to the radial distribution for various taggable species with different energies so as to check for any radial bias in the data as a function of energy. This sort of analysis is still possible even with a deformed vessel, but only if one knows the shape — then the radial fit can be done in slices in the polar angle θ where the vessel radius is approximately constant. Fortunately, the camera system can provide this information.

5.3.1 Obtaining the Vessel Shape

If one considers the physical structure of Borexino and its vessels, it is very difficult to imagine a situation in which the vessels could develop an asymmetry in the ϕ coordinate. The θ coordinate is different however: the fact that we have two different liquids with different densities on either side of the inner vessel means that there is a buoyant force which makes the inner vessel naturally want to float due to the lower density of PC+PPO with respect to PC+DMP. The vessel was initially slightly under-filled with scintillator by about 2% — the mass is known to be 278.292 metric tons (314675 liters) from integrating mass and volume flowmeters installed on the inner vessel fill pipe [108].

In addition to the under-filling, further volume reduction occurred due to cooling of the scintillator — the detector footing functions as a heat sink and during the months following the completion of the filling, a temperature gradient developed between the northern and southern parts of the detector as well as a global cooling. Because of the cooling, some additional scintillator was pumped into the top of the detector in the summer of 2007,

however, this scintillator appeared to have a much higher radon concentration (either due to emanation by storage vessels or perhaps a small air leak in the filling line) than the rest of the volume, so, further volume additions were forbidden. As a consequence of the increasing temperature gradient, much of the water that had been dissolved in the scintillator began to condense out onto the surface of the nylon vessels in the lower hemisphere — this condensation caused a haze to appear on the vessels which can be seen by comparing photos from May and October of 2007 presented in Figure 5.12. In an effort to combat this effect, in the winter of 2007, the water purification plant was used to circulate water in the water tank, heating it in the water plant and then reintroducing it to the tank, and, after a few weeks of operation, the haze disappeared. Nevertheless, the temperature (and therefore inner vessel volume) changes caused the radius of the inner vessel to fluctuate appreciably during this time.

Only for a *very* under-filled or otherwise deformed vessel can one get large distortions in the ϕ coordinate, so, if a profile of the vessel in θ can be obtained, then one can just revolve this shape about the vertical (z) axis and integrate it to obtain the enclosed volume. For a slightly deformed vessel there are small non-uniformities introduced by what we refer to as *tessellation* or *scalloping*, which is the bulging outwards of the nylon between the longitudinal ropes which circumscribe the vessel. When the inner vessel floats upwards due to the minor under-fill, the nylon pushes itself up and out between the ropes and in principle destroys the azimuthal symmetry, but unless it is very severely distorted, the effect is rather small compared to a perfectly spherical volume.

If we have a set of pictures taken with the halogen lights on to illuminate the detector, then we can in theory produce the profile shape of the vessel from this set of images. Since one camera is not sufficient to determine the three dimensional location of an image, we would need to identify objects on the vessel which are visible in two or more cameras, this is a very




Figure 5.12: The figure on the left was taken with camera 4 on May 15, 2007 on the day the filling of the vessels was completed; the photo on the right was taken on October 17, 2007 with the same camera and clearly shows the presence of the haze caused by the lower temperature which developed in the southern hemisphere as a result of heat loss to the earth. Also slightly evident is that the inner vessel volume has decreased since the initial filling and thus displaces less buffer liquid and experiences a larger buoyant force and is starting to shift upwards as is evidenced by the smaller radius near the south pole. It is important to note that these two images have not been tweak corrected (they are just the raw images downloaded from the cameras), so, they are not normalized to a reference image and so cannot be compared 1:1.

non trivial task since the vessels are quite close to being absolutely transparent. We could also attempt to use the point light sources described in Section 3.2.1. An alternative method is to exploit the fact that it is possible to identify the outer edge of the inner vessel by eye — particularly in areas near the lights from the other cameras due to the glare from the vessel outer surface. Thus, the user traverses a zoomed-in photo and whenever it is possible to identify the edge of the vessel, a set of cross-hairs are centered on the position and this point can be marked and the (x, y) pixel coordinates of this point are saved to a file. A screen-shot of the aforementioned process can be found in Figure 5.13.




Figure 5.13: Screen capture during the vessel shape determination process. The user traverses the photo — in this case it is a photo taken with camera 4 on September 1, 2008 — and centers the white cross-hairs over the outer surface of the inner vessel and saves the point (thus producing the green 'x' indicators which are also visible in the photo). The (x, y) pixel values are saved in a file and are later used in the analysis to extract the shape of the vessel. The process is being done here for the inner vessel, but can also be done for the surface of the outer vessel as well.

In order to determine the radius of the vessel at that point, we use the pixel information to construct a ray from that camera through the pixels we have identified. Once we have this ray, the algorithm is to first generate a sphere of a very small fixed radius and check to see if the ray we have constructed is tangent to this sphere, if not, the radius of the sphere is increased and the process repeats with progressively higher radii until we find the radius of the sphere which has this ray as its tangent. The minimum radius for which this condition is met is nominally the radius of vessel at that point, however, this only really holds for a

spherical, non-tessellated, vessel centered on the origin.

As mentioned on page 180, the natural tendency for an under-filled vessel in the Borexino configuration is to float upwards due to the buoyant force — the end caps are held in place nominally at $z = \pm 425\text{cm}$, but the film is otherwise free to lift. When the vessel lifts upwards and starts to distort, the tangent method is not as accurate anymore because it relies on using the intersection of a ray with a sphere centered on the *detector* origin. If the vessel floats upwards, we should use as the origin the center of the displaced vessel for computing the tangent method — this situation is depicted below in Figure 5.14. This effect would also be present for a downward shift of the vessel if it were to occur, however, this is harder to imagine and thus we only discuss the upward shift.

Early versions of the vessel shape analysis required one to iteratively determine the upward shift of the reference sphere in order to accurately correct for the offset effect. A more recent update to the code uses a reference sphere centered on the origin only to extract an approximate shape profile. This reference profile is then used in place of the reference sphere and the shape is scaled up until the tangent point is found for a particular θ coordinate.

After finding tangent points in each image, one can produce an (r, θ) plot similar to the one shown in Figure 5.15. It is clear that there is a fair amount of scatter for a given θ value, this is due to the fact that the tessellation present on a deformed vessel means that the radius reported by the tangent method can vary depending on the angle from which it is viewed. From this data set we then draw a reference curve through the data points by hand and the radius of this reference curve is read off in 5° increments in θ . In actuality, drawing the reference curve by hand is only really necessary for the first shape analysis, for subsequent analyses, the reference curve from the previous analysis is overlaid onto the camera data and the user changes the data points as necessary until a good alignment is obtained. A photo of a profile curve overlaid on shape data can be found in Figure 5.16.




Figure 5.14: Schematic view of the pitfall of not using the center of the displaced vessel as the origin of the reference sphere for calculating the vessel radius. The thick red line shows the vessel shape very similar to its most deformed state, and the green green line shows the ray from the camera (which resides on the stainless steel sphere which is indicated in blue) that is tangent to the inner vessel surface at point P_{true} . Using a reference sphere centered on the detector origin results in a ray that intersects the reference sphere at point P_{recon} which clearly has a different radius than P_{true} .

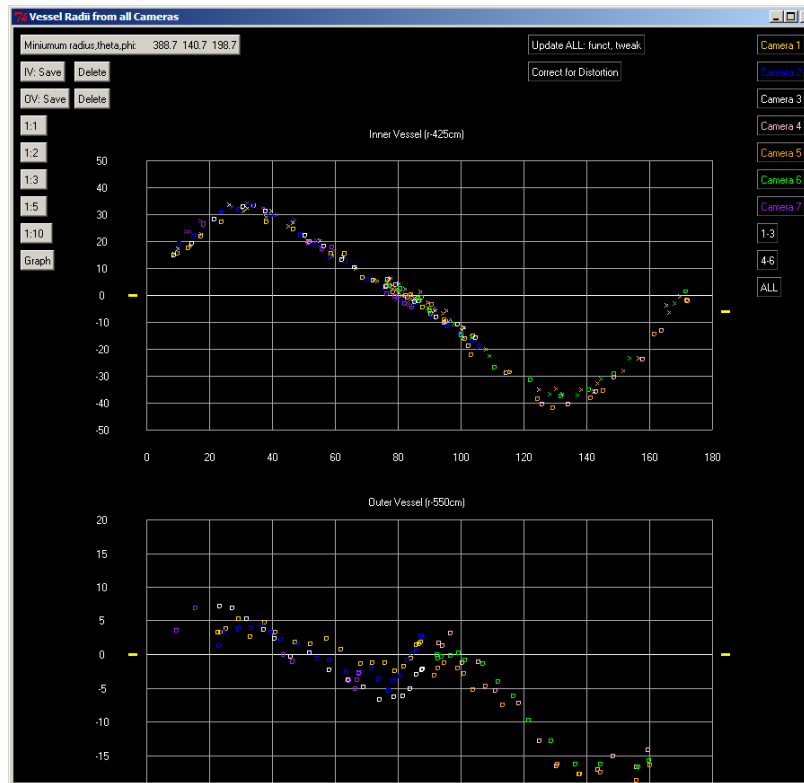


Figure 5.15: Screen-shot of the shape profile produced as a result of the tangent ray method for determining the vessel shape from a set of photos taken on September 1, 2008. The top plot shows the profile of the inner vessel and the lower plot shows the profile for the outer vessel — the x axis displays θ in degrees, and the y axis is the deviation from the nominal vessel radius of 425cm for the inner vessel and 550cm for the outer vessel. The points from each camera are distinguished by the different color indicators, and the x indicators correspond to those points on the “right” half of the vessel in each photo, and the \square for the points on the “left” half of the vessel.

Volume Calculation

Once we have determined a reference curve for the vessel shape by the aforementioned methods, we generate a table of (r, θ) pairs with $\Delta\theta = 5^\circ$ and then revolve this shape about the z axis and numerically integrate it to get the enclosed volume. Several methods are employed to do this, and the analysis is often done independently by more than one person to ensure that all results are consistent. The method employed by R.B. Vogelaar is a




Figure 5.16: Presented in red are the data points obtained from the cameras; the areas with the highest density of points are those that are visible by a majority of the cameras. In the coordinate system chosen, $\theta = 0^\circ$ is the north pole, and $\theta = 180^\circ$ is the south pole. The blue curve on top of the camera points is the user-determined shape profile obtained by drawing an interpolated curve between all of the points. A table of (r, θ) pairs is generated every 5 degrees and is used to perform the numerical integration of the vessel shape. The shape shown here is for the state of the vessel on September 17, 2008 and shows a vessel distorted by almost 50 cm at the two poles.

numerical integration using the Trapezoidal and Simpson's rules; S. Manecki fits a 13th order polynomial to the data points and integrates this shape analytically; the method employed by the author is a numerical integration using Simpson's Rule and will be discussed further in what follows.

The numerical integration is performed in the standard way for Simpson's Rule (the first and last points — at the poles — require special treatment which is discussed below) where for each triad of points — “centered” on a multiple of 10° (i.e.: $\theta = 5, 10, 15; 75, 80, 85;$ etc.) — the volume of the slice is calculated as:

$$V_n = \frac{1}{6} (\theta_{n+1} - \theta_{n-1}) \frac{\pi}{180} \{\alpha_{n-1} + 4\alpha_n + \alpha_{n+1}\} \quad (5.7)$$

Where α_n is given by equation 5.8.

$$\alpha_n = \frac{2\pi}{3} \left(\frac{r_n}{100} \right)^3 \sin \theta_n \quad (5.8)$$

The values of V_n are summed for all of the points and added to the volumes of the two pole slices which are calculated as (the subscripts np and sp denote “North Pole” and “South Pole” respectively):

$$V_{np} = \frac{2\pi}{3} \left(\frac{r(\theta = 5^\circ)}{100} \right)^3 \{1 - \cos(5^\circ)\} \quad (5.9)$$

$$V_{sp} = \frac{2\pi}{3} \left(\frac{r(\theta = 175^\circ)}{100} \right)^3 \{1 + \cos(175^\circ)\} \quad (5.10)$$

When the running sum of the values computed using equation 5.7 is added to the contributions from equations 5.9 and 5.10, we are left with our value V_{Simp} (in units of m^3) for the volume of the inner vessel using Simpson’s technique. With other numerical integration techniques it would be necessary to correct for any over/under estimation by using said technique (which can easily be accomplished by computing the volume of a 4.25 meter radius sphere numerically and then dividing that into the geometric volume to give the overall scaling), however, Simpson’s Rule is an exact method for polynomials of third degree or lower, therefore, no correction is required with this method.

For an estimate of the error in determining the volume, we assume that a human is able to accurately determine the position of the edge of the vessel to within 2 cm (one pixel is

approximately 1 cm at the center of the detector), and then send our 36 (r, θ) values into a separate Perl script which adds a number uniformly distributed on $[-2, 2]$ to each value of r and then computes the volume via Simpson's Rule. This process is repeated 10000 times and the average deviation from the calculated volume is computed and reported. The process is actually done separately for the case of random numbers distributed on $[0, 2]$ and $[-2, 0]$ since the latter has a smaller effect on the volume than the former due to the larger volume associated with a point at higher radius. For the vessel shape determined from photos taken on May 15, 2007 the volume was determined to be $322.08_{-2.256}^{+2.270}$ m³ — an error of about 1.5% for the “absolute” value; however, the value of ± 2 cm for the position determination error is a very generous estimate, and it is likely much smaller at vessel radii. The impact of any systematic error introduced by variations in how each person defines the “edge” of the vessel was found to be very small (less than a quarter of a cubic meter) by comparing data sets created independently by the author, R.B. Vogelaar, and S. Manecki. There is however an unknown systematic offset due to the tessellation of the vessel which is not possible to evaluate, since we do not know the dimensions of the deformation due to the tessellation, however, in principle one could simulate a thin film vessel and extract the dimensions in that way.

The relative error between each measurement is, however, much smaller, since one begins with overlaying the reference shape from the last shape analysis and manipulating the points which define the curve until the best match to the current data is obtained. This method helps to reduce the error on the volume since it does not require the construction of an entirely new reference shape once the points from the camera have been obtained. Moreover, for a vessel whose shape evolves slowly with time, the impact of additional tessellation or other deformity is smaller on a relative scale than on an absolute one. With these considerations in hand, the Virginia Tech group conservatively sets an estimate of one cubic meter for the

error on relative volume measurements.

The Evolution In Time Of The Inner Vessel

For a time period of approximately one year after the data taking begin in May of 2007, the calculation of the inner vessel shape and volume was useful primarily for minor comparisons or corrections for various aspects of the data analysis. Throughout much of 2008, the vessel shape was seen to make a series of more drastic changes — relative to what had been observed in 2007, and finally in September of 2008, an abnormal quantity of PPO was found in a sample of the inner buffer solution; at approximately the same time, the data analysis groups removed an $R < 4.5$ meter analysis cut, and a large amount of scintillation events were seen in the northern half of the inner buffer region. The only possible explanation for these events was that a hole had developed in the inner vessel and the contents were slowly leaking out into the buffer region. As more fluid leaked out, the shape and local stresses across the inner vessel film at the site of the leak became larger, the leak began to accelerate, and the ability to determine the vessel shape and volumes became of paramount importance for planning detector operations aimed at reducing the leak rate.

It is well known that PMT dark rates can increase by many orders of magnitude following exposure to large amounts of light — such as that required to image the inner vessel with the cameras — due to excitation of long lived phosphorescence states in the photocathode bulb and dynodes. Moreover, the Borexino collaboration has determined that the cycling of power to the Borexino PMTs dramatically increases the rate at which PMTs develop problems¹⁰, therefore, in general, the collaboration tries to avoid turning off the high voltage as much as possible; and when it is done, it remains extinguished for a period of two days. To maximize data-taking time, taking vessel-monitoring pictures was discouraged during 2007 and 2008,

¹⁰This was also observed in the CTF campaigns.

and thus, photos (and vessel shape data) were only available in May, July, and November of 2007; and January, April, July and August of 2008 *before* the leak was discovered — afterwards, pictures were taken at least once a week, and sometimes separated only by 72 hours.

The camera-determined vessel data for several different time periods is presented below in Figure 5.17 — the plot clearly shows the acceleration of the volume reduction due to the leak in the first half of 2008. A more precise time period can be determined by looking for an increase in the number of events occurring inside of the buffer volume (a plot from this analysis can be found in Figure 5.18) — performing this analysis yields a best-guess for the time period in which the leak developed of April 5-10, 2008 and at a location of $(\theta, \phi) = (30^\circ, 239^\circ)$. Several important features can be identified from Figure 5.17, namely:

- The vessel was under-inflated initially since the shape deviates from sphericity by ± 15 cm.
- During the months from May to November, the volume decreased as a result of scintillator cooling and the vessel began to shift upwards as a result of this.
- A large difference between the November 2007 and August 2008 data is due to the leak which developed in April. During this time, the temperature in Hall C was oscillating considerably as a result of the commissioning of the OPERA experiment (also mounted in Hall C)
- The vessel reached its maximum level of deformation in the spring of 2009. Several refilling campaigns performed at the beginning of June 2009 brought the vessel back towards a more spherical shape.




Figure 5.17: Vessel shape profiles from a few relevant time periods. The data series from 2007 show the effect due to cooling of the scintillator, the points from 2008 show the acceleration due to the leak, and the last two points show the effect of several re-filling campaigns.

- The fact that the shape profiles intersect the 0 line at $\theta < 90^\circ$ is another confirmation of the upwards shift of the vessel.




Figure 5.18: A plot of the events with energies above ~ 160 keV and with a z coordinate above the north pole position ($z = 4.25$ m) from early April until August of 2008. The y axis is the differential rate from the prior week, and the x axis is the time period. The departure from the straight line is a clear indication of when the leak developed.

In section 3.2.2 it was mentioned that as a result of the leak, the collaboration decided that the best recourse is to reduce the buoyant force on the leak, thus reducing the pressure differential across the location of the leak and therefore the leak rate. Various thermodynamic methods of reducing the density were considered and ultimately found to be very difficult to control as volumes the size of Borexino carry a significant amount of thermal “inertia”. The option with the best prospects was to remove some of the DMP in the buffer region by distilling it out of the scintillator — normally present at a level of 5 grams per liter, a dedicated campaign in the spring and summer of 2009 reduced the DMP concentration to a level of 3 grams per liter.

In Figure 5.19, the evolution of the volume as evaluated by the camera system is presented. The slope of the line between successive measurements yields the rate at which liquid is lost from the inner vessel; from the beginning of the leak until the first refilling at the end of 2008, the averaged leak rate is approximately $1.5 \text{ m}^3/\text{month}$. The first refilling caused an initially lower leak rate, however, the additional volume was small enough that the leak rate quickly

resumed its value prior to the refilling. A second refilling occurred in the summer of 2009 after additional scintillator had been purified — following these operations, the leak rate has reduced to $0.8 \text{ m}^3/\text{month}$. It was recently decided to reduce the concentration of DMP further to 2 grams per Liter and the process was completed in February of 2010. The vessel volume analyses performed since then (which includes only two sets of photos) is consistent with a leak rate of zero and an upper limit of $0.5 \text{ m}^3/\text{month}$.




Figure 5.19: Plot showing the time evolution of the volume of the inner vessel as evaluated by the Borexino calibration camera system. The development of the leak is clear from this plot, as are several campaigns of refilling which occurred in the winter of 2008 and the spring/summer of 2009. The leak was first discovered in the late summer of 2008 and resulted in a flurry of photos being taken to evaluate the progression of the leak. The leak rate is approximately constant at $\sim 1.5 \text{ m}^3/\text{month}$ from its development in April of 2008 until the summer of 2009. The concentration of DMP in the buffer was reduced during the summer of 2009, after which the leak rate has slowed by a factor of about three. It will be shown later that the refilling campaigns coincidentally occurred in close proximity to calibration campaigns making the de-convolution of contamination introduced by either operation quite difficult.

Searching for Evidence of Vessel Creep

As the vessel was becoming more and more deformed due to the leak and temperature fluctuations, there was significant risk that the vessel would stretch beyond its elastic limit and permanently deform — this process is known as *creep*. One can predict that the areas that experience the most stress due to vessel deformations would be in those areas where the deformation is largest, particularly in the very northern and southern regions. A detailed finite element analysis simulation of the Borexino vessels was performed by L. Cadonati[109] where various values of vessel parameters were simulated so as to explore the parameter space where the vessel can safely be operated. Figure 5.20 is a result of one such simulation, assuming a density difference of 0.5% across the vessel (which corresponds to roughly a 5°C temperature difference). From this figure, it is easy to see that the result of the density difference is an increased “bulging” strain in the northern hemisphere, and a large force trying to lift the inner vessel off of the southern endcap — both of which are very consistent with a vessel which has floated upwards. Such a complex system of films, fluids, and ropes is best analyzed by simulating all of the relevant parameters and comparing the analytically generated shape with the one determined by the cameras.

The “simulate and check” method is a very useful and valid way of performing such checks in complex systems such as this, however, one would periodically like to check the results with actual experiments. As a proxy for a direct measurement of vessel creep, we can exploit the camera system if one can identify two points on the vessel in multiple photographs and then calculate the path length between these points — if the vessel has crept, this distance will increase. To that end, the fact that the vessels have ropes running along the latitude and longitude was exploited to define the reference points as intersections of these ropes. Several different candidate points were located in multiple sets of images; attempts were also made to find candidate points in the northern and southern hemispheres — the former is under




Figure 5.20: Result of a finite element analysis of the inner vessel performed by L. Cadonati assuming a 0.5% density difference between the inner and outer vessel. The arc length is measured from the south to the north, therefore the left side of the plot corresponds to the southern endcap and the right the northern. The top represents the “hoop” stress — or the stress in the ϕ direction — and the bottom represents “meridional” stress — or stress in the θ direction. With this density difference, the vessel floats upwards, and tries to bulge outwards in the top — thus the higher hoop stress in the northern part; the tendency to float upwards means that the film is trying to lift off of the southern endcap which is verified by the large meridional stress represented.

much higher hoop stress (again see Figure 5.20) and would be more likely to exhibit creep. Figure 5.21 below shows two such candidate points defined by the intersection of meridional and circumferential ropes.

After these two points had been located in two images, the distance between the two was calculated, however, the straight line distance is not very meaningful, and what we require is the path length between these two points along the vessel. For the purposes of this calculation, we assumed that at the θ value of the two candidate points, the vessel had a constant radius, and then the arc length was simply the length of the circular segment




Figure 5.21: Photo taken from camera 6 on May 15, 2007 showing two of the reference points used in searching for evidence of creep in the nylon film. The intersection of the horizontal and vertical ropes located underneath of the red dot surrounded by the white ‘x’ is one of the points, the other is the intersection of the two ropes to the lower right of the white cross-hairs. These two points were located in later photographs and the distance between the points calculated.

connecting the two points. For a severely deformed vessel, this assumption is on less solid ground since for points in the northern hemisphere, the tessellation or scalloping of the vessel will cause the radius to no longer be single valued on the z plane at the particular value of θ — this effect will be weaker in the southern hemisphere where the tessellation is much less.

In the end, after evaluating several candidate points between photos taken on May 15, 2007 and those taken on September 2, 2008 (near its point of maximum deformation), the maximum arc length difference observed was only 0.53 cm for the points in the northern hemisphere where one would expect to see the difference the most. On average, the points

are separated by an arc length of $\sim 70\text{cm}$, therefore, there is effectively no evidence for vessel creep by this method since this difference of $< 1\%$ can likely be entirely attributed to the fact that we are not actually finding points on the vessels, but rather, on the ropes which themselves are loosely secured to the vessel.

A similar check was performed by picking candidate pairs for a constant value of ϕ , but at different values of θ . In this case, the radius is not single valued, but follows the vessel shape profile and the path length must be determined by summing over small steps. The results showed negligible path length differences as well. With these data in hand, it does not appear that there is any evidence for creep in the inner vessel film.

Chapter 6

A Calibration Source Deployment System For Borexino

In Chapter 4, some motivating factors for performing a source calibration were presented; and then in Chapter 5, a method for finding the precise location of a source inside of the detector was introduced. The focus of this chapter is centered around how one actually places a source inside of the detector, specifically the hardware involved and the method used.

In Section 4.4.3, some of the more relevant mechanical design constraints were presented and chief among them was the very narrow opening through which the source must be introduced — the five meter long pipe which leads from the gate valve at the top of the detector to the entrance to the inner vessel has an inner diameter of only 4 inches (10.2 cm). The insertion system for the KamLAND experiment, which was heavily based on the Borexino design, uses a long rod which is inserted essentially vertically and then tilted off-axis[110]. The Virginia Tech component of the Borexino collaboration investigated several types of neutrally buoyant flexible rods for deploying the source, as well as several other ideas.

Ultimately, the idea which was adopted was to model the system after the human arm and have it comprised of two straight sections with a hinge capable of bending the lower portion up to 90° to provide the off-axis maneuvering. Such a system could be made quite narrow and inserted straight down on-axis, and once the hinge exited the access pipe, the lower portion could be bent upwards. Of course, the challenge was to do this invasive process without ruining the extreme radiopurity and optical qualities of the scintillator. In the sections that follow, the hardware for deploying and maneuvering the sources¹, as well as the computer interface for controlling the system will be presented. The Borexino system is based on the CTF deployment system (which was also designed by the Virginia Tech group). The system is easiest to describe if we consider separately the components which actually enter the detector, and those that exist for cleanliness or operational reasons.

6.1 In-Detector Hardware

In its most basic form, only four components actually enter the detector: the rods themselves, the Teflon tether tube which is used to pull the source off-axis, a special coupler to hold the source, and the source itself. Since these components will be in direct contact with the inner-vessel scintillator for the duration of any calibration campaign, they must be *extremely* clean. The Borexino operational group adopted a cleanliness standard of Class 30 (as defined by Military Standard 1246C [111]) for any component in contact with the inner-vessel scintillator. All of the rods, the tether, and associated hardware were cleaned thoroughly in the Borexino cleaning module (see Section 3.2.4) until particulate counting of the rinse water showed a cleanliness level at, or below, class 30. Since the sources are so much smaller in size, they were cleaned in a separate process using an ultrasonic bath of

¹The sources themselves will be described in the following chapter.

detergent and rinse water.

6.1.1 Insertion Rods

The source deployment system design adopted requires a set of identical mating rods that can be coupled together in a clean-room atop the detector and lowered into the desired depth before pulling the hinge off-axis. To reduce unwanted stress on seals and other components, and to minimize operator fatigue, the rods are designed to be essentially neutrally buoyant in the Borexino scintillator. They are in fact designed to be very slightly negatively buoyant so that the weight of rods below the hinge produce a torque that will make it naturally want to fall back to vertical.

The rods are hollow pieces of 1.5” diameter 316L stainless steel tubing (0.065” wall thickness) that have custom couplers orbital-welded onto the ends. The couplers are tongue and groove type to ensure the rotational alignment (in ϕ) of all of the rods — a photo can be seen in Figure 6.1. Both male and female couplers have two bolt holes which are partially threaded — the male coupler bolt hole is threaded below the countersink, the female coupler bolt hole is a simple clearance hole for half of its depth, and threaded for the other half. The bolts themselves are partially threaded, with the rest of the shank turned down to a smaller radius. This seemingly odd configuration is actually a safety measure — after the threaded portion of the bolt passes through the first set of threads, the reduced shank means the bolt is no longer tightly constrained and is free to wobble slightly inside until it is lined up with the threaded portion of the female hole; while this makes it slightly more difficult to fasten two rods together, it is *very* difficult for a bolt to back out completely by itself (there are two such bolts on each coupler, making it even more unlikely for trouble to occur). Finally, a very thin sleeve with two bolt holes fits around the couplers; after the bolts are attached, the

sleeve is rotated so that the bolt holes no longer line up with the bolts, making it effectively impossible for two rods to decouple.



Figure 6.1: A spare set of rod coupler components.

When assembled and welded, the tube and couplers form a rod which is exactly one meter long — measured from the midpoint of the bolts at each end. The female coupler has one additional bolt hole in the outer face which was used to provide purge gas during the welding process, as well as an access hole through which to place pieces of metal rod to tune the weight of the rods. An initial testing of the system in water was foreseen, and it was important to design the system so that the buoyancy of the rods can be adjusted for different fluids. For use in pseudocumene, each rod contains a 90cm long piece of 3.5mm outer diameter stainless steel rod for buoyancy matching purposes. In order to add even more weight to aid in returning the hinge to the Z axis for very light sources, one rod² contains two such pieces. After the additional weight was added to the rods, the bolt holes were sealed by replacing the bolt and potting the countersunk area around it with an epoxy resin³; the rods were then leak checked.

For identification purposes, each of the 18 rods is stamped with a number. The rods were

²The rod with the extra weight is rod number 13.

³Epotek 353ND epoxy — this is the same one used on the Borexino PMT calibration fibers.

all tested to determine how best to assemble them; in theory they are all identical, but small variations in rods, bolts, sleeves, etc., dictated that a definitive assembly plan was warranted to determine which rod should “go next” when deploying a source. The order of the coupling of the rods during most calibrations was: Source, 2, 3, 4, Hinge, 5, 7, 8, 9, 11, 14, 1, 15, 18, 16. Rods 2, 3, and 4 each couple well to both the source holder and the hinge, so, if a shorter portion of the arm below the hinge is required, one or more of them can be removed without causing problems. Straying from this predefined sequence is not advised, since smooth mechanical coupling of the rods will not be absolutely ensured otherwise.

Since all of the rods are electropolished, care must be taken when choosing the fasteners for mating two rods because using the same material for a fastener and the piece(s) to be fastened is almost guaranteed to result in seizing of the fastener. Early tests of the system utilized both silver and nickel plated bolts, however, the plating process increased the overall dimensions of the threads⁴, resulting in a tighter fit and caused minor abrasion of the plated material. Eventually, we adopted the use of bolts which had been sputtered with molybdenum disulfide (MoS₂). The MoS₂ coating was much thinner than the others we had tried, and did not cause the fit problems we experienced with other coatings.

Hinge rod

Figure 6.2 is a CAD drawing of the hinge rod holding a source — the tether is not shown, but would connect to the coupler just above the source and proceeds all the way up into the access pipe and to the clean-room atop the detector. In order to prevent the possibility of the hinge being bent upwards past horizontal and becoming jammed, the design of the coupler prohibits motion past 90°. The shoulder bolt about which the hinge pivots is nickel

⁴The manufacturer verified that the bolts were undersized to begin with, however, the results varied greatly among batches.

plated in order to reduce friction, and contains several wave washers stacked underneath of its head to take up any angular play.




Figure 6.2: CAD drawing of the hinge rod, source, and source coupler. The hinge pieces are coupled together by, and pivot around, a shoulder bolt. Not shown is the Teflon tether tube used to pull the assembly off-axis, however, it would attach to the coupler just above the source if it were drawn here.

Cleaning

After fabrication and welding, the rods were electropolished in order to produce a very smooth surface, unlikely to trap particulate matter. After being shipped to Italy, they were loaded into a custom built chamber (essentially a long piece of stainless steel pipe with conflat flanges and high throughput valves at each end), four at a time, and connected to the Borexino cleaning module. The rods were exposed to a hot, turbulent flow of Micro 90 critical cleaning detergent for six or more hours, and then rinsed with a continuous stream of deionized water (18.3 M Ω resistance) for several hours until measurements of the conductivity of the waste water failed to indicate the presence of residual detergent. After rinsing, the rods were exposed to a hot EDTA chelating solution for several hours to remove

additional metal ions from the surface of the components (of particular concern are ions of polonium, lead, bismuth, etc. from the uranium and thorium decay chains). After rinsing for several more hours, a sample of the waste water was collected and passed through a filter for particulate counting. When each batch of four rods was verified to have met at least Class 30, the rinsing was ceased and they were left to dry under a nitrogen flow for several hours. The entire chamber was then taken into a class 100 clean-room (the lower CTF clean-room) where they were removed from the chamber and bagged in class 10 plastic bags for transport up to the class 100 clean-room atop the detector: clean-room 4 (CR4).

The bolts were cleaned in a separate process in an ultrasonic bath of Micro 90 detergent, rinsed, particulate counted, and bagged in a clean room.

6.1.2 Tether & Fiber Optic

So far, we have described the “bones” of the arm model of the deployment system (the rods) but, we have left out the “tendon” that allows us to bend the arm (the tether). If we attach one end of the tether to any portion of the rods below the hinge, and pull on the other end, the arm will obviously move upwards, and return to vertical if tension is released. In the early versions of the CTF insertion system, the rods were moved off-axis by pulling on a tether made of monofilament fishing line, however, the CTF lacked a precise source location system like we have in Borexino, so, it was not necessary to provide a light source. For locating the source in Borexino, a fiber optic cable transmits red laser light to a diffuser bulb mounted close to the source, and the tether tube acts as a conduit for this fiber optic.

Tether Tube

The tether tube is a 100 foot (30 meters) long continuous piece of $\frac{1}{4}$ " (6.4 mm) outer diameter, $\frac{3}{16}$ " (4.8 mm) inner diameter, ultra clear PFA Teflon tubing. Obviously, the end of the tether that goes into the detector must be closed to prevent the tube from flooding, but must also be transparent so that the diffuser light can get to the cameras. The termination for the detector end is a blind piece of $\frac{3}{16}$ " outer diameter Pyrex tubing (like a test tube) fused to a piece of thin wall stainless steel tubing — these are commonly available from scientific glassware vendors. The stainless steel tubing is several inches long and is pushed into the detector end of the tether tube — since the outer diameter of the stainless steel tubing is the same as the inner diameter of the tether tubing, it lends itself well to a force fit. In order to further secure this termination, thin stainless steel wire is used to tightly wrap the tether tube in two places. A space in the source coupler (section 6.1.3) traps the wire-wrapped portions so the whole system is constrained from slippage. A photo of the detector termination with the illuminated diffuser can be seen in Figure 6.3.

A piece of tubing 100 feet long is rather unwieldy if not properly stored; the presence of the fiber optic cable inside makes it even more important to restrain it from kinking or twisting. A tether storage drum mounts on the side of the glovebox and the tether is stored in alternating clockwise and counterclockwise loops to avoid any twisting of the fiber — see Figure 6.4.

Fiber Optic Cable

The fiber optic cable inside of the tether tube was chosen to preserve the ability of using it to transmit both red laser light for illumination of the diffuser for the cameras, as well as ultraviolet light from a laser for studying light propagation and scattering in the inner




Figure 6.3: The “detector” end of the tether showing its transparent termination. Starting from the left, we see the bottom edge of the source coupler (Section 6.1.3), the Teflon tether tube with the thin wall stainless steel tubing inserted into it, and the transition from stainless steel to Pyrex. Inside of the Pyrex tubing, it is possible (albeit difficult) to see the bare fiber optic cable, and finally the white Teflon diffuser at the far right end of the figure. The tether tube is very tightly wrapped with thin stainless steel wire and trapped inside of the source coupler to ensure it cannot slip out.

vessel. Important characteristics of the fiber optic cable can be found in Table 6.1.

The camera system gives the precise location of the source required for an accurate calibration, however, it is highly desirable for the system operators to know approximately how to put the source there — in other words, how many rods above and below the hinge are required, and how much of the tether tube has to be retracted in order to bend the arm to that location. The calibration software provides such a prediction for a given source position, and setting the rod orientation is a simple matter of counting the rods already inserted and measuring the amount of the rod still sticking up out of the detector — but we still require a method to determine how much of the tether has been retracted. To that end, a 30 meter long tape measure with 1 cm graduations has been attached to the fiber optic and can be read through the transparent Teflon tube — this white tape is visible in Figure 6.4 attached to the black fiber optic (each meter is marked in bright yellow). An electronic method is also available which employs a rotary encoder that records the quantity of the tether that passes




Figure 6.4: The drum mounted on the side of the glovebox which holds the 100 foot long Teflon tether tube. Since the tether contains a fiber optic which should not be twisted, it is stored in alternating clockwise and counter-clockwise loops which removes any net twisting. Two ears on each side of the drum keep the tether contained, but can be pulled to the side to allow the operators to take or return a length of tubing from the drum. The initial version employed two augers which fed the tubing into and out of the drum, however, operation of the augers proved to be difficult through several layers of glove material.

through a set of compression wheels, however, it has not yet been used in a calibration as the tape measure has proved to be more than adequate.

6.1.3 Source Coupler

The final component of the “wet” end of the deployment system (save for the sources themselves) is the coupler which is used to bring together the rods, tether, and the source. The coupler also carries the majority of the responsibility for providing the weight to bring the hinge back to vertical when tension on the tether is released, so, it is intentionally designed to be somewhat bulky and heavy. A photo of the assembled source coupler during a calibration can be seen in Figure 6.6, and a CAD drawing showing the space between the cover plate and body can be seen in Figure 6.5.

Manufacturer & Model	CeramOptec Optran UV
Numerical Aperture	0.22 ± 0.02
Core	400 microns; Pure fused silica
Cladding	424 microns; Doped silica
Coating	450 microns; Polyimide
Jacket	3mm; PVC with Kevlar strength members
Connectors	(1) SMA-905, (1) bare polished
Length	30 meters
Part Number	SMA1C/UV400/424/450P/OPVC-3.0mm/30.0M/VT

Table 6.1: Characteristics for the fiber optic cable used in the Borexino insertion system. This fiber normally just carries red laser light to a diffuser in the detector which is used by the camera system to precisely locate its position. However, it is also desirable to have a fiber which can carry UV light to provide an in-situ source of simulated scintillation light to study light propagation, scattering, and PMT timing calibration parameters.

Since we are using such a narrow walled Teflon tube, we have to avoid sharp bends whenever possible since a tube of these dimensions quickly buckles or kinks if it is forced to form to a small radius. To that end, the source coupler contains an arced inlay which is designed to guide the tether smoothly into the coupler without inducing such a bend. The inlay eventually transitions to a straight section which has a small cutout of a slightly larger diameter to accommodate and trap the wire-wrapped portion of the tether. The very bottom of the coupler contains a set of holes where the source support wires are anchored.

At the top of the coupler is our standard female coupler (present on all of the rods), however, the collar is different. Since we must always keep the pressure differential between the cross and glovebox, the rods pass through a sliding seal which maintains this pressure. With the standard collar, it is possible that an operator could accidentally pull the source coupler up into the rod seal — since it is not round, the pressure would be lost and the source could be damaged in passing through the seal. To guard against this possibility, the collar on the source coupler is thicker than the normal ones so that it cannot be drawn up into the seal.

Since we again have the problem with fastening two electropolished surfaces together, the




Figure 6.5: Assembly drawing of the source coupler for Borexino. The top end (right hand side in this photo) is identical to the female end of the rods, however, the sleeve that goes around it is made of very thick Teflon — this is to prevent it from being completely pulled through the rod seal and causing a pressure loss. The curved inlay in the two rectangular pieces provides a smooth path for the tether to enter the holder to prevent kinking. The inlay eventually turns to a straight section with a cut out — this cut out houses the wire-wrapped portion of the tether and secures it from being withdrawn. A set of holes at the bottom of the coupler are used to secure a set of wires which hold the source below it. A photo of the actual assembled coupler can be seen in Figure 6.6.

bolts used to fasten the cover plate of the source holder onto the body are nickel plated stainless steel.

6.2 Glovebox, Cross, and Associated Hardware

In the previous section we described the “wet” components of the system — the ones that actually go into the detector — and in this section the “dry” portion of the system is presented. The two main components of the dry side of the system are the glovebox and the six-way cross.

Maintaining the radiopurity and optical qualities of the scintillator means reducing, as much as possible, the exposure to air and any particulate matter. As a first order of protection, the entire system is housed in a class 100 clean room (CR4). Additionally, all manipulation




Figure 6.6: The source coupler during a source change in the January 2009 calibration campaign. The tether, fiber optic, and rule can be seen entering the source holder from the left. The thick Teflon collar is used to prevent the possibility of extracting the source completely through the pressure seal that the rods pass through. At the lower right of the figure is the AmBe neutron source in its holder, ready to be attached to the wires at the upper right corner of the photo.

of the system is performed inside of a Low Argon and Krypton Nitrogen (LAKN) filled glove box. The glove box also acts as storage for any components which go inside of the detector to minimize radon daughter plate out.

Section 4.4.3 discussed the requirement that the source deployment system be able to maintain the pressure in the detector (set by the height of columns of scintillator in three head tanks) during the operation of the system. The pressure difference relative to the clean room

is approximately 95 mbar⁵ — this would be over 2000 pounds of force on one wall of our glovebox! Thus, an intermediate volume which can accommodate this pressure is placed between the glovebox and the inner-vessel: a spherical-body six-way cross used for vacuum systems.

6.2.1 Glovebox

Commercially available gloveboxes were not able to meet the space, cleanliness, and functionality requirements for the source deployment system. The glovebox needs to be at least one meter tall to accommodate the insertion rods, however, its overall height is limited by the height of the ceiling in CR4. An additional constraint is provided by the fact that one must also have enough space to couple two rods together *above* the rod seal.

Top Hat

In actuality, only the portion of the glovebox directly above the gate valve needs to be tall enough to allow the coupling of two rods — it was therefore decided to modify the ceiling of CR4 to include a recess to accommodate the “top hat” on the glovebox which provides this extra space. The top hat is constructed from a piece of standard schedule-40 pipe and a cap welded to a flange. The top hat is attached to the top glovebox panel with large bolts and the seal is provided by a sheet of Buna-N rubber. It was originally envisioned that a rod could be prevented from falling down by being attached to a piece of string looped through an eye bolt in the top hat, however, this is not currently in use, but the eye bolt still exists.

⁵This value does fluctuate depending on the temperature of the underground laboratory, but usually lies in the region (90, 100). The head tanks contain a set of weirs set at a height corresponding to 900 mm — any increase past this value causes the excess pseudocumene to flow down the weirs into a recovery tank, D201. The pressure differentials across the nylon vessels is very small — only a few mm of water — the majority of the ~ 100 mbar pressure lies on the stainless steel sphere.

Collar

The height constraint also forbids having the floor of the glovebox mount to the top of the cross — the available height from the top of the cross to the ceiling is less than one meter. Since the top of the cross sticks up several inches above the rest of its body, we can lower the bottom panel of the glovebox by this distance to gain the extra height required. The only issue with this solution is that now the seal between the glovebox and cross has to come from somewhere else — this is accomplished by placing a collar on the bottom panel which the top flange sits inside of and seals to via an o-ring on its periphery. The bottom of the collar contains a shallow groove for a piece of square Viton o-ring cord stock to provide a seal to the lower glovebox panel.

Glovebox Panels

The glovebox panels are made from $\frac{1}{16}$ " (1.6mm) thick sheets of stainless steel and were bent into a shallow pan / cookie sheet shape so that the flat portion of each panel measures 24" x 43" (61cm x 109cm). The tabs on the bent over portions contain bolt holes every 4" (10cm) and the top and bottom panels contain identically spaced holes on the flat section for fastening to the side panels. The back of all the holes contain self clinching 5mm threaded nuts (PEM nuts) for easier assembly of the panels and acrylic windows.

Glovebox Frame

Mounting of the glovebox to the rest of the system is done via a frame attached to the underside of the bottom panel. The frame is constructed from 1.5" (3.8cm) square tubing that has been welded together, and is fastened to the bottom panel via nuts welded to the underside of it. The frame is itself attached to the cross via two large U-bolts which fit

around the tubular arms of the six-way cross.

A CAD drawing of the glovebox panels, top hat, collar, and frame can be seen in Figure 6.7.




Figure 6.7: The glovebox panels, top hat, collar, and frame. The space constraints imposed by the height of the clean-room require that extra height be obtained by two means: adding a top hat (golden color) which extends into a recess in the ceiling; and lowering the bottom panel so that it does not seal to the top of the cross, but to the periphery of a collar (blue) installed on the base of the bottom panel. The entire glovebox is attached to a frame made of square tubing — slightly visible through the collar. All of the panels are identically sized at 24" x 43" (61cm x 109cm). Not shown is the hole in the panel on the right for mounting the tether drum.

Acrylic Panels

The front and back panels on the glovebox are made from sheets of static-dissipative acrylic, 43" x 44.75" (109cm x 114cm) with bolt holes drilled to match the holes on the glovebox panels. Each acrylic panel has holes for three glove ports, two of them are at a height of 12" (30cm) above the bottom of the panel, with the third port at a height of 14.5" (37cm) above the outer port and centered horizontally with it. The third port is on the tether drum side

(south east facing) of the glovebox to make the operation of the tether drum easier, however, it turned out to be operable from the two lower ports. The two “extra” gloves do provide a very valuable function as buffer volumes to keep the glovebox pressure from exceeding the high and low set points when operators insert or withdraw their arms from the other gloves.

Gloves

A rather complicated situation is that of the gloves — we obviously need a glove material which is compatible with pseudocumene, and also one that has low emanation and diffusion rates of radon. A disparagingly small amount of compatibility information is available for glove materials and pseudocumene. After an exhaustive search and testing period, the candidate materials were determined to be:

- Viton — Commonly used for o-rings. However, radon emanation and diffusion rates for Viton have been measured to be prohibitively high. Moreover, a material as stiff as Viton makes for a rather tiring material to work with since it resists movement.
- Nitrile — While Nitrile is just a name for any organic compound with a $C\equiv N$ group (the “Nitrile” group), *some* types of nitrile rubbers used to make gloves are compatible with pseudocumene. We tested many types of nitrile gloves for both chemical compatibility and Radon emanation — unfortunately, most exhibited problems with either pseudocumene compatibility or emanation of radon.
- Norfoil / SilverShield / 4H / New Barrier ⁶ — These are thin laminate film materials commonly used in hazardous material cleanup and are resistant to a very wide array of materials. The problem with gloves of this type is the lack of dexterity one has when

⁶SilverShield is the name as marketed by North Hand Protection; 4H by Safety4; and New Barrier by Ansell Edmont

wearing them. Since it is not a rubber, but rather two pieces of plastic sealed together in a hand shape, it does not conform to ones hand in any way, and tend to be quite slippery.

Ultimately, the decision was made to abandon the attempt to use one single glove material since it proved quite difficult to find one that was compatible with pseudocumene and had low emanation and diffusion rates for radon. Butyl rubber is widely available in glovebox-length gloves, is very inexpensive, and has one of the lowest radon emanation rates of all the elastomer materials tested by the collaboration; it however has poor resistance to pseudocumene. By placing a smaller — not glovebox-length — and thinner, glove of a PC compatible material we were able to find a middle ground between radon emanation and chemical compatibility.

The choice was made to outfit each glove port with a standard Butyl rubber glovebox glove and then to cover each⁷ hand portion with a nitrile glove (Certiclean Class 10 Nitrile — by Kimberly Clark). These gloves were measured by H. Simgen and G. Zuzel of MPIK-Heidelberg to have a Rn emanation rate of approximately 2 mBq per glove, which is a tolerable level. They have the advantage of being disposable and their cleanliness level is already certified to be Class 10. The downside is that it can be quite difficult to insert ones hands into the gloves, since first one has to put their hands into the butyl gloves, and then put another set of gloves on over top of them.

A photo of the assembled glovebox in the clean room can be found in Figure 6.8.

⁷In practice, the two upper, rarely used, glove ports do not normally have a nitrile glove on them.




Figure 6.8: The glovebox assembled in CR4 as seen from the southwest corner. The tether drum can be seen mounted on the wall of the glovebox and the gas box to the lower left of the photo.

6.2.2 Cross

While the glovebox supplies a slightly positive pressure atmosphere in which to work, it cannot maintain the 90 mbar pressure required to regulate the fluid level in the head tanks, therefore, a separate volume was required for this purpose. This extra volume is provided by a six-way cross — the type traditionally used in ultra high vacuum systems. The top of the gate valve to which the cross mounts is a 4.5” standard conflat flange, so, MDC vacuum company produced a custom six-way cross with five 10” conflat flanges (all four horizontal sides plus the top), and one 4.5” rotate-able conflat flange on the bottom. A CAD view of the cross can be seen in Figure 6.9.




Figure 6.9: Exploded view of the six-way vacuum cross used as the pressure vessel to maintain the ~ 90 mbar pressure head when the gate valve is open. The cross is produced from a spherical 12" diameter body with 8" tube sections and 10" conflat flanges on five sides, and a 4.5" flange on the bottom. Two adjacent sides have blank flanges (one of them has several tube fittings for various process control instruments), and the other two adjacent sides have 3.375" quartz viewports with removable acrylic covers for protection against breakage.

Of the four horizontal flanges on the cross, we require that two of them be removed frequently in order to change the source, however, the other two will remain static. The two static flanges are simple 10" blank flanges, one of them contains several NPT fittings (sealed with anaerobic PTFE pipe thread sealant) for various process control equipment. Before closing the gate valve, one would like to look at the source (under low level light of course) to verify that it is clear of the closure path to avoid breaking it, thus, the two removable flanges were designed with quartz viewports, acrylic covers for the viewports to guard against breakage of the quartz, and light-blocking covers.

Since the cross also experiences vacuum during the purging process, it must be able to maintain positive pressures up to ~ 100 mbar, and vacuum down to 100 torr. To ensure this, the cross has been He leak checked after installation and is pressure tested prior to every insertion of a source.

Cleanliness

In theory, the cross should never be exposed to scintillator, however, it is a region *directly* connected to the inner vessel when the gate valve is open, so, it is subjected to the same cleanliness constraints as the rest of the in-detector components. After production, the cross was bead-blasted by the manufacturer and then sent for electropolishing. Upon arriving in Italy, several external welds were performed which were followed by pickling and passivation of the weld site. Before being installed in the clean-room, the flanges were attached and the entire cross connected to the cleaning module and underwent the same cleaning treatment described in Section 6.1.1.

Top Flange

The rods, tether, and other items must pass through the top flange on their way into the detector. The periphery of the top flange is also where the seal between the glovebox and the cross occurs — recall from section 6.2.1 that the bottom panel was lowered and the top flange now seals to the inside of the collar. Figure 6.10 shows a section view of the top flange assembly with the rod, tether, and sweep arm seals atop it as well as a photo of the actual components during assembly of the system. Installation of the top flange is a non trivial matter and must be done very carefully in order to compress the peripheral o-ring evenly without having it roll out of its groove.

One problem with having the seal to the glovebox occur at the periphery of the top flange is that the bolt holes used to fasten the top flange to the cross are essentially 24 leaks to the clean-room. A solution is provided by using soft nylon washers under the head of each bolt — these soft washers form an adequate seal between the top flange and the underside of the bolts, thus eliminating the leak path.




Figure 6.10: The top flange of the cross with the three sliding seals through which the sweep arm, rods, and tether pass. Note the groove in the periphery of the flange — this is for the large o-ring that seals to the collar on the glovebox bottom panel. The actual installation of the top flange in the collar is shown on the right during assembly and installation — when this photo was taken, the cross had already been thoroughly cleaned, so, the plugs were placed in the seals to avoid contaminating the cross. The incoming purge gas for the rod seal can be seen at the center of the photo. Since the photo was taken, all of the compression caps for the three seals have been replaced by delrin after experiencing problems with the threads on some of the seals.

Rod Seal

In order to maintain the pressure difference between the cross and glovebox while still being able to add or subtract rods from the detector, we created a purged sliding seal (center seal in Figure 6.10). The voids in between two coupled rods (in the threads and heads of the bolts, between the mating faces of the couplers, etc.) can contain trapped gasses that we would like to be able to purge, thus, the seal contains a honeycombed inner sleeve and purge ports which allow purge gas to flow throughout the region between the two o-rings in the seal. The internal components of the seal are made of delrin and teflon to reduce friction and avoid scratching the rods (which would generate particulates) — the honeycombed sleeve is made from electropolished stainless steel and is bored to a larger diameter to avoid contacting the rod. A photo of the internal components of the rod seal assembled onto an insertion rod during testing of the system can be seen in Figure 6.11

Compression of the internal o-rings is provided by tightening three screws that bind together a cap and threaded sleeve which mate to the rod seal body. Originally, the compression was provided by turning a fully threaded cap, however, we experienced problems with the threads seizing during testing and it was replaced with one made from delrin.

An issue arises when one considers the potential for abrasion of the o-rings as the rod couplers pass through the seal. While the leading and trailing edges of the sleeves (as well as the bolt cutout on the sleeves) were smoothed as much as possible, they still have a tendency to bite into the o-rings slightly. The hinge also poses a similar problem when the rotating coupler passes through the rod seal. To minimize the degree with which the o-rings were compromised, we employ teflon covered viton o-rings in the rod seal — the slipperiness of the teflon makes a drastic difference in how smoothly the couplers transition through the seal. Moreover, the o-rings are pre-stretched over a mandrel for several days prior to being installed




Figure 6.11: Internal components of the rod seal assembled onto an insertion rod during testing of the system. From left to right: the compression cap, first compressor ring, sealing o-ring, second compressor ring, honeycombed purge sleeve, third compressor ring, and the bottom sealing ring. The purge sleeve allows LAKN to reach voids between the rods when the couplers pass between the sealing rings. Note that the o-rings shown in this picture are not the ones installed in the system — see body text for more information.

in the seal. We do not experience as severe a problem when the rods are being withdrawn from the detector when they are already coated in pseudocumene (the lubricating properties of which are very helpful in this aspect), thus, a bottle of ultra-pure scintillator is kept in the glovebox to wipe down the rods and tether in order to wet the o-rings.

Tube Seal

In analogy with the rod seal, one must also be able to insert and remove the tether from the detector without losing pressure in the cross. To that end, the tube seal (to the right of the rod seal in Figure 6.10) is basically a smaller version of the rod seal except it is not purged because the tube is smooth and has no voids that can trap any gases. The body and inserts for the tube seal are made from stainless steel and the cap is made from delrin. To allow for

larger diameter assemblies to be attached to the bottom of the tether and still be removed, the hole in the top flange underneath the tube seal is much larger than the $\frac{1}{4}$ " tube that passes through it — otherwise, one would have to cut the tether to remove it.

Sweep Arm & Seal

Before closing the gate valve after withdrawing a source from the detector, one would like to verify that the source is in fact up into the cross, above the gate valve. Since the high voltage to the PMTs is on when this occurs, one cannot just remove a light blocking cover and look through the viewport, so, a tactile method is provided by a sweeper arm. The arm is a straight piece of $\frac{1}{16}$ " stainless steel rod attached to a $\frac{1}{4}$ " stainless steel rod that goes up through this seal to a handle. The sweep arm resides just above the plate portion of the bottom tubulation of the cross — slightly over 2" above the gate valve. Free movement of this handle as it is turned between two fiducial marks indicates that the gate valve is unobstructed and can safely be closed. Figure 6.12 shows the inside of the cross during a source change and highlights the sweep arm as well as the liquid level monitor (described further in Appendix A.1).

Details on the process control of the source deployment system can be found in Appendix A.




Figure 6.12: Photo of the inside of the cross prior to inserting the neutron source in June of 2009. Indicated with yellow arrows are the sweeper arm (used to ensure that the gate valve closure path is clear), and the float of the liquid level monitor (see Section A.1).

Chapter 7

Calibration Source Production

The final ingredient required for a calibration is of course the source itself. In the one year period from October of 2008 to August of 2009, four separate calibration campaigns were performed in the Borexino detector using a variety of sources, both radioactive and optical. The focus of this chapter is the design, fabrication, and quality control of the sources produced at Virginia Tech for these calibrations.

7.1 Calibration Source Design

To calibrate the detector with a radioactive source, the scintillation light produced by the source should mimic that of the bulk scintillator as closely as possible. This requires that careful thought be given to the source containment vessels. First and foremost, if the scintillation light that the source produces will come from within the vial itself (this is the case for alpha sources since the alphas cannot penetrate the walls), the scintillator inside of the vial must be as identical as possible to the bulk scintillator. To be useful for an energy calibration, we have the additional constraint, as explained in Section 4.4.3, that the entire system

block as little light as possible — given the relative dimensions of the rods to the detector, the biggest potential for light blocking comes from the source coupler. Other issues include how to seal a vial so that it will not leak, how to attach it to the rest of the insertion system, how to clean the vials, and what modifications are necessary to obtain the best performance from different types of sources (optical, α , β , etc.).

7.1.1 Containment Vials

The emission spectrum of PPO was shown in Figure 3.22 in Chapter 3 and has its peak at about 365nm. Thus, whatever container is used to hold the source must transmit ultraviolet light throughout as much of the PPO emission spectrum as possible. After having gone through so much trouble to achieve the purity levels present in the detector, the source vial design should ensure that the contents do not leak into the detector — as a second level of defense, we should only use isotopes which can be removed by the purification system, otherwise, we risk contaminating the detector irreparably. Finally, the vials should attach to the deployment system in a very secure, and preferably spring-loaded, method.

Vial Design

The vials used for the Borexino calibrations are based on those used in the later CTF calibration campaigns — a one inch diameter quartz sphere. The primary differences from the CTF design are in the sealing and attachment to the deployment system — for the CTF sources, an externally-threaded Teflon ring was placed around the neck of the vial and threaded into a stainless steel spring (see Figure 7.1). One of the main differences between the CTF and Borexino is the lack of a location system in the CTF; having a large, essentially opaque, spring mounted above the source would block the cameras from seeing the light

source — so, the Borexino sources are designed to be mounted via four thin stainless steel wires which provide an almost completely unobstructed view of the source.

Perhaps the biggest difficulty with the CTF sources was creating a hermetic seal — the CTF versions employed sources with necks that transitioned from glass to metal which were then crimped closed and usually potted with an epoxy resin. The glass - metal transition was particularly weak and it was very difficult to crimp the neck closed without snapping the source at the joint; the Borexino sources were designed to obviate the need for crimping.




Figure 7.1: Design of the source attachment scheme for the calibration of CTF 3 which occurred in the spring of 2005. A split Teflon ring compressed around the source neck and was threaded into a spring to provide a spring-loaded mounting method.

The design and dimensions of the containment vials for the calibration sources employed in the four calibrations of Borexino are presented in Figure 7.2.




Figure 7.2: Dimensions of the calibration source vials (in inches) used in the four calibrations of Borexino which occurred between October, 2008 and August, 2009. The vials are one inch diameter quartz spheres, with a neck that transitions to Pyrex due to its lower melting temperature. The flare on the neck is present to provide a surface that will prevent the entire neck from slipping out of the coupler. The dimensions given are only approximate as it is essentially impossible to make a uniformly shaped piece of glassware out of quartz.

It is important to note that the working temperature for quartz is almost 900°F hotter than for borosilicate glass, and is thus *very* difficult to blow into a uniform shape, so, the dimensions of the source vial given in Figure 7.2 are only approximate and thickness variations on the order of $\frac{1}{32}$ " have been measured on broken vials. The total height of the source vial is intentionally not given in the figure — the vials are produced with a much longer neck for connecting to the source loading station (see Section 7.2.1), however, this extra length is removed when the vials are sealed (the sealing site is indicated in the figure). Slight variations in the distance from the base of the neck to the outward flare are also common because this area contains a graded transition from quartz to borosilicate (via several short pieces of progressively “softer” glasses) and this distance is not precisely repeatable. For this reason, every time a source is attached to the coupler, we record the precise distance from the diffuser to the center of the source, thus, the length of the grade is not an absolutely critical dimension.

Sealing

The experience with the CTF sources and the difficulties encountered in sealing them led us to search for additional methods of sealing the vials. A method we had long considered, but deemed to be potentially dangerous, was to seal the glass of the vial itself; however, the failure rate of crimp seals gave occasion to reconsider the option. Several of the features presented in Figure 7.2 were added to make the sealing of the vials safer and easier:

- During the initial testing of the sealing of the vials, the transition to Pyrex was removed, since it was thought not to be completely necessary (the glass-to-metal transitions previously used only existed as metal to Pyrex, so, the transition was necessary) and it was thought that the small diameter of the tube to be melted would not justify the extra complication with the grade. After experiencing problems with the repeatability

of all-quartz vials, we decided to reintroduce the grade to Pyrex, and to seal the Pyrex portion of the tubing.

- The other major design change was the constriction in the neck. By constricting the wall of the tubing and performing the sealing at that location, the seal can be made with less time in the flame and resulted in a more repeatable sealing method.

The production of the sources will be described further in Section 7.2 and a main theme of it is to prevent the scintillator from coming into contact with oxygen during the preparation. To that end, the vials are attached, via an Ultra-Torr¹ vacuum fitting, to a $\frac{1}{4}$ " bellows valve via a short stub of stainless steel tubing — this valve is then used to connect the vial to the rest of the loading system. After the vials are filled and ready to be sealed (they are usually placed in the quality control detector described in Section 7.3.1 before being deemed ready to be sealed), the vial is immersed in a dewar filled with liquid nitrogen and frozen, the valve is then opened and the source vial is evacuated with a rotary vane pump for several minutes.

The frozen, evacuated source vial is removed from the system and the vial transported to the glass blower to perform the sealing. The presence of vacuum in the vial ensures that no gases are present which could expand and rupture the vial when heat is applied. Keeping the source frozen is also a safety precaution since the expansion of pseudocumene vapors, or the pseudocumene itself, could cause the vial to explode. When the glass blower is ready to seal the vial, it is removed from the dewar, and the sealing site is placed in a lean flame to slowly warm the constriction, then the full flame is applied while the source is rotated — once the glass becomes soft, the atmospheric pressure pushes the constriction closed as the glass blower pulls the two ends away from each other. An example of a sealed source is shown in Figure 7.3.

¹Ultra Torr is the trade name for the Swagelok brand of fittings which seal on the outer diameter of a tube via an o-ring crushed into the base of the fitting by a special ferrule.




Figure 7.3: The ^{14}C - ^{222}Rn source used in the first off-axis calibration in January 2009, shown shortly after sealing of the source. After filling, the vials are frozen in a liquid nitrogen bath, evacuated, and then transported to the glass blower for sealing while still frozen.

Even with all of the precautions taken to try to ensure the efficacy of the sealing method, the success rate is still around 80%. A few failures have been traced back to valves which have started to fail for one reason or another. Another area where many sources fail during sealing is in the Ultra-Torr fitting used to couple the Pyrex portion of the source to the bellows valve — if the glass spends too much time in the flame, the temperature of the neck around which the o-ring in the Ultra Torr fitting sits can get hot enough to cause the o-ring seal to fail. Still other times, the problem is just with the glass itself — micro-fractures in the glass which were not removed during the post fabrication annealing process spread to larger cracks when placed in the flame. Recently, we have added a tee with a Bordon tube vacuum gauge between the vial and the valve so that it is possible to monitor for loss of vacuum during preparation or sealing.

Attachment To The Source Coupler

Affixing the source vials to the source coupler must be done in a way to maximize the visibility of the diffuser mounted just above the vial, so, the CTF method of mounting the sources (see Figure 7.1) is not sufficient for Borexino. Two pieces of stainless steel wire (0.040", 304, spring temper) are bent into an elaborate "U" shape and the two upper ends hook through holes in the source coupler. The wires extend essentially straight down for about 9 cm and then bend in sharply towards the source vial before bending back down to vertical with a final upward lip at the lowest spot. The ~ 1 cm long area above the upward bend is where the source is actually attached by wrapping it tightly with a very fine stainless steel wire — the two sharp bends at either end of the wire-wrapping prevent the source from slipping out in either direction. For added security, the sources are secured by two independent wire-wrappings, one just above the other. The strength of this type of source attachment was vetted by successfully attaching a two kilogram mass to the source coupler in this fashion. Photos of the source attachment are presented in Figure 7.4.

Vial Preparation

Before loading the vials with anything, we would like them to be as clean as possible so that the only activity present is that which we introduce ourselves. When we receive the vials from the glass blower, many times they have a few soiled areas — perhaps where organic deposits had resided when heat was applied — and they could contain a variety of other impurities.




Figure 7.4: Top: Photo of a calibration source (Rn without scintillator) inside of the cross before being deployed. The entire length of the suspension wires can be seen in this manner. Bottom: A closeup view of a source being wire-wrapped onto the suspension wires. From this photo it is easy to see how the two sharp bends (coupled with the flare in the source neck) prevent the source from slipping out in either direction.

The interior of the vials are first partially filled with acetone and shaken by hand, emptied, and then filled completely with acetone and placed in a warm ultrasonic bath for ten minutes. The acetone is rinsed out and the process is repeated with isopropanol, and then rinsed with deionized water for several minutes. The vials are then filled with a 5% nitric acid solution and placed in a water bath at 80°C for two hours. At the conclusion of the acid bath, the vials are rinsed for ten minutes with a continuous stream of deionized water and then blown dry with UHP nitrogen for several hours; once dry, the vials are sealed in plastic bags until needed.

After the vials are filled and sealed, they are wiped down with acetone and then carefully wrapped and shipped to Italy by certified technicians with the Environmental Health and Safety Services office at Virginia Tech. When they are received in Italy, the sources are cleaned in a warm ultrasonic bath of Micro 90 — or a similar critical cleaning detergent — to clean the exterior of the source. After the ultrasonic bath, they are rinsed with deionized water for several minutes, dried and double bagged for transportation to the clean room underground.

7.1.2 Design Variations For Different Particle Types

The design of the sources presented in the previous section is applicable to most types of sources, however, some variations in the design are necessary. The primary motivation for modifications to the basic design are to ensure that the particles emitted by the source deposit all of their energy in scintillator — assuming that the scintillator in the vial is identical to that in the bulk of the detector, we just demand that it be in scintillator and not in the wall of the vial. Radiation which interacts partially with the vial and partially with the scintillator leads to a distortion in the energy spectrum of the source — this renders a source practically

useless for an energy calibration. Another reason to modify the source design is if we wish to filter the output of the source — for instance, to remove low energy gammas emitted in alpha decays, or filtering out charged particles if we only wish to investigate gammas.

β Sources

If we were to simply dissolve a β -emitting isotope in some scintillator (which itself is not a negligible task) the spherical geometry dictates that a large fraction of the betas will be produced near the wall of the vial. The longer range of the betas means that there is a higher probability that it will interact with the vial wall and deposit some of its energy into it (without producing scintillation light) on its way out of the vial — this will cause a distortion in the observed energy spectrum thus making it more difficult to use for an energy calibration.

To attempt to quantify this effect, a Geant4 simulation was written where betas are generated uniformly inside of the source vial (taken to be a uniform sphere with the dimensions given in the previous section) and the amount of energy deposited in the scintillator was tallied throughout the life of each generated particle. As expected, more of the energy is deposited in the scintillator for betas of lower energies (i.e. less of a chance of encountering the wall). To begin, only monoenergetic betas with energies of (200, 500, 750, 1000) keV were generated, and the average fraction of their energy deposited in the scintillator was (98.52, 94.08, 90.45, 88.27)% respectively.

For a monoenergetic beta source (such as a conversion electron source), this effect can be simulated and corrected for (with some additional systematic error), however, a “standard” beta source emitting a spectrum of energies is a more complicated matter. To see the spectral distortion caused, the simulation was run again, but the monoenergetic betas were replaced

with a spectrum of betas distributed according to the β decay energy distribution function (equation 7.1) with a Q value of 687 keV (to simulate the dominant decay mode of ^{85}Kr). The spectrum of betas that were simulated (blue) along with the spectrum of the energy which was deposited in the scintillator (red) are shown in Figure 7.5.

$$\frac{dN}{dE'} = C_0 (E' - Q)^2 (E' + m_e c^2) \sqrt{E' (E' + 2m_e c^2)} \quad (7.1)$$




Figure 7.5: Simulation of ^{85}Kr β events in a source vial to examine the spectral distortion caused by interactions with the vial wall. The blue curve shows the energy spectrum of the betas that were actually simulated, and the red curve shows the amount of energy deposited in the scintillator — the difference is due to interactions in the wall of the vial which do not produce scintillation light.

The spectral distortion can be reduced if we “fix” the betas to originate from approximately the center of the source vial. A method for doing this is to deposit the isotope inside of a micropipette, and simply place this pipette inside of the source vial filled with scintillator — if the pipette is thin enough, then a negligible amount of energy is deposited into the wall of the pipette and the vast majority of the betas will interact completely within the source vial. Development of this method began several years ago at Virginia Tech, but has never been used with a source in Borexino. The pipettes that were used for development have an outer

diameter of $100\mu\text{m}$ with a wall thickness of $10\mu\text{m}$ and are made entirely of quartz. A smaller capillary tube is used to deposit a small aliquot of beta emitting isotope into the pipette, the ends of which are then flame sealed to an overall length of one inch — we attempt to place the isotope solution at the center of the capillary tube $\pm 5\text{mm}$.

To check the efficacy of this method the same Geant4 simulation described above was run again after adding the pipette to the geometry and constraining the particles to be emitted from the inner surface of this pipette. Again, the simulation recorded the fraction of the total energy deposited in the scintillator for monoenergetic betas with energies of (200, 500, 750, 1000)keV — now the fraction of the energy deposited in the scintillator has increased to (97.24, 99.24, 99.51, 99.64)% respectively. The correlation between energy deposition and energy is now reversed, although this is to be expected since now the primary “absorber” is the pipette instead of the vial wall and this volume affects the lower energies more.

Thus, we see that by constraining the isotope to lie on the inside of a very thin microcapillary, we can minimize the spectral distortion caused by non-scintillating interactions with the vial walls. This method also has the additional benefit that we need not have the isotope in a form that can be dissolved in pseudocumene — which is rarely an easy task for non-gaseous isotopes.

γ Sources

Since all gammas produced in nuclear decays are accompanied by the emission of another type of radiation (or at least in close succession with another form of radiation), obtaining a clean sample of gammas requires some modification to the generic source design. In most instances (the exceptions being the decay of metastable states), the gamma and alpha or

beta produce scintillation light at the same time, so, the detector sees the calorimetric sum of the two energy depositions. To filter out the charged particles emitted with the gammas, we can exploit the longer range of the gamma and ask that scintillation light only be produced by the scintillator in the bulk of the detector — in other words, we remove the scintillator from the vial. The alphas and betas produced in these decays will range out in the water or the vial walls and the gammas are the only particles able to exit the vial and produce scintillation light.

Two methods have been employed to produce gamma sources in Borexino — the first is to obtain the isotope in an aqueous solution (usually as an acid) and dilute it with deionized water and simply pipette it into the vial. The water of course does not scintillate when the charged particles are produced, but the gammas are free to exit the vial. The second method employed was to not put any liquid in the vial at all — this is possible if the gamma-producing isotopes are parents of a gaseous substance. In the calibration campaign in January 2009, a vial was loaded with a small amount of radon gas, and the rich gamma spectra of ^{214}Bi and, to a lesser extent, ^{214}Pb produced many gamma lines for calibration purposes. It should be noted that placing a material with a different index of refraction inside of the vial does produce a slight lensing effect for the light that passes through the source on its way to the PMTs — this may skew the position reconstruction slightly as it alters the time of flight for photons, however, these sources are really only useful for energy calibration, thus lensing is acceptable as long as the light collection is still the same.

Neutron Sources

The collaboration purchased an americium beryllium (AmBe) neutron source² and have used it for calibrating the CTF detector in the past, and it is a very desirable source to use in Borexino. The Am and Be metals are centered in a mold of plastic to make handling safer, the overall diameter of which is 18mm. However, for the AmBe source, we have the inverse problem of that for the gamma sources discussed in the previous section, as the alpha decay of ²⁴¹Am also produces a 59.5 keV x-ray with a 36% branching ratio[80]. For the source owned by the collaboration, this would produce 180 kBq of activity from this x-ray, therefore, we must shield it[112].

If we wish to reduce the activity of this x-ray to the level of 0.1 Bq, we require a reduction in the intensity by a factor of 5.5×10^7 . The usual choice for shielding x-rays is lead due to its high density Z value (it also does not have a large neutron capture cross section). The mass absorption coefficient for lead at an energy of 60 keV is 5.021 ([113]), standard exponential attenuation of gamma rays dictates a thickness of:

$$x = \frac{\ln\left(\frac{I_0}{I}\right)}{\left(\frac{\mu}{\rho}\right)\rho} = 2.1\text{mm} \quad (7.2)$$

So, to be safe, the entire neutron source was surrounded by 3 mm of lead on all sides. Obviously, the neutron source cannot be placed inside one of our source vials, so a special holder was fabricated for the neutron source — the body is white delrin, and the lead shielding is epoxied into the interior faces of the holder; stainless steel hardware and a Viton o-ring seal the holder. Prior to being used in Borexino, the neutron source holder was placed inside of a pressure vessel and exposed to several bars of Helium for thirty minutes — after

²A core of americium is surrounded by beryllium — some of the alphas from the decay of americium-241 are absorbed by the beryllium-9 to form ¹²C and a neutron.

removing it from the pressure vessel, the sealing area was inspected with a Helium leak detector in sniffer mode to verify that there were no leaks. The design and dimension of the holder can be seen in Figure 7.6.




Figure 7.6: Dimensioned CAD drawing of the neutron source holder (units are mm). The body is made from white delrin, and contains a 3mm thick lead cylinder for shielding the 60 keV X-rays emitted in the decay of ^{241}Am . Drawing courtesy of S. Parmeggiano.

It is worthwhile to note that the shape of the neck of the neutron source holder as depicted in Figure 7.6 is very similar to that of the quartz source vials — a similar diameter and the same outward flare so it will mate without any modifications to the system. Photos of the neutron source coupler can be seen in Figure 7.7.




Figure 7.7: Top: The neutron source being placed inside of the special neutron source holder. Bottom: The source being tied onto the suspension wires — note that the source coupler has the same basic shape as the quartz vials.

Optical Sources

Finally, we have the possibility to use the fiber optic inside of the tether to channel light from an ultraviolet laser mounted on top of the detector. The only modification necessary here is that the diffuser we use for locating the sources with the cameras is not sufficiently uniform for studying light propagation with the PMTs. The diffuser (visible in Figure 6.3 in Chapter 6) was designed to be narrow in order to minimize light blocking. However, it does not scatter well towards the two poles of the detector. Since we do not have the light blocking restraint when using an optical source (because we do not have a radioactive source producing scintillation light), we can attach a larger diffuser around the existing one to attempt to scatter the light uniformly into 4π .

The lasers which currently exist on top of the detector have wavelengths of 394, 355 and 266 nm. The first laser is used primarily for the PMT calibration (its wavelength is very close to the maximum quantum efficiency of the PMTs), it is also within the PPO emission band. The second of these, 355nm, lies in the region of overlap between the PPO emission and absorption curves — the PC is largely transparent at this wavelength and so it is useful for studying absorption and re-emission (i.e. inelastic scattering) on the PPO molecules. Finally, the wavelength of 266nm is strongly absorbed by pseudocumene, thus, it is useful to simulate excitation of the scintillator into the singlet levels.

In the calibrations performed so far, the emphasis has been on checking the timing calibration of the photomultipliers, so, we used the 394nm laser since the DAQ is already set to trigger on the “Sync Out” pulse from this laser (it is used during each normal DAQ run to monitor the energy response of each PMT). Two laser calibrations have been performed — the first at the conclusion of the July 2009 calibration campaign, and the other in September of 2009. The first deployment of the laser source was done with an inadequate diffuser, and

the second was performed with a larger, more uniform diffuser. In theory, the uniformity of the light is not absolutely necessary; but in order for the calibration to be useful, we wish to achieve only one photoelectron per PMT (see Section 3.3.2 for a description of why the PDF for multiple photoelectron events is different than for single photoelectron events). Achieving one photoelectron per PMT is possible, even for a non-uniform diffuser, by turning down the intensity of the laser — of course, this increases the time required for the test — provided that *some* light reaches each PMT. Since any contamination introduced during a calibration is a function of the amount of time taken to complete the calibration, we used a larger diffuser for the second calibration, and did reduce the intensity further — the result was a much more useful calibration. The diffusers were simply delrin spheres with a hole to slip over the existing diffuser.

“Pin” Sources

An idea inspired by the pin sources used in each rod coupling of the KamLAND 4π calibration system [110] is to place a radioactive source inside of the shoulder bolt in the hinge of the insertion system. The location of the diffuser as determined by the camera system combined with the knowledge of how many rods have been inserted completely determines the position of the shoulder bolt at the center of the hinge. If one can place a source *inside* of the shoulder bolt, we can use this information to look for radial biases in the data as one alters the number of rods below the hinge.

We collaborated with QSA Global Inc. to design such a bolt, however, it was never purchased due to time constraints. An appropriately sized bolt would be drilled out, and an ion exchange bead coated in a ^{137}Cs solution (roughly 100 Bq of activity) would be placed inside of this recess and then welded shut and leak checked. In theory, it would also be possible to do this for the bolts used to couple the insertion rods, and we could have sources placed one

meter apart.

7.2 Preparation of the Sources

Over the course of the Borexino and CTF experiments, several groups have attempted to produce unquenched radon sources by several methods ([27], [114]), but all have exhibited some degree of quenching. Many of the previous methods relied on trapping radon obtained from air or water inside of an activated charcoal filter, and then heating the charcoal and allowing the radon to emanate out and dissolve in a sample of scintillator which was then loaded into a vial. A version of this method was employed by the author to create the sources for the CTF 3 calibration in March of 2005, however, they were found to be quenched and the charcoal trap based method was abandoned completely. The remainder of this section will focus on describing the method by which the unquenched radon sources were created, as well as the ^{14}C calibration sources.

7.2.1 Source Loading Station

In all prior attempts to produce calibration sources for Borexino or CTF, the scintillator used to produce the sources was quenched to begin with, or the scintillator in the source was exposed to oxygen (or another quenching agent, although oxygen is the most common) during the course of loading. To combat this problem, a dedicated source loading station was constructed to facilitate the loading and filling of the source vials in a manner that prevented oxygen contamination. Earlier versions of such a system, based on the charcoal trap idea, were constructed with an abundance of tube fittings and inexpensive ball or plug valves — the new system was designed to minimize these components as much as possible.

A piping and instrumentation diagram for the source loading station is presented in Figure 7.8.

Design

To reduce the number of tube fittings, the lines in the system were orbital-welded together and then pickled and passivated to remove any scale or potential for corrosion. All of the valves used are either stainless steel bellows valves (valves V1-3 and V8) or needle valves (valves V4-7) where regulation of a gas stream is required. Three rotameters (R1-R3) are used in various locations to set and monitor flow rates through the radon source, or fritted bubbler in the scintillator storage flask. A rotary vane pump provides the vacuum required for fabricating the sources. The pump is connected through a liquid nitrogen vapor trap to prevent pseudocumene or its vapors from entering it, and the pressure is monitored by means of a stainless steel Bourdon tube pressure gauge. Scintillator is stored in two storage flasks, with the second one present solely to act as a one-way valve (the exhaust from the first flask is immersed in the scintillator of the second flask), the first one contains a fritted glass bubbler to sparge the scintillator and contains the tube from which scintillator is drawn off. The exhaust of the entire system is safely fed to an elephant trunk take-off from a fume hood. An annotated view of the system is presented in Figure 7.9

System Purity

As leak tightness of the system is of paramount importance to produce unquenched sources, we invested in high quality bellows valves (Swagelok BN series for V1,2,3 and H series for V8 due to its smaller size). After orbital-welding and surface treatment of the weld site (pickling and passivation), all components were assembled and He leak checked down to

10^{-10} mbar·l/s. After the initial assembly, but before producing any sources for a calibration campaign, the system was flushed with a hot Citranox detergent solution to remove any PPO residue and other impurities — jumpers were placed across the radon source, vacuum pump, and other areas where detergent was not permitted to enter (although all fittings on these components were removed, cleaned, and reassembled).

All of the stainless steel and copper lines in the system are heated with a heat tape to allow baking of the system under vacuum after such a cleaning. Several thermocouples installed on the system monitored the temperature during baking to ensure it did not exceed 200°C (the maximum temperature allowed for the bellows valves). After baking, all of the valves were tested to ensure they functioned properly and had not been damaged by the baking. The teflon tubing sections in the system were flushed with acetone, then isopropanol, and finally rinsed with deionized water and then blown out with high purity nitrogen.

Radon Source

To create a source with radon in it, we obviously require a source of radon itself, preferably from a sealed source for safety reasons. Such sources are commercially available for calibrating radon detectors, and two such sources were procured by our collaborators at Princeton University for testing a radon filter for the cleanroom in which the vessels were constructed[58]. We obtained one of these sources after it had served its purpose in Princeton. The source is made by Pylon Electronics and is model Rn-1025 with an equilibrium activity of 21.4 kBq and provides a continuous activity of 2692 mBq/minute[115]. The source contains powdered radium salts between two fine glass filters encased in a sealed aluminum housing with two needle valves and can be operated under vacuum or at flow rates of up to 10 LPM at 3 bar of pressure. This source, visible in Figure 7.9, provides a very safe, highly accurate method of loading a source vial with radon without risking exposure to oxygen.

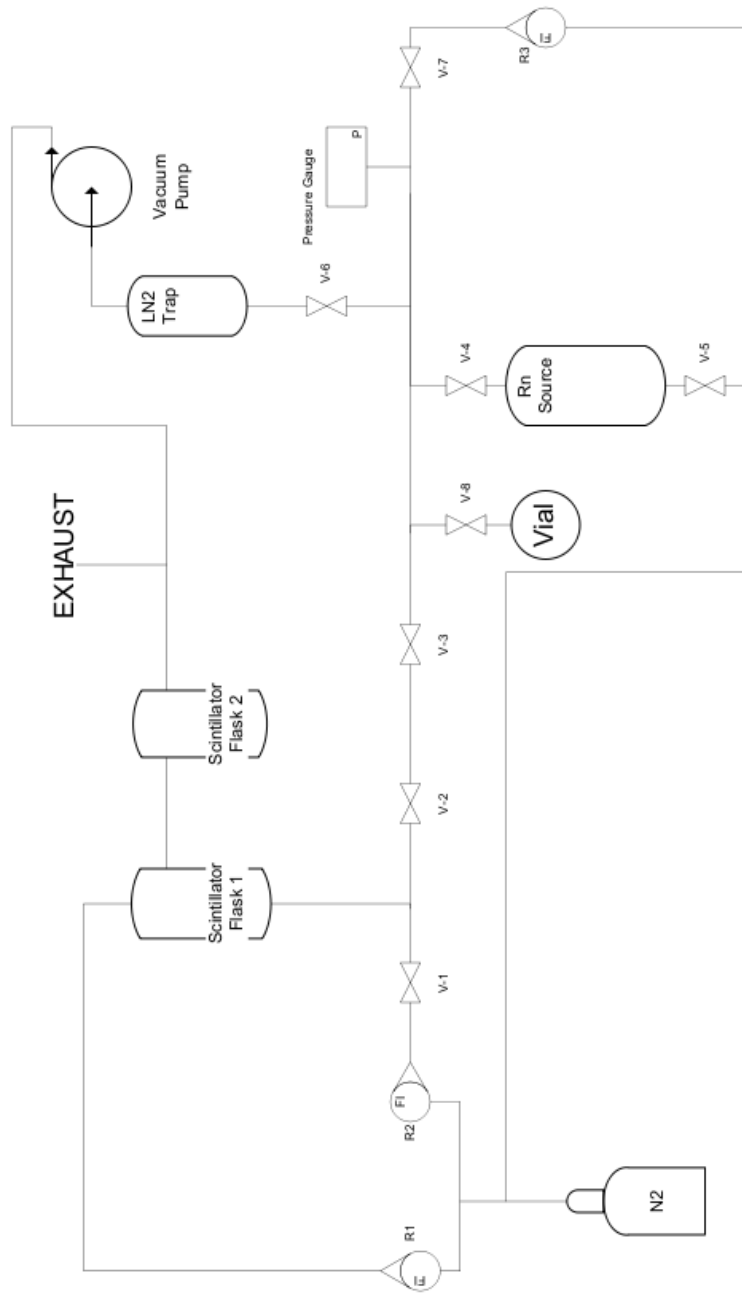


Figure 7.8: Piping and instrumentation diagram of the source loading station. See the text for more details about the components and its operation.




Figure 7.9: Annotated photo of the source loading station at Virginia Tech. The photo was taken prior to producing the calibration sources for the January 2009 off-axis calibrations. See the text for more details about each component and the operation of the system.

7.2.2 Production of ^{14}C Sources

While the Borexino trigger condition sets a hardware limit of approximately 50 keV for the experiment, the presence of the ^{14}C events produced by the carbon in the scintillator itself require a software threshold for analysis to be set around 200 keV. In theory, this still allows a region of ~ 220 keV to observe the edge of the pp neutrino spectrum, however, the pile-up spectrum of ^{14}C almost perfectly mimics the pp neutrino edge as can be seen from the simulated spectrum in Figure 8.18. However, it may be possible to place a limit on the pp neutrino flux with Borexino if we know the ^{14}C spectral shape precisely. Unfortunately, this shape is somewhat poorly constrained due to the presence of several shape factors in the equation for the beta decay of ^{14}C (see [116] for details on the ^{14}C measurement performed in the CTF). If we can insert an intense, localized source of ^{14}C into Borexino, we may be able to make significant strides towards determining both the spectral shape of ^{14}C , and setting a limit on the pp flux with Borexino.

Source Configuration and Activity

From the simulations of the source vials (results presented in Section 7.1.2) we see that the response of lower energy betas is actually better if we do not constrain them in a pipette in the center of the vial; this is a very important result as it directly impacts how we might create such a source. The same simulation was performed using betas with a Q value of 156 keV³ distributed uniformly inside of the spherical source vial and the total energy deposited in the system was plotted along with the total energy deposited into the scintillator — the results are presented in Figure 7.10. The average amount of energy deposited into the scintillator was 99.78%, thus, preparing a source with the ^{14}C uniformly distributed in the

³The distribution from which the beta energies were taken did not include the Coulomb correction or the correction for the spectral shape factor, but these should introduce a relatively minor effect overall.

vial is more than sufficient.

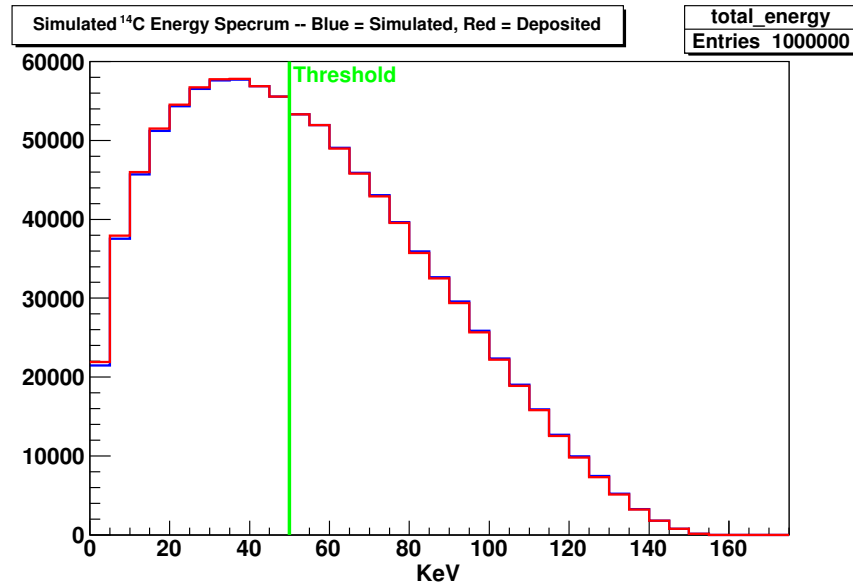


Figure 7.10: Results of a simulation of ¹⁴C betas uniformly distributed inside of the spherical source vial. The blue curve is the energy of the betas that were produced, and the red curve shows the amount of energy that the betas deposited into the scintillator — the difference is the amount deposited into the wall of the vial. It is clear from this plot that the spectral distortion is effectively zero for energies above the DAQ threshold (approximately 50 keV). It is this result that allows us to prepare a ¹⁴C calibration source where the betas can be uniformly distributed inside of the source as opposed to suspended in a microcapillary tube at the center of the vial in the manner described in Section 7.1.2.

Another issue to consider for the ¹⁴C sources is that the vast majority of the spectrum lies below the detector threshold, thus, to have an appreciable rate in the detector, we require a source with a comparatively larger activity than for radon or other isotopes. In order to not cause stability problems with the Borexino DAQ, the source rate above the trigger should be kept around 100 Bq; this gives a rate of a few Bq in the region between 60 and 80 npe — this region is usually referred to as the ¹⁴C window in data analysis. Working from the energy distribution for ¹⁴C beta decay, we see that for a rate of ~ 100 Bq in the range above 50 keV, we require a source with a total activity of 300 – 400 Bq.

For the on-axis calibration and first off-axis calibration, we decided to keep the ¹⁴C activity

smaller by a factor of about 5 and we produced sources with a total activity in the range of 60 Bq.

Obtaining The ^{14}C

To create a calibration source with such an activity, we must have a concentrated form of ^{14}C which can be dissolved uniformly in the scintillator that we place in the source vial. Enriched forms of various carbon-containing gases, such as CO_2 , methane, etc., are commercially available. However, they primarily come as very small amounts of highly enriched gas and it is therefore quite difficult to accurately meter out the amount necessary to produce the 300 Bq source. Many liquids containing enriched amounts of ^{14}C are available as they are routinely used as tracers in biological and pharmacological applications. However, it is difficult to dissolve liquids in a non-polar aromatic hydrocarbon such as pseudocumene. A common method of calibrating liquid scintillation detectors is to purchase commercially available ^{14}C -Toluene — toluene in which the fraction of ^{14}C to ^{12}C has been increased by enrichment — with a specified activity and tune the response of the liquid scintillation detector to this solution. Fortunately, pseudocumene (trimethylbenzene) is closely related to toluene (methylbenzene) and two such similar molecules tend to be quite miscible. In fact, the ^{14}C -Toluene solution is such a commonly used calibration solution that it is available for a very small fraction of the cost of other, less common, ^{14}C enriched liquids.

For the sources used in Borexino, we purchased a 10 mL bottle of ^{14}C -Toluene with a specific activity of 6.6 kBq/mL from American Radiolabeled Chemicals (catalog number ARC0282). From a solution with this concentration, we can create a calibration source with a ^{14}C activity of 330 Bq by simply using a pipette to transfer 50 μL of this solution into the source vial. To produce the lower activity (60 Bq) sources, we first prepared an intermediate solution by diluting 1 mL of the stock solution in 22 mL of pure pseudocumene — we then used a

pipette to transfer 200 μL of this solution into a source vial.

Filling the Vial

Due to safety restrictions imposed by the radiation safety office at Virginia Tech, we performed the pipetting of the concentrated solution inside of a disposable glove-bag — effectively a glove-box without a frame. A clean 50 μL pipette was placed inside of the glove bag along with a cleaned source vial, bellows valve and Ultra-Torr fitting, and the ^{14}C -Toluene solution, the glove-bag was then sealed and filled with Nitrogen — see Figure 7.11 for photos of the glove-bag and the items placed inside of it. Once the bag had been pressurized, a needle valve on an exhaust line was opened slightly to allow continuous purging of the bag to flush out the oxygen and any toluene vapors. After the bag had been purged for an hour, the 50 μL of concentrated ^{14}C -Toluene solution was pipetted into the bottom of the clean source vial and then attached to the Ultra-Torr fitting and bellows valve to seal it off from the atmosphere. After allowing the bag to purge for several more minutes, it was opened and the source vial retrieved and held in a vertical orientation while the source loading station (see Section 7.2.1) was prepared.

Before filling the vial, the piece of Teflon tubing that connects the scintillator storage flask to the back of V2 (see Figure 7.8 for piping and instrumentation definitions) had to be completely filled with scintillator, otherwise, the vial will not fill completely. A spare source vial and valve were connected to the source loading station and evacuated for several minutes. After evacuation, valve V2 was opened to allow the vacuum in the piping and vial to pull scintillator from the storage flask into the loading station. This process was repeated several times (emptying the vial after each time) since the volume of the vial is less than the volume of the tubing we are trying to fill with scintillator. Once the vial is able to be completely filled with one attempt, we know that we have removed all traces of gas from the system

and we can proceed to fill the actual source vial.

After connecting the vial containing the 50 μL of ^{14}C solution to the loading station, the vial was evacuated and then valve V2 was opened to allow it to fill with scintillator. Care must be taken to not allow too much scintillator to enter the vial, otherwise it can fill the neck past the sealing point. After the vial has been filled to the desired level, the valve on the source vial (V8) is closed and the vial temporarily removed from the system — the reason being that there is still some scintillator in the tubing of the source loading station. This excess scintillator is removed with a high pressure stream of nitrogen. Before re-connecting the vial to the loading station, the side of the valve in contact with the atmosphere is thoroughly dried to remove any scintillator present above the bellows seal⁴.

The only remaining task now is to degas the solution since it may contain dissolved radon, oxygen, or other contaminants, and we also require that it be under vacuum before sealing it (see Section 7.1.1). The entire vial is immersed in a liquid nitrogen bath for several minutes until the solution is completely frozen, the vial is then evacuated for several minutes while frozen. The valve on the source is then closed and the liquid nitrogen bath removed and the solution is allowed to warm — as it warms, the dissolved gas rapidly exits the solution due to the dependence of solubility on temperature and pressure. The process is then repeated three or four more times, until only a small number of gas bubbles are seen exiting the vial when it is warmed. After the final degassing, the vial is frozen once more and evacuated before being transported to the glass blower for sealing.

⁴This “excess” scintillator very often causes the vial to over fill when the valve is opened again — moreover, it has been exposed to oxygen and its introduction would quench the scintillator in the vial.



Figure 7.11: Top: The disposable glove-bag inside of which the ^{14}C -Toluene solution was loaded into the vials. Bottom: The bottle containing the concentrated ^{14}C -Toluene solution, the $100\ \mu\text{L}$ pipette, and the source vial.

7.2.3 Production of Rn Sources

The “workhorse” calibration source for the CTF and Borexino has been a radon source because one can obtain α , β , and γ events covering energies up to ~ 3 MeV with just one source. It is thus quite easy to study the effect that different particle types and energies have on position reconstruction — and vice versa: to study what effect the position of the source has on the energy detected. Moreover, it can be used to obtain reference pulse shapes for α and β events to be used in the calculation of the Gatti parameter (see equation 3.33 and the accompanying discussion). To obtain as much information as possible from the calibration campaign, we typically also add ~ 30 Bq of ^{14}C to the radon sources via the method described in Section 7.2.2.

Loading Method

Previous production methods all relied on an intermediate volume where the radon was collected and then transferred to the source vial — the VT source loading station was designed to obviate the need for the intermediate volume and to build up the radon in the source vial itself and then simply fill it with scintillator.

To trap the radon in the vial, we exploit the fact that radon gas liquefies at -61°C , and solidifies at -71°C , while the boiling point of liquid nitrogen is -196° . If we immerse the vial in a liquid nitrogen bath and establish a flow of radon loaded nitrogen through this vial, some of the radon will freeze out onto the interior surface of the vial, and the nitrogen will pass through unaffected. Moreover, if the source is kept frozen, we can apply a vacuum to the vial and the vast majority of the radon will not be removed — we can then use this vacuum to pull in scintillator from the storage flask in the same manner as was done for the ^{14}C sources as described in Section 7.2.2. By altering the amount of time that the radon

loaded nitrogen is allowed to flow through the vial, we can adjust the activity of the source.

Extra Hardware

The requirement that we be able to purge the vial with radon loaded nitrogen, backfill it, and then seal it without disconnecting it from the system (and thus risking oxygen contamination) poses somewhat of a challenge as we would like to force the gas to flow inside of the spherical portion of the vial to ensure it contacts the walls at liquid nitrogen temperature. The solution to this problem was to install a tee just above the neck of the source vial — the straight section of which acts as the exhaust and has a piece of $\frac{1}{16}$ " diameter stainless steel tubing that passes through an Ultra-Torr fitting so that it can be moved above the sealing point of the vial — the branch portion of the tee allows us to attach it to the loading station. By placing the end of the exhaust tube just above the bottom of the vial, and introducing the radon loaded gas through the branch portion of the tee, we force the gas to enter the spherical portion of the vial and have the opportunity to freeze out onto the surface. By sealing around the outside of the exhaust tube with an Ultra-Torr fitting, we maintain the possibility to move the exhaust tube without opening the system since the vial obviously cannot be heat sealed with the exhaust tube inside of it. A photo of this valve assembly, without the vial attached, is shown in Figure 7.12 and a closeup view of the vial with the exhaust tube in place can be found in Figure 7.13.

Filling The Vial

As mentioned earlier, the most common type of radon source produced was a compound ^{14}C - ^{222}Rn source, thus, we begin the loading process after pipetting approximately $100\mu\text{L}$ of the ^{14}C solution (prepared such that this amount of solution gives 30 Bq of ^{14}C) into the



Figure 7.12: The valve assembly used for loading the radon sources — the vial (not shown) attaches to the Ultra-Torr fitting at the right of the figure. This assembly connects to the source loading station via the valve at the bottom of the photo — the radon loaded nitrogen enters here as well and flows down into the vial and exits via the $\frac{1}{16}$ " tube placed just above the bottom of the vial, the exhaust gas flows up this tube and out the valve at the far left of the figure. The $\frac{1}{16}$ " Ultra-Torr fitting at the top (left in the figure) of the tee allows the exhaust tube to be retracted above the sealing point without opening the system. To ensure leak tightness, the valves employed are bellows type — Swagelok H series due to their smaller size and mass.

vial and attaching it to the valve assembly shown in Figure 7.12.

For safety reasons, the needle valves V4 and V5 on the radon source remain closed when the source is not in use, however, this causes the concentration of the radon in the source to build up and therefore produce more than the ≈ 3 mBq/min that it is calibrated to produce. In order to have repeatable results, we purge the source for several hours (usually overnight) prior to producing calibration sources so we can be sure that the gas stream contains a constant amount of Radon — a jumper tube is connected to the source loading station in lieu of the source vial, the other end of which is connected to a fume hood.

After the flow through source has been purged, a spare vial and valve are connected to the loading station and the line from the scintillator storage flask is primed as described on page 252. At this point, the real vial and valves are connected to the source loading station and a teflon tube is connected to the top of the valve on the exhaust tube, the other end of which



Figure 7.13: A closeup view of a ^{14}C - ^{222}Rn source just before radon loading. The exhaust tube can be seen here just above the bottom of the vial where a small pool of the ^{14}C solution has been placed. After immersing the vial in liquid nitrogen, radon loaded nitrogen flows down through the neck of the vial, and freezes out onto the surface of the vial, and the excess gas flows up the exhaust tube and to a fume hood. Before sealing the vial, the exhaust tube is withdrawn up to a level above the sealing point of the vial.

is routed to the fume hood. To ensure that the vial does not contain any air, we start a flow of unloaded nitrogen through the vial at a flow rate of 10 mL/min (as set by rotameter R3) for five minutes. Finally, the vial is immersed in the bath of liquid nitrogen and the flow of radon loaded nitrogen precisely set to 10 mL/min — it is during this time that the radon actually builds up in the vial.

The duration of the radon loaded nitrogen flow is determined by the desired activity of the source — due to the fact that they have to be shipped internationally, the sources are usually produced 1.5 - 2 weeks before their intended insertion date. For example, the source used in the second off-axis calibration was produced with a 10 mL/min flow rate of radon loaded nitrogen for 70 minutes — ten days later when the source was inserted, it had an activity of just under 300 Bq (although there was some question as to how well the source had been purged the night before because the rate was higher than anticipated).

After the time required to allow the buildup of the radon has elapsed, the Ultra-Torr fitting on the exhaust tube is loosened slightly and the tube pulled back so that it lies above the sealing point on the vial — since the vial is under positive pressure, any small leak introduced by loosening the fitting is an outwards one. The valve on the exhaust tube is then closed, the nitrogen flow is ceased, and the source vial is evacuated (since it is still in the liquid nitrogen bath, the radon is a solid and not affected by the vacuum — at least to first order). After pumping for several minutes, the vacuum is stopped and the vial allowed to warm and then the line to the scintillator flask is opened to allow the scintillator to fill the vial; again caution must be exercised to ensure the vial is not over filled. After re-freezing, the vial is evacuated, checked in the characterization detector and transported to the glass blower for sealing. Figure 7.14 shows two photos of sources used in the first off-axis calibration during fabrication.

7.2.4 Gamma Source Production

Two varieties of gamma sources were produced for the calibration campaigns: isotopes dissolved in water and loaded into a source vial, and an isotope deposited into an empty source vial. By not placing scintillator inside of the vial, the charged particles emitted with the gammas of interest do not produce scintillation light since they do not penetrate the vial wall — the gammas however, can pass through and interact in the Borexino scintillator.




Figure 7.14: Left: A ^{14}C - ^{222}Rn source during loading with radon. The spherical portion of the vial is immersed in a liquid nitrogen bath so that the radon loaded nitrogen passing through the vial will freeze out onto the inner surface of the vial. Right: The “radon in vacuum” source after loading and removal from the liquid nitrogen bath — note that the exhaust tube has been withdrawn above the sealing point.

Isotopes in Water

Throughout the calibration campaigns performed so far, the collaboration has made use of quite a few gamma sources fabricated by dissolving the isotope in water. Most of the isotopes are produced as a salt and dissolved in a weak hydrochloric acid solution — for example, the ^{139}Ce source was purchased as CeCl dissolved in a 0.5M HCl acid solution. Usually, the isotopes are provided with specific activities far greater than what is required for the source, and they can be diluted with deionized water to obtain the desired activity. Table 7.1 shows the isotopes and energies used during the various calibrations.

Isotope	Energy (keV)	Activity (Bq)
^{57}Co	122	3.45 ± 0.07
^{139}Ce	165	3.12 ± 0.04
^{203}Hg	279	2.78 ± 0.06
^{85}Sr	514	
^{54}Mn	834	
^{65}Zn	1115	
^{60}Co	1173 & 1332	
^{40}K	1460	

Table 7.1: The gamma sources used during the calibration campaigns, and their energies. All of the sources are useful in measuring the quenching of gamma rays (what is actually observed is the quenching of the many electrons which the gammas impart energy to as they Compton scatter), and provide monoenergetic lines for energy calibration. The first three sources were also used to study the trigger efficiency and were first measured in underground Ge detectors to determine their true activity (listed in the third column), and then inserted into Borexino — any difference in activity is due to the inefficiency of the Borexino trigger.

Radon in Vacuum

For the January 2009 off-axis calibration campaign, a source with many gamma lines was requested and we decided to use a vial in which we deposited some radon gas — the rich

gamma decay spectra of ^{214}Bi and ^{214}Pb provide many gamma lines, and the lack of scintillator blocks the betas produced by the decays in this chain. The source was fabricated following the same procedure as for the radon sources described in Section 7.2.3, except that no ^{14}C solution was placed in the bottom of the vial first, and obviously we did not backfill with scintillator.

Unfortunately, after the vial was sealed and it had warmed, a few spots of liquid appeared on the inside of the vial — it is thought that perhaps some pseudocumene vapor picked up by the nitrogen flowing through the plumbing of the source loading station was deposited into the vial with the radon. Fortunately, the scintillation yield of pure pseudocumene is very low⁵, and the impact of the few small drops on the output of the sources is negligible.

7.3 Characterization of Calibration Sources

Historically, the production of un-quenched calibration sources has been rather difficult, thus, it would behoove those preparing the sources to have a method of knowing the relative quenching of the sources before sending them to Italy to be used in a calibration. Moreover, we did not know, a priori, the correlation between the activity ultimately achieved in a source after flowing radon through a vial at a given flow rate for a specified interval of time, so, having a method to determine the activity of the sources after production is very important. The calibration of the detector is a 24 hour operation requiring, at the very least, four shifters at all times, so, the operations must be scheduled well in advance so that proper manpower may be on hand — being able to determine the activity of the sources weeks before the operation is very important in order to plan the operation and personnel accordingly. To that end, a “quality control” detector was constructed and a method for

⁵Being a solid, the PPO would have been left behind in the tubing of the loading station, thus, only pure pseudocumene could make it into the vial as a vapor.

measuring the relative quenching of the sources was developed at Virginia Tech.

7.3.1 Quality Control Detector

The purpose of this detector is to provide an estimate of the activity of a source and an energy spectrum which can be exploited to measure the relative quenching between the sources. Since quenching affects the light output from the scintillator, not the actual energy deposition, we must use a detector that looks at the scintillation light (i.e. a PMT based detector), not a detector which would measure the radiation emitted by the source (i.e. a germanium detector counting gammas emitted from the source). The design, electronics, and data acquisition of this PMT based detector are described in the following pages.

Detector Design

The interior of the quality control detector constructed at Virginia Tech is shown in Figure 7.15. Two Electron Tubes 9813B high-gain photomultiplier tubes with C638A voltage dividers sit inside of a recess in a Lexan cylinder designed to hold a source vial (the photocathodes are optically coupled to the Lexan with optical coupling grease). The Lexan holder contains a bore just larger than the source vial diameter and with a depth such that the vial origin is collinear with the axis of the photomultiplier tubes. The bore in the Lexan is filled with mineral oil to match the indices of refraction and the outside of the Lexan cylinder is wrapped in a highly reflective mylar film to increase light collection. The entire assembly sits inside of a dark box with feed throughs for the high voltage and signal cables as well as two $\frac{1}{4}$ " Swagelok connections. To avoid unnecessarily sealing a quenched, or otherwise defective, source, they are always placed in the detector after loading (and before sealing), thus, the dark box must be able to accommodate the valve assembly that is placed on the

sources for loading (see Figure 7.12), so, the lid contains a 14” “top-hat” to accommodate the height of the retracted exhaust tube.




Figure 7.15: The inside of the dark box housing the quality control detector for evaluating calibration sources at Virginia Tech. The blue cylinders at either end are the Electron Tubes C638A voltage dividers connecting to the 9813B photomultiplier tubes. The PMTs view the center of a Lexan coupler where the source vials are placed. The Lexan coupler is wrapped in a highly reflective mylar film to increase light collection. The lid of the box (not shown) contains a long top hat to accommodate the valves present on the sources before sealing.

Electronics

The pulse processing of the signals is performed by NIM modules in an electronics rack residing next to the dark box; a CAMAC crate houses the charge ADC used to acquire the energy spectrum and is controlled by a computer in the same crate. A photo of the NIM modules and annotated wiring scheme is visible in Figure 7.16. To reject dark noise, the gate of the ADC is only issued when signals are coincident between the two PMTs. The wiring of the system shown in Figure 7.16 shows a setup where the system issues a gate on the second of two signals which occur in a time interval of less than 1 ms — which for a weak enough Radon source, will trigger predominantly on the ^{214}Po alpha decay.




Figure 7.16: Photo of the pulse processing electronics for the quality control detector — the cabling was removed for clarity and the connections redrawn by hand. From left to right, the modules are: Fan-In/Fan-Out, x10 amplifier, quad discriminator, triple coincidence unit, an unused delay module, and a timer / gate generator. The setup shown in the photo will issue a gate on the second of two events that occur within a 1 millisecond time interval — i.e. to trigger on the ^{214}Po event in a radon source.

The signals from the PMTs arrive at a Fan-In/Fan-Out where multiple copies are produced — the first copy goes to a delay and then to the ADC (a Lecroy 2249W 12 channel, 11 bit charge ADC) — the second copy goes to a discriminator set to reject the vast majority of the dark noise pulses; the output of the discriminator is then sent to the coincidence unit. If one simply wishes to acquire all signals which are coincident between the two PMTs, the gate signal can be generated at this point. To trigger just on the second of two pulses in a short interval, the output of the coincidence unit is fed to a gate/delay generator — which creates a 1 ms long gate after initially waiting for 100 ns to block re-triggering — this 1ms long gate signal is fed back to another block of the coincidence unit along with the signals from the discriminators for each channel. When this three-way coincidence occurs between both PMTs and the gate signal established by the presence of a preceding coincidence, the ADC gate signal is generated.

Since all of the modules used to generate the gate signal require a finite amount of time, the PMT signal that is sent to the ADC must be delayed so that it falls inside of the ADC gate. The required delay is measured by replacing the ADC with an oscilloscope and triggering on the gate signal and simply counting the time interval required to place the PMT signals inside of the gate — this typically amounts to about 50 ns.

The system was also configured to compute the tail to total ratio (TTR — see Section 3.3.2) of events by routing another copy of the PMT signals to the ADC, but delayed slightly less than the other copy so that the first ~ 30 ns of the pulse do not lie inside of the gate — in other words, only the tail of the pulse (defined to be anything after 30 ns) was integrated. During the processing of the data afterwards, the ratio of the truncated signal to that of the full signal would be a measure of the contribution of the slow component of the pulse to the total, which is higher for alphas than for betas. Unfortunately, the signal provided directly from the PMTs was too fast, and it was not possible to precisely acquire only the tail

of the signal — shaping of the signal with a preamplifier would have made it considerably easier. However, the alpha-beta discrimination in this detector was not a requirement, so, the attempt was abandoned due to more pressing issues.

A tail pulse generator in the NIM crate is often used in place of the PMT input signals as a “calibration” pulse for evaluating the timing and efficiency of the system. To check the ability of the system to resolve two pulses arriving in close temporal proximity to each other, the tail pulse generator was used to generate double pulses with an adjustable delay. The true number of pulses sent through the system was determined by routing the sync-out signal from the pulse generator to a scaler; the number of detected pulses is obtained from the output of the coincidence unit as counted by another scaler. The system maintained an efficiency of 100% for pulses with spacings greater than $3.5 \mu\text{s}$, so, source rates up to ~ 285 kBq can be acquired without introducing any trigger inefficiencies.

The other component which is rate-limited is the charge integration performed by the ADC. The same method was employed to determine the efficiency of the DAQ, and the maximum continuous rate of acquisition is approximately 2 kBq, however, the system can handle short bursts of activities up to several hundred kBq. The efficiency of the DAQ is not so important because the activity of the source is measured by the crate electronics and not the acquisition of the energy spectrum, however, it does tell us that if the source is too strong, we cannot obtain an energy spectrum from it.

DAQ

The ADC readout is performed via a crate controller which interfaces with an ISA card in the computer which controls it — ISA was the predecessor to the PCI interface and is obsolete, thus, it requires an equally old computer to control it. The acquisition is performed via

the MIDAS (Maximum Integration Data Acquisition System) DAQ software from TRIUMF and PSI [117]. The default data read-out code bundled with MIDAS was heavily profiled and modified to improve performance since the computer controlling it is only an 870 MHz computer running a very modern OS (CentOS 5 Linux), so, speed improvements are essential where possible.

The computer reads out the ADCs whenever a LAM (Look-At-Me) interrupt is issued by the CAMAC controller (which, by default, occurs after the completion of the conversion of the analog signal into a digital number by the ADC). During the readout process, the DAQ subtracts the pedestal value from the value reported by the ADC — the pedestals were obtained by terminating the signal inputs with $50\ \Omega$, and then pulsing the gate with a square wave generator and fitting the result to a Gaussian. After subtracting the pedestals, the results are saved into a ROOT file, one histogram for each channel.

To correct for gain changes, periodic spectra of the dark noise of each channel were obtained by replacing the gate signal with a square wave from a pulse generator (the ADC was required to have a value greater than 10 channels before anything was binned, so, the spectrum was not simply a huge peak at zero). The acquired spectra were then fit to a Gaussian plus an exponential and the mean of the Gaussian was recorded. Variations in the mean of this Gaussian (the center of the single photoelectron signal) are indications that the gain of the PMT has changed — however, periodic measurements recorded over the course of several months did not exhibit variations at a level larger than 1% therefore they are neglected in the treatment that follows.

To have a sense of the temporal ordering of the pulses, a scaler is also installed in the CAMAC crate and is fed with the output of a square wave generator whose frequency can be adjusted. If desired, the output of the scaler can be acquired whenever the ADCs are read, the value of the scaler at that time can be used to determine the timing of the events

if the widths and separations of the pulses arriving at the scaler are known. Since reading out this extra device requires more time, it is generally disabled to allow the acquisition of higher rate sources.

7.3.2 Relative Light Yield Measurement

We desire a non-invasive method for determining the relative quenching of scintillator samples, both loaded and non-loaded. While spectrophotometry affords the ability to test the transmission of light through the samples, which is of course useful to know, it does not reveal anything about the light emitted by the scintillator itself. For weak radon sources, it is possible to efficiently tag the ^{214}Po alpha line — two sources whose alpha peaks do not overlap indicates that one of the sources is quenched with respect to the other. However, for alpha sources with a very high activity, the tagging method discussed in the previous section is not efficient and the spectrum obtained is a very broad peak which contains contributions from ^{222}Rn , ^{218}Po , and ^{214}Po . Further, this method only works for sources containing Rn and does not help with, for example, a ^{14}C source — spectral fitting of this type of source is also not very useful as the spectral shape is poorly constrained (see Section 7.2.2).

A non-invasive method is to expose the vial to gamma rays from an externally mounted source and use the position of the Compton edge as a probe of the quenching of the source. The small size of the source vials and the small atomic number of the “absorber” (in this case, the scintillator) means that the probability of the gamma to undergo photoelectric absorption is exceedingly small, and thus does not allow us to exploit its monoenergetic nature. The method does have one drawback: it is difficult to use on a very active source with a rich spectrum, such as a Radon source. The presence of multiple isotopes in the vial often obscures the Compton edge. However, this can be partially overcome by using

a higher-energy gamma-ray source such that the Compton edge lies above the bulk of the source spectrum.

The difficulty in using this method is how to properly identify the Compton edge. A common method is to simply take the value of the energy at which the spectrum has fallen to half of its initial height, but, this method only works if there are no other contributions to the spectrum, or if there are, they must be known very well and accounted for with little error in the fit. A more precise method is to perform a fit to the expected spectral shape, and this is what was done at Virginia Tech.

Fitting Function

The success of this method relies on being able to locate the edge of the Compton shoulder. While shifts can be determined visually, the location of the Compton shoulder in experimental plots is very difficult to determine by eye. The ideal method is to fit the spectrum in the region of interest. A good approximation of the functional form of the experimental spectra (neglecting multiple scattering, photoelectric effect, partial containment of secondary electrons, etc.) is obtained by convolving the Klein Nishina differential cross section for photon - electron scattering with an energy resolution function. The Klein-Nishina formula (with respect to photon energy) is given by [118]:

$$\frac{d\sigma}{dT} = \frac{\pi r^2}{m_e c^2 \rho^2} \left(2 + \frac{s^2}{\rho^2 (1-s)^2} + \frac{s}{1-s} \left(s - \frac{2}{\rho} \right) \right) \quad (7.3)$$

$$\rho = \frac{T_\gamma}{m_e c^2} \quad (7.4)$$

$$s = \frac{T}{T_\gamma} \quad (7.5)$$

To emulate the detector response, $\Phi(T)$, it must be convolved with an energy resolution function, which we take to be a Gaussian with energy resolution ξ :

$$\Phi(T) = A_0 \int_0^{T_C} \left(\frac{d\sigma}{dT'} \right) \cdot e^{-\frac{(T-T')^2}{2\xi^2}} dT' \quad (7.6)$$

In equation 7.6, T_C is the energy corresponding to the Compton edge, and A_0 is an overall normalization constant. Since the actual variable recorded by the DAQ is the ADC channel number, x , we take the energy to be given by $x\eta$ where η is the conversion factor from ADC channel number to MeV. After performing the integration of equation 7.6, we obtain the fitting function:

$$\Phi = \frac{-A_0 \rho \xi m_e c^2}{T_\gamma^4 (T_\gamma - \eta x)^2} \{c_1^2 + c_2 + c_3 \cdot c_4\} \left(\text{Erf} \left[\frac{x\eta}{\sqrt{2}\xi} \right] - \text{Erf} \left[\frac{x\eta - T_C}{\sqrt{2}\xi} \right] \right) \quad (7.7)$$

$$c_1 = -m_e c^2 \eta x \quad (7.8)$$

$$c_2 = 2x\eta m_e c^2 T_\gamma (T_\gamma - x\eta) \quad (7.9)$$

$$c_3 = T_\gamma (x\eta - T_\gamma) \quad (7.10)$$

$$c_4 = x^2 \eta^2 + 2T_\gamma (T_\gamma - x\eta) \quad (7.11)$$

The parameters in the fit then are the following:

- T_γ — The energy of the gamma imparting energy to the electrons in the scintillator (in MeV).
- ξ — The energy resolution of the detector, roughly 7% at 500 keV.
- η — The conversion factor from ADC channels to energy, a quasi light yield (labeled

LY in the fit results shown later). The conversion turns out to be roughly equal to 1.7 channels / keV.

- T_C — This is the energy of the Compton edge and is the parameter that we are ultimately seeking to determine with these fits.
- A_0 — The overall amplitude of the fit, and is not of importance for the results.

The fit is performed while restricting all components to be positive. The Compton edge position T_γ is restricted to the range $[0.8, 1.2] \times T_\gamma^0$, where T_γ^0 is the theoretical position of the Compton edge given by equation 7.12.

The energy resolution is confined to lie in the region $[0, 0.1]$ over the range of the fit, and the conversion from ADC channels to energy is restricted to $[1, 3]$ channels / keV.

$$T_\gamma^0 = \frac{E_\gamma}{1 + \frac{m_e c^2}{2E_\gamma}} \quad (7.12)$$

An important fact is that the fit does not consider whether or not a source is quenched to begin with — it simply allows one to compare the Compton edge positions between two different sources. An absolute energy calibration of this detector is very difficult and was not ultimately required — by taking spectra of sources that are later inserted into Borexino, we can work backwards from the Borexino measurement to roughly determine what a truly unquenched source looks like in the quality control detector.

Fit Results on an Unloaded Vial

Before attempting to use this method on a source, we would like to verify that the fit can actually distinguish between a quenched and non-quenched source. In order to ensure we

have one source quenched more than the other (since we cannot determine how quenched our stock solution is with this method), we can intentionally quench the source by bubbling pure oxygen through the vial.

The method employed for this study was to use a vial assembly similar to that used to construct the radon sources (see Section 7.2.3) where the gas flow was sent through a tube immersed in the solution (the tube used for this study was fixed and made of Teflon) — a photo of this assembly is presented in Figure 7.17. The valves on the source vial were attached to opaque tubing inside of the dark box that connected to bulkhead Swagelok feed throughs installed in the lid to afford the possibility of sparging the vial while it was inside of the detector.

After installing the vial in the dark box, it was sparged with high purity nitrogen for thirty minutes to remove any dissolved oxygen. A ^{137}Cs gamma source ($t_{1/2} = 30.07$ years, $E_\gamma = 661$ keV) with a nominal activity of $1\mu\text{Ci}$ was mounted on the outer surface of the dark box in a manner such that it was directly in line with the center of the source vial. 30000 events were acquired and the spectrum was fit with the function described in equation 7.7 to yield the “unquenched” Compton edge.

The purge line to the source was then connected to a gas cylinder containing pure oxygen and the source was sparged with this pure oxygen stream for ten minutes. Another spectrum of 30000 events was acquired and then fit to obtain the “quenched” Compton edge. Plots of the results for the two cases (quenched and unquenched) are presented in Figure 7.18 — when fitting the quenched case, the “light yield” (conversion from ADC channels to energy) was fixed to the result of the first fit.

We can see from 7.18 that the fitting function accurately describes the acquired spectrum (the χ^2 value of the top plot is 1.06) and that the method is able to show the effect of




Figure 7.17: The vial assembly used for testing the Compton edge fitting procedure. For this study, the tube inserted in the scintillator is actually used to supply gas to the scintillator to sparge it in situ rather than as an exhaust tube. The downstream side of the plug valves shown in this photo are connected to opaque tubing which passes through light-tight feed-thrus in the lid of the dark box to connect to a gas supply for the inlet, and a charcoal filter for the exhaust.

quenching. An example of the fit performed after exposure to a ^{22}Na source (which produces two gammas of energies 511 keV and 1.27 MeV) is shown in Figure 7.19.




Figure 7.18: Energy spectra obtained from a source vial exposed to a ^{137}Cs gamma source. Top: Spectrum after sparging with nitrogen for thirty minutes to remove any dissolved oxygen. Bottom: Spectrum after sparging the source with pure oxygen for ten minutes to intentionally quench the scintillator. It is clear from the fit results that the exposure to oxygen quenched the source by about 8%.




Figure 7.19: Fit results of an energy spectrum obtained after exposing the source vial to a ^{22}Na source emitting two gammas of 511 keV and 1.27 MeV. The two spectra are identical, but the bottom pane shows the makeup of the fit — two Compton spectra summed together.

Method for Higher Activity Sources

A major limitation of the source characterization detector (described in Section 7.3.1) is the maximum data acquisition rate of a few kBq (limited by the speed of the computer controlling the DAQ). A source with a very high rate poses two separate problems: the first is that it affects the stability of the DAQ, the second is that it often swamps the signal from our gamma sources, and trying to use a gamma source of higher activity will further cause rate problems. For the case of Radon sources, the alpha peaks from ^{222}Rn , ^{218}Po , ^{214}Po , and ^{210}Po are rather close together — 5.490, 6.002, 7.687, and 5.305 MeV respectively (quenched down to roughly 420, 480, 750, and 400 keV) — and appear as one broad peak in the energy spectrum. The position of this peak unfortunately washes over the Compton edges of ^{137}Cs , ^{22}Na , and ^{60}Co as shown in the plot contained in Figure 7.20

Since the direct measurement of the Compton edge is impossible for a radon source with an activity of several hundred Bq⁶, we perform an indirect measurement with a second source. If we produce a source with an activity in the range of 10 Bq, the contribution from the externally mounted gamma source will be dominant over the radon in the vial and we can extract our Compton edge in this manner. Then, if we remove the source and count for a longer interval of time, we can obtain the spectrum from the radon in the source — this we can compare directly to other Radon sources, and the measurement of the Compton edge allows us to compare it to other non-radon sources. A Gaussian fit to the large peak from the alphas in the radon chain will yield its mean; variations in this mean value between two sources will reveal the presence of any quenching. If the peak of the radon source coincides with that of the weak radon source, and the Compton edge of the weak radon source coincides with that of the ^{14}C source (or any source not containing radon), we can be sure that all

⁶In order for the source to have an activity of ~ 100 Bq when inserted into the detector, it must be produced with an activity larger by a factor of > 4 because it will take at least two half lives (8 days) to reach Italy.




Figure 7.20: Experimental data from a radon source (in blue), and a vial exposed to three different gamma sources (see legend in plot). The presence of the very broad peak in the energy spectrum of the radon source (actually a sum of peaks from four alphas) obscures the Compton edges of each of the three gamma sources shown, making it impossible to do a direct measurement of the quenching of the source by exploiting the position of the Compton edge.

such sources are not quenched relative to each other.

7.3.3 Measurements of Sources Used In Calibrations

The method described in the previous section does not allow one to perform an absolute measurement of the quenching of the sources, but, the Borexino detector does. Therefore, we can only assure that the first calibration sources are not quenched relative to each other. After analyzing the results from these sources in the Borexino detector, we will know their absolute quenching and then can use that information to determine the degree to which any subsequent sources are quenched. For the deployment locations of the sources used in the

calibrations, see Figure 8.1 and the appendices at the end of this dissertation.

On-Axis Calibration

Between October 5th and 10th of 2008, four sources (^{14}C - ^{222}Rn , ^{54}Mn , ^{85}Sr , and ^{14}C) were deployed along the vertical axis in what was the first of four calibration campaigns performed. At the time, the off-axis system was still awaiting the cleaning of several critical components, but a calibration was desired because the leak in the inner vessel had been recently discovered and the future of the experiment was unclear. The two gamma sources used were produced at the laboratory in Italy with some extra isotope solution that was left over from a calibration of the CTF, however, the rates were very low: 3 Bq and 0.3 Bq for the Mn and Sr sources respectively.

The two sources produced at Virginia Tech were ^{14}C - ^{222}Rn and the pure ^{14}C and were prepared using pseudocumene from a drum purchased several years ago, filtered through an alumina column, mixed with 1.5 g/L of PPO, and then sparged for several hours to remove any dissolved oxygen.

After producing the weak radon source, spectra were recorded by the quality control detector with and without an externally mounted ^{60}Co source (see Figure 7.20 to see why ^{60}Co is used). After producing the ^{14}C - ^{222}Rn source to be used in the calibration, an energy spectrum was obtained and compared with the weak radon source — the results are presented in Figure 7.21. The overlap of the peaks in both sources tells us that these two sources have the same light output characteristics (i.e. quenching).

After producing the pure ^{14}C source, we wish to ensure that its light output is, at the very least, as good as the compound source, so, a spectrum is taken with an externally mounted

^{60}Co source⁷. The spectrum from this source was then placed on top of the spectrum from the weak radon source with the ^{60}Co source — the results are presented in Figure 7.22. The Compton edges in both spectra lie effectively right on top of each other, although, the presence of events above the peak in the radon source causes a distortion in the tail, thus, the fit was restricted to the very narrow range around the knee.




Figure 7.21: Comparison of the energy spectrum from the radon source used in the on-axis calibration (in red), and the weak radon source (blue) created to allow checking both radon and non-radon sources for quenching. The fact that the two peaks overlap with each other indicates that the two sources are not quenched relative to each other.




Figure 7.22: This plot shows two sources exposed to ^{60}Co — the weak radon source (blue) and the ^{14}C source used in the on-axis calibration (red — note the tail of the ^{14}C spectrum on the left of the histogram). The fits to the Compton edges are not shown here for clarity, but they agree to within a few percent, thus the ^{14}C source is not quenched relative to the weak radon source and, by the results of Figure 7.21, not quenched relative to the compound ^{14}C - ^{222}Rn source.

With these results in hand, we were assured that the sources being used in the calibration were at least identical to each other in terms of their light yield, however, it was unknown whether it matched the light yield of the Borexino scintillator. After inserting the sources into the detector, it was immediately apparent that they were in fact quenched by $\sim 30\%$. While the 30% quenching of the on-axis sources was disappointing, the quenching studies

⁷Since the ^{14}C decay is very low in energy, it is possible to avoid the difficulties with the stability of the detector for large counting rates by increasing the levels on the discriminators, however, for the 30 Bq source discussed here, this was unnecessary.

performed on those sources and the fact that they were produced by different methods allowed us to eliminate many potential explanations for the quenching. Two explanations persisted: either the scintillator used was quenched from the beginning, or, the design of the containment vials induces the apparent quenching.

The latter is explicable if the radiative transfer from the pseudocumene molecules to the PPO molecules plays a more appreciable role in the energy migration than previously thought. Since the concentration of PPO is small compared to that of pure pseudocumene, direct excitation of the PPO is also negligible. Literature⁸ suggests that the vast majority of the energy transfer from solvent to solute takes place non-radiatively (primarily via thermal diffusion of the excited molecules or inducing resonances in other nearby molecules) — this is the reason why the fluorescence emission from pure pseudocumene is very inefficient. If, however, radiative migration of the energy plays a more important role than anticipated, it would require the PPO to be excited by absorbing a photon emitted during the de-excitation of a PC molecule. Since the emission spectrum of PC predominantly lies at lower wavelengths where the transmission through quartz is very inefficient — this could lead to a loss of light and imitate “quenching”.

The effect of the quartz absorbing primary scintillation light from pseudocumene before it is able to be absorbed by a PPO molecule should be proportional to the surface-to-volume ratio of the vial. Thus, we can test whether or not this is an important effect by examining spectra produced by light contained inside of vials with different surface to volume ratios. A new vial was constructed with a cylindrical shape with an outer diameter of 1” and a surface to volume ratio larger than the spherical vials by a factor of 1.9. After loading each with the same scintillator, spectra were taken under exposure to ^{60}Co and the Compton edges were extracted by the spectral fit. A shift of less than 2% was found, and this is within the

⁸See [64] and the articles cited therein

error of the fitting algorithm and is definitely not large enough to cause the observed 30% quenching effect.

This left the scintillator as the likely culprit. A common misconception is that one can simply sparge scintillator to remove dissolved oxygen and produce an un-quenched sample of scintillator, this is however not the case. Oxygen is a known quenching agent, and if left in contact with pseudocumene for a sufficient period of time, molecules will react with the oxygen to form dimethylbenzaldehyde or other oxygen-containing molecules — these are chemical changes and are not removed by sparging with a gas (fortunately, DMBA has a higher boiling point than pseudocumene, thus, distillation can separate the two chemicals). An equally plausible explanation is that the pseudocumene used for the on-axis sources was fine, but the PPO used to prepare the scintillator may have degraded over time. Nevertheless, since the Borexino scintillator is our “golden standard”, we carefully obtained samples of the scintillator directly from the distillation system during late 2008, and again in the spring of 2009. The Borexino scintillator was then used to load a source vial at Virginia Tech and a spectrum was taken after exposing to a gamma source and we achieved the results shown in Figure 7.23 — a 30% increase in the light output as evaluated by the shift in the Compton edge.

First Off-Axis Calibration

By the beginning of 2009, the full off-axis calibration system was ready and was put to extensive use from January 16 thru February 6 to deploy four different sources in over 95 different positions throughout the detector. The sources deployed were: another ^{14}C - ^{222}Rn compound source, a ^{14}C source with a nominal activity of a few Bq, a ^{222}Rn in vacuum source, and the AmBe neutron source. The compound source was intentionally loaded with a higher concentration of radon due to the large number of points where the source had to




Figure 7.23: Energy spectra from source vials filled with scintillator taken from the Borexino distillation system (blue) and the scintillator used to produce the on-axis calibration sources (red). This plot clearly shows that the scintillator used to make the on-axis sources was quenched by 30% relative to the Borexino scintillator. The blue curve is actually a spectrum of the ^{14}C source used in the first off-axis calibration with an externally mounted ^{60}Co source — the tail of the ^{14}C spectrum can be seen in the first few bins of the plot.

be placed during the campaign — the activity was approximately 100 Bq when inserted and fortunately did not cause problems with the stability of the Borexino DAQ. The compound source also contained double the amount of ^{14}C as the on-axis source — approximately 60 Bq. A plot of the energy spectrum obtained by Borexino for both the on- and off-axis sources is shown in Figure 7.24 below — the effect of the quenching is obvious, as is the increased amount of ^{14}C for the off-axis source.

Another source deployed during the first off-axis calibration was the “Radon in Vacuum” source where the gammas produced in the decays of ^{214}Pb and ^{214}Bi produce several prominent gamma lines which can be calibrated to. Due to the size limitations of the detector at Virginia Tech, we could only check whether or not this source had any activity in it at all —




Figure 7.24: Energy spectrum (in the npe variable) obtained by Borexino for the on (red) and off (blue) axis compound ^{14}C - ^{222}Rn sources deployed approximately at the center (runs 8817 and 9462 respectively). The data is restricted to those events which are reconstructed within 50 cm of the center of the detector. The small falling edge at around 50 npe is the ^{14}C component of the source (whose concentration was doubled for the off-axis source), the first peak is the sum of ^{222}Rn and ^{218}Po (5.5 and 6.0 MeV respectively), and the second peak is that from ^{214}Po (thus the 2:1 height ratio between the first and second peaks — the presence of a small amount of ^{210}Po in the first peak causes this ratio to be slightly higher), and finally, the broad spectrum extending up to 1400 npe is that of ^{214}Bi which contains numerous gamma rays and results in a very complex spectrum.

a Ge detector would have been the optimum method of checking the activity of the source, however, ours had been sent out for repair at this time. To test for activity, we placed a vial with scintillator inside of the detector and counted the coincidence rate to obtain the “background” activity of the vial, and then placed the Rn in vacuum vial against the outer surface of the box and recorded the activity. After doing so, the count rate increased by a factor of two (which was not within statistical fluctuations), thus, the source was deemed to indeed have Rn inside of it and was considered ready to be deployed in the detector.

Second Off-Axis Calibration

Following the first off-axis calibration, the vessel was refilled several times in order to attempt to bring its shape back towards spherical, and there was a desire to test if the shape of the vessel had any impact on the position reconstruction. Moreover, due to time constraints during the first off-axis calibration, we were not able to deploy the AmBe source in as many locations as we would have liked. To provide these data, a third calibration was performed which required one month of around-the-clock operation — several scheduling issues required that this calibration be broken into two campaigns, each two weeks long occurring at the end of June and end of July respectively. A compound ^{14}C - ^{222}Rn source was to be deployed in some of the same places as in the first off-axis calibration (deviations from the previous results would indicate an effect caused by the shape of the vessel), and then in 15° increments in theta at the fiducial volume radius to provide data as to where events of different energies reconstruct when placed at the fiducial radius — a total of more than 100 data points. To expedite the process, an even more active source was desired, and the one produced had an initial activity of approximately 300 Bq. The scintillator used to produce this source was from the same bottle of scintillator as for the first off-axis calibration, and had been sitting for several months after being opened — thus, some quenching would not be completely surprising. Indeed, examining the spectra obtained from placing these sources in Borexino indicates a small degree of quenching that can be seen in Figure 7.25.

The rest of the calibration campaign was focused on deploying the gamma sources described in Section 7.1.2 for trigger efficiency and quenching (ionization, not impurity) studies, and deploying the AmBe neutron source in numerous locations. Some results from these calibrations are presented in the following chapter.

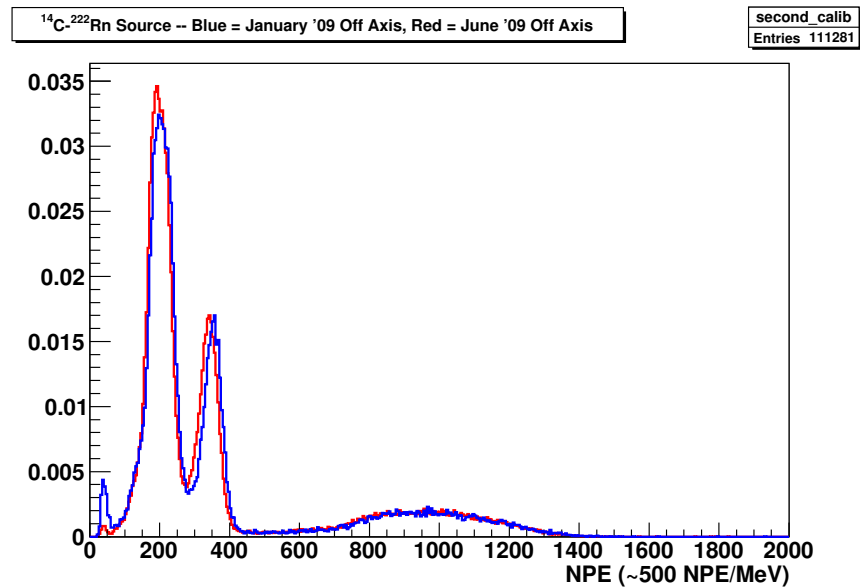


Figure 7.25: Comparison of spectra obtained in Borexino with the compound ^{14}C - ^{222}Rn sources used during the January 2009 off-axis calibration (blue) and the June 2009 off-axis calibration (red). The data are for sources at the center of the detector (run 9462 and run 10309 respectively). From looking at the ^{214}Po alpha peak (centered around 350 npe), we can see that the source for the June calibration appears to be quenched by approximately 5% relative to the on-axis calibration; the agreement is better for the lower energy alpha peak and the ^{214}Bi spectrum however.

Chapter 8

Results From The Calibrations

The four calibration campaigns performed between 2008 and 2009 resulted in the deployment of 19 different sources, in 295 locations inside of the detector (see Figure 8.1), providing more than 35 days worth of DAQ time. With such a vast quantity of calibration data on disk, it necessarily requires that a multitude of people devote their attention to extracting as much information as possible from these data. The rest of this section is devoted to the subset of that work performed by the author — concentrating mostly on determining the level to which the calibrations introduced radioactive contamination to the detector, as well as some results on the position and energy reconstruction.

8.1 Performance of The Calibration System

During the commissioning of any complex system like the one described in Chapters 5 and 6, problems are inevitable; indeed some were found, however, the system performed mostly in the way it was intended. To ensure there were no gross problems with the on-axis variant of the system used in October of 2008, a short test of the system was performed in the

CTF during the early fall of 2008. The test consisted of lowering the tether tube with a counterweight and unloaded source vial into the CTF through a version of the Borexino tube seal modified to fit on the CTF top flange. No increase in the radon contamination was seen in the CTF following this test — but note that the CTF is a less sensitive instrument than Borexino. In addition, the test did not employ a glovebox or sophisticated gas handling and monitoring system, however, a null result in the CTF was considered to be a sufficient go-ahead for using it in Borexino.

8.1.1 On-Axis Calibration

The on-axis calibration was done primarily to obtain *some* calibration data that could be used for studying errors on the data we had recorded up to that point. At the time, there was great concern that the leak in the vessel would become larger and ruin the experiment before it was calibrated. The system used for the on-axis calibration was put together in a somewhat rushed manner, and we encountered a problem with locating the source. The hardware used for the on-axis test can be seen in Figure 8.2. The general idea is a counterweight with Swagelok connections machined onto each end — the tether tube attaches to the top one, and the diffuser to the bottom. Inside of the counterweight, the Kevlar structural members of the fiber optic are potted with an epoxy to form a plug that prevents the counterweight from falling should the connection to the tether slip or break. The diffuser used in the on-axis test was one of our quartz vials filled with cleaned $\frac{1}{16}$ " diameter borosilicate spheres (used as packing in distillation columns) inside of which the fiber was placed. Figure 8.2 shows the hardware used in the CTF test where the diffuser was a piece of FEP Teflon since we did not care about light blocking for that test.




Figure 8.1: The location in $\rho - z$ space where the calibration sources were deployed during the off-axis calibrations.



Figure 8.2: The hardware used for the on-axis test in the CTF. From right to left: cone to guide the assembly back into the fill pipe, counterweight, diffuser, and source vial. When used in Borexino, the diffuser was replaced with another source vial filled with glass beads.

The illumination method used in the on-axis test was to be a red laser pointer mounted in a fixture with the fiber optic cable. The batteries in the laser pointer were removed and the connections hard-wired to a power supply controlled by a relay inside of the camera control box in order to synchronize the illumination of the diffuser with the photo taking operation. Unfortunately, the alignment fixture was not adequate and the laser easily fell out of alignment with the fiber optic, and little, if any, light was transmitted to the diffuser, thus it was impossible to locate the sources. We were able to finally see the source with the cameras after switching to a green laser (see Section 5.2.1 for the reason why the cameras are more sensitive to green light than red), however, the PMTs of the detector are themselves highly sensitive to green light, so much so that they had to be turned off when a photo was taken. The general procedure followed was to lower the source to a position, turn off the PMTs and take a photo, turn the PMTs on and take data, move the source to a new position and take data, turn off the PMTs and take a photo, etc. A photo of the illuminated source inside of the detector can be seen in Figure 8.3.

8.1.2 Off-Axis Calibrations

While undergoing a final commissioning of the insertion system prior to the first off-axis calibration in January 2009, we experienced some difficulties with the impedance of the gas regulators for the insertion system — in particular, R1 and R3 were both set near the bottom of their operating range, thus representing a significant impedance to gas flow. The series combination of the two regulators provided enough impedance that it was difficult to obtain enough flow when both the rod seal and cross were being purged; consequently, when one of two were closed, the needle valve on the output of the other needed adjustment as now the flow was too high. Prior to this point, the system was meant to be operated in a constant state of purging the cross. However, the discovery of this problem made this setup

prohibitive. To deal with this issue, we decided to implement automatic regulation of the cross pressure which allowed the cross purge flow to be set very low and this helped to cure the problem of the regulator impedances.

During the third day of off-axis calibrations, the computer which controlled the gas flow and camera systems froze because of a memory leak in the camera control software and required rebooting. Unfortunately, the computer froze with the cross feed valve closed and the exhaust valve open, so the pressure in the cross bled off and it flooded with scintillator. The source was withdrawn and the calibration suspended to allow some of the gauges to dry. During this time, the control of the gas control system was off-loaded onto another computer. As another safety measure, manual toggle switches were installed inside of the gas box on V5 and V7 (see Figure A.1 for the P&ID) to allow the shifters to close the valves and isolate the cross regardless of the state of the computer.

System Operation

The system is operated by three people in CR4 (plus one or more shifters in the counting room to control the DAQ): one calibration expert, and two other shifters, one of whom assists the expert in the glovebox, and the other shifter operates the computers. The system can actually be operated with just one person in the glovebox and one on the computer if a small movement of the source is required (one that does not require addition or subtraction of a rod). When the source is being moved, photos are taken frequently to ensure that the source is following the expected movements — for instance, when lowering the source back to vertical before extracting it, a photo is always taken to ensure that the hinge is in fact rotating down towards vertical. Also, when inserting or withdrawing a source, a photo is always taken after a rod is added or subtracted so that if operators lose track of their steps, we have some recourse since the status of the system can always be determined from the

series of photos taken.

The source is attached to a straight rod and lowered into the detector, attaching additional rods and substituting in the hinge when necessary (usually the third or fourth rod). After the center of the hinge has reached the required location, the tether tube is extracted by the amount dictated by the calibration software, and a photo is taken. Normally, the source is fairly close (within 10 cm) to the desired location. However, the software does not account for the fact that the fill pipe is not completely vertical, and that the Teflon tube does stretch slightly, so if the position of the source is absolutely critical, refinement of the position is then made. After taking the photo, the coordinates of the source are sent by radio to the shifters in the counting room where they are entered in the master table (they can also be regenerated at any time if one knows the DAQ run which corresponds to a given photo). Because the transparent Teflon tether tube does funnel some light into the detector, the lights in the cleanroom are usually extinguished and opaque plastic is placed over the acrylic windows of the glovebox. A photo of the neutron source deployed during the calibration in June 2009 can be seen in Figure 8.4

Status of the System

At the time of this writing, the off-axis system has been unused for approximately six months, but is being constantly purged with a low flow of LAKN. In theory, a calibration operation can be performed with as little as one day of preparation by a calibration system expert. However, more than six weeks of constant use has taken its toll on the system. In particular, the tether tube does contain some kinks and deformation mostly from over-tightening of the tube seal and not properly lubricating the o-rings of the tube seal before operation. The fiber optic is fine (however, it was broken once — fortunately very near the source, so, it was an easy repair and did not require removal of a large length of the fiber optic so).




Figure 8.3: Photo of a source deployed with the on-axis system taken from camera 7 on October 10, 2008. The bright spot just above the source in the full image is a reflection of the house lights from the tether tube — the zoomed photo shows the cone, counterweight, and diffuser illuminated with the green laser light.

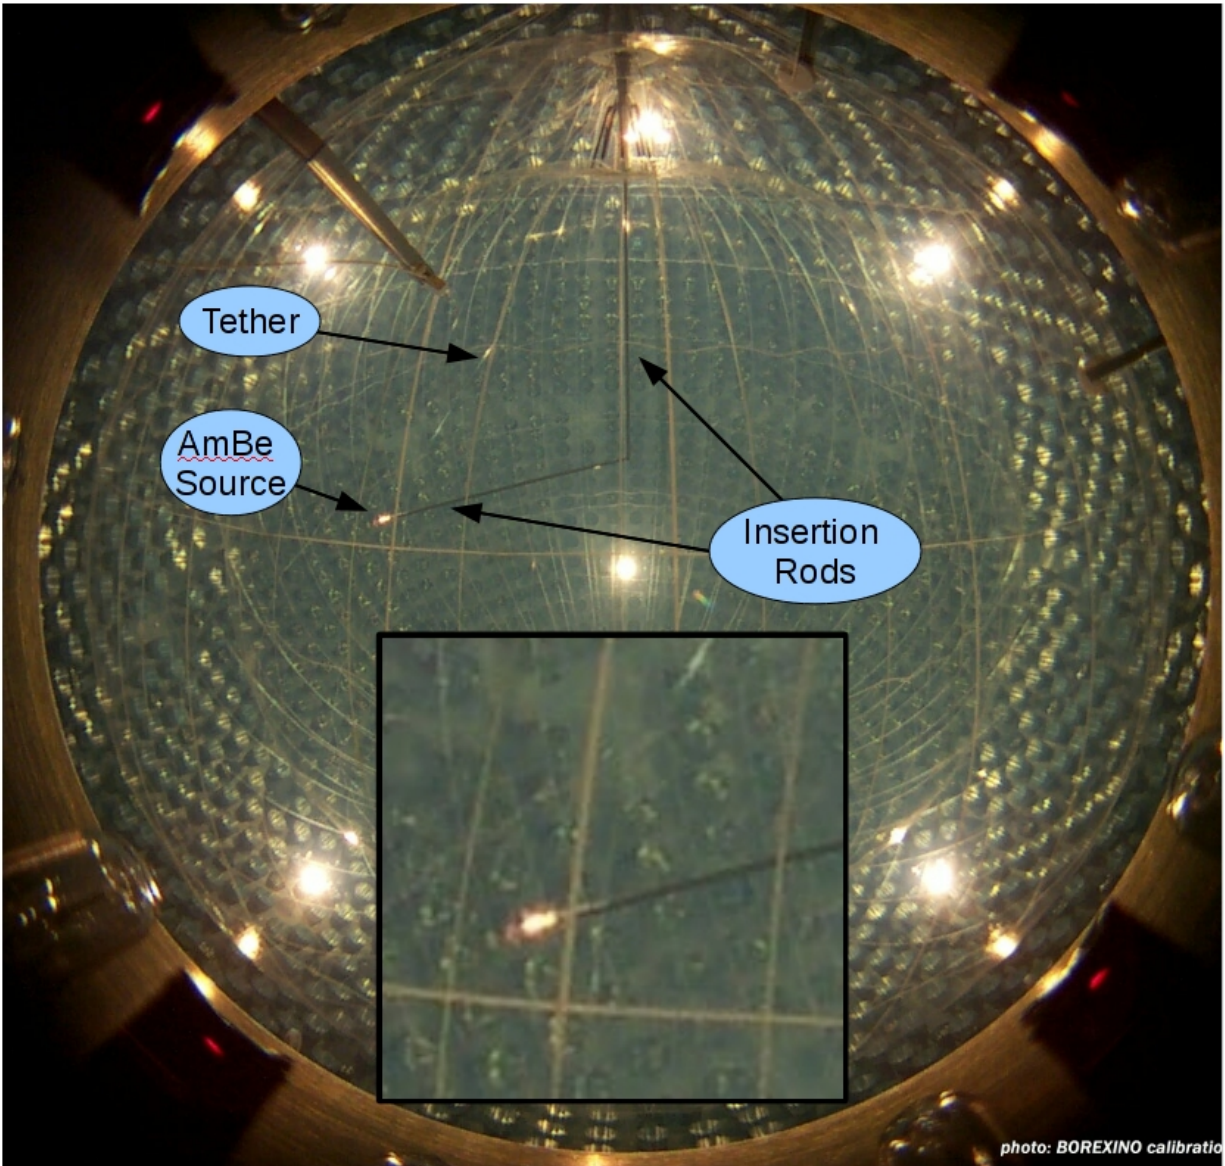


Figure 8.4: AmBe neutron source deployed at the conclusion of the calibration in June 2009 as seen by camera 1. The insertion rods are clearly visible entering the detector through the inner vessel fill pipe — the tether tube is visible at several locations where the halogen lights reflect from it. The inset is zoomed-in to better show the source — the illuminated area is a combination of the red laser light and reflection off of the source coupler.

8.2 Radioactive Contamination Introduced During Calibrations

The Borexino detector started data taking in May of 2007. However, the first calibrations did not occur until a year and a half later — the primary apprehension was due to the potential for contamination. The background levels achieved after filling were better than anticipated (by an order of magnitude in some cases, see Table 3.6) and it is not unreasonable to assume that an invasive operation such as a calibration could spoil that.

In the following pages, the results of analyses of the ^{238}U , ^{232}Th , and ^{85}Kr backgrounds are presented. In addition, some data are presented regarding the appearance, and growth, of an unidentified background component since the initial filling of the detector.

8.2.1 ^{238}U Contamination

The decay chain of ^{238}U was presented in Figure 3.33, the taggable feature of which was discussed on page 107. Briefly, this chain can be tagged with high efficiency by exploiting the fact that ^{214}Bi beta decays to produce ^{214}Po which has a half life of $164.3\mu\text{s}$ which then decays by alpha emission to ^{210}Pb . Since the ^{222}Rn in the calibration sources deployed during these calibrations produces the same coincidence, we can exploit the calibration data to obtain high statistics samples from a precise location in order to extract some of these efficiencies.

If we assume a light yield of 500 npe/MeV , and the ionization quenching formula given in equation 3.20, then we can assume that the ^{214}Po alpha peak will have its mean approximately at $0.745\text{MeV} \cdot 500 \frac{\text{npe}}{\text{MeV}} = 373\text{npe}$. For a perfect detector, we would also expect the width of the Gaussian to be given by the square root of the observed number of photoelec-

trons, which is approximately 20 npe. Thus, if we set our energy cut for ^{214}Po at 300 – 500 npe, we allow significant room for non-ideal behavior of the detector and variations from our assumed light yield.

The efficiency of the energy cut on ^{214}Bi is a slightly more complicated matter since the decay scheme is very complex and is best determined via Monte Carlo. The Borexino Geant4 simulation `g4bx` was used to simulate 50000 ^{214}Bi events which were then processed with the Borexino electronics simulation code (`bxelec`). Finally, the result of the simulation was processed with *Mach4* to yield the spectrum presented in Figure 8.5. In order to veto false coincidences with ^{14}C we apply an energy cut for identification of ^{214}Bi of 100 – 1600 npe. The fraction of the simulated spectrum which lies above the 100 npe energy cut is 99.3%.




Figure 8.5: Simulated energy spectrum of ^{214}Bi processed with the Borexino electronics simulation code followed by reconstruction with *Mach4*. From the simulated spectrum we can determine the efficiency for a [100, 1600] npe energy cut for ^{214}Bi identification — integration of the spectrum in the shaded area indicates an efficiency of 99.3%.

Finally, the decay of ^{214}Po proceeds with a mean life of $237\mu\text{s}$ — however, fast coincidences like those from the mass-212 BiPo decay would also be partially identified by these cuts, so we use a time cut of $1.5\mu\text{s}$ — 1ms, the lower end of which is 3.5 times the mean life for the decay of ^{212}Po . The efficiency of the time cut is 97.9%.

The other cuts which we wish to impose are distance cuts between the Bi and Po candidate events, and cuts on the alpha beta discrimination variables — these are best evaluated with calibration source data. Figure 8.6 depicts the distance between candidate ^{214}Bi and ^{214}Po events identified with the cuts listed in the previous paragraph for a calibration source run taken at the center of the detector (an additional cut was made such that all candidate events had to lie within one meter of the source position). From simply counting the fraction of the events below a given cut value we can determine the acceptance of such a cut — making a cut at 80 cm retains 98.5% of the coincidence events.




Figure 8.6: Histogram of the distance between candidate ^{214}Bi and ^{214}Po events tagged by cuts on the energy of, and time interval between, the two events. The data set was from a run taken with a ^{222}Rn source deployed at the center of the detector (Run 9462). From this histogram we are able to determine the efficiency of a cut on the distance between two events which are known to originate from approximately the same location — for example, for a distance cut of 80cm, we retain 98.5% of all of the events.

For the α/β discrimination cut, we make a plot of the Gatti parameter (see Section 3.3.2) for ^{214}Po and ^{214}Bi candidate events and fit the result to a sum of two Gaussians — an example of a histogram and fit are shown in Figure 8.7 for the case of ^{214}Po . The presence of two peaks is not surprising since the ^{214}Po energy cut region is a subset (about 10%) of that for the Bi, so, for a very active source like the one used to provide this data, some

“contamination” of the ^{214}Po Gatti plot is expected, and vice versa. For the ^{214}Po Gatti histogram, the fit to the larger peak (comprising the alpha particles) yields $\mu = 0.04041$ and $\sigma = 0.01122$ — thus a cut on the Gatti parameter in the region $[0, 0.01]$ would encompass 99.98% of the α 's, and includes less than 3% of the β 's. The same analysis applied to the ^{214}Bi Gatti distribution (plot not shown) yields an efficiency for a cut in the region $[-0.01, 0]$ of 99.98% and a misidentification of Po events of just over 3%.




Figure 8.7: Histogram of the Gatti parameter distribution for events identified as ^{214}Po candidates for a ^{222}Rn calibration source at the center of the detector. The fit (green line, visible in the valley between the two peaks) is a simple sum of two Gaussians — one to represent the ^{214}Po events (red), and the other to represent the distribution of ^{214}Bi events (blue) that passed the identification cuts. Since the energy cut for ^{214}Bi includes the region used to identify ^{214}Po some misidentification is expected, especially for a very active source.

If we now use these cuts to search for coincidences in the valid, non-calibration, runs between May 2007 and October 2009 (livetime of 589.74 days) in the entire fiducial volume, we can use the acquired energy spectra to determine the efficiency of our energy cuts and thus the total efficiency. We select coincidences with the added restrictions that the ^{214}Po candidate lies within the fiducial volume and the ^{214}Bi candidate lies within 3.2 meters. Plotting the ^{214}Po energy distribution and performing a Gaussian fit yields the results presented in Figure 8.8 — integration of this Gaussian within our energy cuts yields an efficiency for the cut of

99.9%.




Figure 8.8: Energy histogram of ^{214}Po candidate events found during the 589-day data set comprising the time period from May 2007 to October 2009. The events are restricted to those that lie within the fiducial volume. Integration of a Gaussian with the indicated mean and sigma yields an efficiency for the [300, 500] npe energy cut of 99.9%.

Thus, the total efficiency for our tagging method is given by the product of all of those discussed in the preceding paragraphs: timing (97.9%), ^{214}Po energy (99.91%), ^{214}Bi energy (99.3%), Gatti cuts for both events (99.98% each), spatial separation (98.5%), and finally, that neither event is a muon (99% each) — the total tag efficiency is then 93.7%.

The result of the ^{238}U contamination analysis script are plotted in Figure 8.9 for those events which fell within a 3m cut — the average level of contamination over this entire time is 2.21×10^{-17} g/g. The contamination level for various time periods relevant to this work are presented in Table 8.1. From looking at the figure and table, we can deduce that the calibration campaigns did increase the contamination of ^{238}U inside of the fiducial volume, however, it did so at a level much less than the refilling campaigns did. Moreover, the calibrations did not cause long term increase in the ^{238}U contamination as can be seen by

comparing the first and last entries in the table¹.




Figure 8.9: The rate of Bi-Po ^{214}Po coincidence events (per week) occurring within a 3.0 meter radial cut from the start of data taking in May 2007 until October of 2009. The thick red lines indicate inner-vessel refilling operations, and the green shaded boxes indicate calibration campaigns. From this plot it can be seen that the calibration campaigns do temporarily increase the number of ^{214}Bi - ^{214}Po coincidences, but they do so at a level markedly less than the refilling of the inner-vessel and the rate is not permanently affected. It should be noted that the y axis shows just the number of coincidences found, and is not corrected for the 93.7% tagging efficiency.

The fact that the contamination level has not permanently increased indicates that any contamination introduced was unsupported (i.e. was not introduced as atoms of ^{238}U), and is mostly due to emanation of ^{222}Rn from the surface of the rods, tether, source, etc. It is true that for each ^{222}Rn atom which decays in the detector, it will eventually lead to a decay of ^{210}Bi — the daughter of the long lived ($t_{1/2} = 22.3$ yr) ^{210}Pb — which makes the CNO analyses more difficult as the spectra are practically indistinguishable. However,

¹Here we are assuming secular equilibrium throughout the ^{238}U chain — the ^{214}Bi - ^{214}Po tag really only tells one the contamination due to isotopes below ^{222}Rn .

Operation	Time Period (m/d/y)	Contamination (10^{-18} g/g)	Error (10^{-18} g/g)
“Stable” Period 1	7/26/07 – 1/26/08	5.05	0.711
On-Axis Calib.	10/05/08 – 10/10/08	46.6	19.0
After On-Axis Calib.	10/10/08 – 10/23/08	14.3	3.81
Off-Axis Calib. 1	1/16/09 – 2/5/09	90.6	16.5
After Off-Axis Calib. 1	2/10/09 – 5/25/09	13.7	1.36
Off-Axis Calib. 2	6/15/09 – 6/28/09	< 6.36 (95%)	
After Off-Axis Calib. 2	7/1/09 – 7/13/09	33.8	6.76
Off-Axis Calib. 3	7/13/09 – 7/29/09	35.5	20.5
After Off-Axis Calib. 3	8/4/09 – 10/4/09	8.44	1.27
“Stable” Period 2	8/20/09 – 10/4/09	5.12	1.18

Table 8.1: ^{238}U contamination levels (corrected for the tagging efficiency of 93.7%) inside of the FV in different time periods. The data presented for the time periods after the calibration campaigns only extend as far as the next refilling. The first entry represents a “stable” period for the contamination: after the initial spike due to the refilling in May of 2007, but before the water loop began in late 2007 (the increase in contamination after this point is likely due to convective mixing caused by the water loop). The last data point is for all of the data starting from three weeks after the last calibration campaign — we note that within the error of the measurement, these data are identical with the first stable period. Thus, the calibrations did not introduce any long term increase in the ^{238}U contamination. This supports the hypothesis that the contamination is unsupported and is likely in the form of ^{222}Rn emanation from the inserted components. The contamination quoted for the calibration periods is obtained from runs taken during the calibration campaign when no source was inserted — usually while changing a source.

we highlight that the refilling operations are a much larger source of these events than the calibrations were.

Some elaboration of some of the features in Figure 8.9 are presented below:

- The main filling with scintillator was completed on May 15, 2007, and we can see the initial rate decreasing for the first few data points
- Scintillator was added to the inner-vessel on May 24 and May 25 of 2007. The scintillator was added to the top of the inner-vessel and after several days it had diffused

into the inner-vessel region, giving rise to the large increase in the ^{222}Rn concentration after these refillings.

- By the middle of the summer, the vast majority of the radon introduced from the filling (and the two volume additions) had reduced to a rather stable level, determined mostly by the emanation of radon from the inner-vessel surface.
- In mid-December of 2007, the water loop was started in order to reduce the haze that had formed on the inner-vessel (see the discussion on page 180 in Section 5.3.1). The heating of the water tank caused convection currents to develop inside of the inner-vessel and thus caused mixing of the inner-vessel scintillator. Thus, the scintillator near the inner-vessel surface with its higher radon concentration was brought into the fiducial volume and induced the increase in the contamination during the winter and spring of 2008.
- The background did increase during, and shortly after, the on-axis calibration. However it reduced to a low level in the two weeks between the end of the calibration and the start of a new batch of refilling campaigns. These refilling operations brought the contamination level up by a factor of three over that introduced by the calibration.
- Following the first off-axis calibration, the contamination increased to a level twice as high as after the on-axis calibration. However, the components inserted in the off-axis calibration have a much larger surface area (by a factor of 35) than the on-axis hardware, and the calibration lasted more than three times as long as the on-axis calibration. Thus, the contamination was vastly reduced relative to expectations based on the duration and surface area of the inserted components.
- Prior to the second off-axis calibration, several refilling operations were performed that introduced a large amount of contamination which was not given sufficient time

to decay before the calibration began. Thus, the contamination measured during and after the calibration is a superposition of that introduced by the calibration and by the refilling.

- After all of the calibrations and refilling were completed, the detector was able to return to a stable state in which the contamination reached a level identical to that in the stable period during the summer of 2007. Therefore, the calibration did not have a lasting impact on the ^{238}U contamination.

8.2.2 ^{232}Th Contamination

The ^{232}Th decay chain (see Figure 3.34 on page 110) actually provides two taggable coincidences — one ^{212}Bi - ^{212}Po (a β followed by an α), and a ^{220}Rn - ^{216}Po double-alpha coincidence². This duality in tagging methods is fortuitous because it affords the ability to probe two different regions of the ^{232}Th decay chain — differences in the results from each method would indicate “unsupported” contamination (most likely introduced as ^{212}Pb).

Tagging Via The ^{212}Bi - ^{212}Po Coincidence

Once ^{212}Bi has been produced, it decays via emission of a β with a Q value of 2.254 MeV and a 64.06% branching ratio to ^{212}Po . The ^{212}Po then decays with a meanlife of 432.8 ns via emission of an α with energy of 8.954 MeV. Following the same approach as in the previous subsection, the ^{212}Po energy is expected to be around 515 npe with an ideal width of 23 npe. With those numbers in mind, we set an energy cut of [425, 600] npe for the identification of ^{212}Po .

²One can exploit the decay of the parent of thoron, ^{224}Ra , to make this a triple-alpha coincidence, although the 5.788 MeV energy of the ^{224}Ra is close enough to that of ^{210}Po (5.407 MeV) to make this coincidence difficult to tag without including numerous false coincidences with ^{210}Po .

The simulated energy spectrum of ^{212}Bi is shown in Figure 8.10 (where we have only considered the β decay channel), and the red shaded area represents the energy region of [100, 1200] npe; a cut in this region gives an efficiency of 94.5%.




Figure 8.10: Simulated spectrum of ^{212}Bi processed with the electronics simulation code and then reconstructed with *Mach4*. The spectrum presented here is only for the β decay channel with a 64.06% branching ratio. Integration of the energy spectrum in the region of [100, 1200] npe yields an efficiency for this cut of 94.5%.

The lifetime of ^{212}Po is so short that the clustering efficiency is too low to make use of the first mean life of the decay — the clustering effectiveness does not reach $\sim 100\%$ until events are spaced at least 400 ns apart. Therefore, the timing cut we apply for the identification of the coincidence is [400, 1500]ns — which gives an efficiency of 36.6%.

Combining these efficiencies with the Gatti parameter and spatial separation efficiencies found in Section 8.2.1 gives a cut efficiency of just 33.37% — multiplying by the 64.06% branching ratio for this channel gives a total efficiency for the tagging method of 21.4%.

The analogous plot to that of Figure 8.9 but for the mass-212 sequence is given in Figure 8.11, and some relevant values are presented in Table 8.2. As with the ^{214}Bi - ^{214}Po case presented in the previous section, the contamination did increase with the calibration and

refilling operations, but was not permanently affected, and there is no evidence that any permanent source of contamination was left in the detector. Moreover, examination of the ^{232}Th decay chain shows that below ^{224}Ra , all daughter isotopes are short lived ($t_{1/2} < 11$ hours), until the stable isotope ^{208}Pb is reached at the termination of the chain, so there is not even a possibility to introduce any long-term contamination if only unsupported ^{212}Bi is introduced.




Figure 8.11: The time evolution of the ^{212}Bi - ^{212}Po contamination in the fiducial volume. The red lines again indicate the refilling operations and the green shaded areas show the calibrations. The data have not been corrected for the counting efficiency. This plot clearly shows the effect of the convection currents induced by the water loop (which began in December, 2007) and the mixing caused by them. From the plot, it is evident that the on-axis calibration (the first green box) introduced quite a bit of ^{232}Th while the others did so at a much smaller level. It is worth noting that the refilling operations also increased the counting rate.

An important thing to check is the expected number of ^{214}Bi - ^{214}Po coincidences tagged by the ^{212}Bi - ^{212}Po cuts. The principle differences in the tag are the energy cuts for the Po

Operation	Time Period (m/d/y)	Contamination (10^{-18} g/g)
Pre-Water Loop	5/15/07 – 12/11/07	< 2.73 (95%)
During Water Loop	12/12/07 – 5/19/08	9.72 ± 3.44
After On-Axis Calib.	10/10/08 – 10/23/08	40.8 ± 23.6 < 106 (95%)
After Refilling #2	10/24/08 – 11/30/08	22.3 ± 11.2 < 51.1 (95%)
After Off-Axis Calib. 1	2/10/09 – 5/25/09	3.62 ± 2.55 < 11.4 (95%)
After Refilling #3	5/31/09 – 6/15/09	50 ± 28.8 < 129 (95%)
After Off-Axis Calib. 2	7/1/09 – 7/13/09	18.0 ± 18.1 < 85.5 (95%)
Since Last Calib.	8/4/09 – 10/4/09	5.12 ± 3.63 < 161 (95%)

Table 8.2: ^{232}Th contamination levels (corrected for the tagging efficiency of 93.7%) inside of the FV in different time periods. The data presented for the time periods after the calibration campaigns only extend as far as the next refilling. From the table it is clear that the ^{232}Th rate does increase following calibrations and refilling operations, however, it falls back down as expected, indicating that no permanent source of contamination was left in the detector.

tagging and, of course, the timing cut. In the previous section we determined that the ^{214}Po events were well-modeled by a Gaussian with a mean of 388.5 npe and a sigma of 28.1 npe. The percentage of this Gaussian that lies within the region of [425, 600] npe (the energy cut for ^{212}Po) is 9.69%. The timing cut efficiency is obtained by integrating an exponential distribution with a time constant of $237\mu\text{s}$ from 400 to 1500 ns, the result is 0.46%. Thus, we would expect $1390 \cdot 0.0969 \cdot 0.0046$, or ≈ 0.6 ^{214}Bi - ^{214}Po coincidences falsely identified as a mass-212 BiPo coincidence — a negligible amount, made even smaller if we were to consider the additional loss of efficiency due to applying the ^{212}Bi energy cut on the ^{214}Bi energy spectrum.

Tagging with the Double Alpha Coincidence

The second possible tagging method for the ^{232}Th chain is provided by the decay of ^{220}Rn (a.k.a thoron) to ^{216}Po — the former decays by emission of a 6.404 MeV α , and the latter decays via emission of 6.906 MeV α with a mean life of 0.209 seconds. We expect quenched alpha energies of 267 and 305 npe with respective widths of 16 and 18 npe, thus, we set an energy window for ^{216}Po of [260, 400]npe and a window for thoron of [210, 340] npe. In addition to the spatial separation (distance < 80 cm) and alpha/beta cuts (Gatti > 0), we set a time window of [0.15, 0.75] seconds, the lower limit of which is set high enough to avoid accidental coincidences from the ^{214}Bi - ^{214}Po coincidence and the double alpha coincidence in the ^{235}U series. The efficiency of the energy cuts are practically 100%, as are the alpha/beta cuts and non-muon cuts — combining all of the cuts gives an overall efficiency of 84.2%.

Applying this search method to the entire data set considered in the previous analyses results in 33 candidate events inside of the fiducial volume. This corresponds to an average contamination inside of the fiducial volume for the ^{232}Th chain of $(1.91 \pm 0.33) \times 10^{-18}$ g/g — which is a factor of ~ 3 less than the average contamination measured by the ^{212}Bi - ^{212}Po coincidence: $(6.36 \pm 1.3) \times 10^{-18}$ g/g!

We have seen from Figure 8.11 that the onset of the water loop caused convective currents to form which brought vessel-emanated ^{232}Th daughters into the fiducial volume. Once produced, thoron decays rather quickly with a half life of about a minute (55 seconds), shortly followed by the decay of ^{216}Po into ^{212}Pb which has a half life of almost 11 hours before decay via the ^{212}Bi - ^{212}Po decays. Thus, it is feasible that the quickly-decaying thoron produced events outside of the fiducial volume, while the ^{212}Pb produced from its decay had a much longer time to continue traveling into the fiducial volume where it decayed — this would cause such an imbalance in the contamination levels measured by the two separate

methods. Because of the energy degradation present for positions far outside of the fiducial volume (see Section 4.2), we do not extend these analyses into the full volume.

8.2.3 ^{85}Kr Contamination Measurement

The most worrisome background for the Borexino experiment is that of ^{85}Kr due to it being a β decay with an endpoint energy of 687 keV, dangerously close to the 662 keV shoulder produced by the ^7Be neutrinos. The decay scheme was discussed briefly in Section 3.4.2 (page 114), but we briefly review it here for completeness. One out of every 231 ^{85}Kr decays proceeds via the so-called “gold plated” decay channel: β emission ($Q = 173$ keV) to a metastable state of ^{85}Rb with a half life of $1.015\mu\text{s}$ which then de-excites with the emission of a 514 keV γ . Determining the ^{85}Kr contamination of the scintillator through this tag is quite difficult due to the very small branching ratio, energy threshold of the detector, fast decay time, and the proximity of the decay energies to those of ^{14}C and ^{210}Po .

Cuts & Efficiencies

To determine the energy cuts for the $Q = 173$ keV β emitted in the decay to the metastable state, the decay was simulated with Geant4 after modifying the data files for the radioactive decay module (RDM) so that only the rare branch would be simulated. The resulting output was sent through the electronics simulation code and finally reconstructed with *Mach4*. The low energy portion of the spectrum is inaccessible due to the detector threshold (~ 50 keV) — and we set the energy threshold even a bit higher because the trigger may not be 100% efficient at that minimum threshold energy. Exploiting data from the low-energy gamma sources inserted during the calibration campaigns allows us to determine the optimum location for the lower energy cut. The ratio of the activities of the ^{57}Co source as measured in Borexino

to that measured in a Ge detector was found to be $> 97\%$ at a mean energy of 43 npe [119]. Thus, we set our ^{85}Kr energy identification cut to be $[40, 120]$ npe — integration of the Monte Carlo generated spectrum in this region suggests an efficiency of 21%.

The 514 keV gamma emitted in the decay of the ^{85m}Rb metastable state will not actually appear at precisely that energy due to the fact that it will Compton scatter one or more times, each time creating a lower energy β which will be quenched according to the well-known Birks parameterization of ionization quenching (equation 3.22 on page 78). Fortunately, the decay of ^{85}Sr populates the same excited state of ^{85}Rb as the gold plated decay of ^{85}Kr thus the same gamma is released in the decay of ^{85}Sr — and we have deployed this source twice in the calibration campaigns. The energy spectrum of all events occurring within 50 cm of the source is shown in Figure 8.12 along with a fit to a Gaussian — the mean of the fit is 212.8 npe and the sigma is 18.6 npe. Placing our energy cuts for this γ at $[160, 280]$ yields an efficiency of essentially 100%.

We also would like to apply an α/β identification cut on the 514 keV gamma. While neither an α nor a β , the gamma imparts all of its energy into electrons, so, it should have a small tail-to-total ratio (TTR) parameter. Again we can exploit the ^{85}Sr gamma source to determine the efficiency of placing a $\text{TTR} < 0.2$ cut on the events. The distribution of the TTR parameter for events occurring within 50 cm of the source is shown in Figure 8.13, and is fit with a sum of two Gaussians. The efficiency of a TTR cut of $[0.0, 0.2]$ for identifying this γ is 98%.

With this decay ($\tau = 1.464\mu\text{s}$) we face the same problem as for the ^{212}Bi - ^{212}Po delayed coincidence in that the time is short enough that we begin to encounter problems with the clustering efficiency if the coincidence gate is opened too early, so, we use the same start time of 400ns and place the stop time at $6\mu\text{s}$ — approximately four mean lives. The efficiency of the timing cut is 74.4%.




Figure 8.12: Energy spectrum of all events occurring within 50 cm of the ^{85}Sr source (deployed at $x=6$, $y=-9.8$, $z=-0.3$ cm, Runs 10577 & 10579). The decay of ^{85}Sr populates the same state of rubidium that the ^{85}Kr decay does, thus, this is the same gamma that we are attempting to tag in the krypton delayed coincidence.




Figure 8.13: The TTR parameter distribution of all events occurring within 50cm of the ^{85}Sr source. The presence of some bulk events with larger TTR values (most likely ^{210}Po) can be seen in the tail on the right, however, they are unlikely to pull the fit appreciably (note the y axis is a log scale). Utilizing an α/β cut in the krypton analysis will allow us to reduce the probability of including false coincidences from ^{210}Po .

Some care must be taken with the spatial separation cut because for events with energies as low as the Kr β that we are attempting to tag, the position resolution is much worse. The position reconstruction algorithm presented in Section 3.3.2 is based on the arrival time of photons and spatial hit pattern of the PMTs, which is a discrete process governed by Poisson statistics, making the error in the position (and therefore the spread of the distribution) inversely proportional to the number of hits. To determine the optimum distance cut to use for the ^{85}Kr analysis, 100 keV betas and 514 keV gammas were simulated and processed with the electronics simulation and reconstruction codes. After processing, the distance between each pair of events was calculated and binned in a histogram. The 80 cm spatial separation cut that was used for higher energy coincidence events now only has an efficiency of 64.9% — widening this cut to 150 cm brings the efficiency up to 91.4%.

Combining all of the cuts gives a total cut efficiency of 14% — multiplying this by the branching ratio gives a paltry 0.0607% efficiency for tagging ^{85}Kr via this coincidence method.

Data Contamination From Other Decays

Since the ^{85}Kr tagging method has such a poor efficiency, the inclusion of one or two “false” coincidences can drastically impact the projected activity, so, we consider the expected level of inclusion of ^{214}Bi - ^{214}Po coincidence events in the data sample, as well as those from ^{212}Bi - ^{212}Po and random coincidences.

For ^{214}Bi - ^{214}Po coincidences, we can expect approximately 2.3% of all such coincidences (with $\tau = 237\mu\text{s}$) to occur inside our time gate of [400, 6000] nanoseconds. The ^{214}Bi energy spectrum does cover the ^{85}Kr beta decay window of [40, 120] npe — fortunately, the large gamma component of the spectrum shifts the bulk of the spectrum towards higher energies such that only 0.88% of the ^{214}Bi spectrum could be considered as a ^{85}Kr candidate. Finally, the ^{214}Po alphas are well-modeled by a Gaussian (see Figure 8.8) which overlaps very minimally with the window we have set for the identification of the ^{85m}Rb gamma ray — the overlap is only 0.005%. Without even considering the TTR parameter cuts that were used for the ^{214}Bi - ^{214}Po events, we expect a contamination of only $1.01 \times 10^{-6}\%$ by ^{214}Bi - ^{214}Po events.

While the mean life of ^{212}Po is much closer to that of ^{85m}Rb than ^{214}Po , the probability of ^{212}Bi - ^{212}Po coincidence events being falsely identified as ^{85}Kr - ^{85m}Rb is basically zero due to the high energy of the ^{212}Po α . The ^{212}Po events identified by the methods presented in Section 8.2.2 are described by a Gaussian with a mean of 495 npe and a sigma of 22.1 npe — the overlap of this Gaussian with the [160, 280] npe energy window for identifying the 514 keV gamma is zero (the means are separated by almost 13σ). Thus, both the mass 212

and mass 214 BiPo coincidences present negligible probabilities of being falsely identified as ^{85}Kr - ^{85m}Rb events.

Finally, we deal with the possibility of random coincidences among events which pass the energy and α/β cuts used to identify Kr and Rb events. To estimate the number of random coincidences, we first count the number of all events in our 592.4-day data set which pass one of the two sets of cuts:

1. “Rb-Like” Events: $\text{TTR} < 0.2$ and $160 < \text{npe} < 280$
2. “Kr-Like” Events: $40 < \text{npe} < 120$

The number of events inside of a 3.5 meter radial cut passing the Kr cuts is 425650432 and the number of events passing the Rb cuts is 745627. The rate of random coincidences is then given by the standard formula: $R_{\text{Rand}} = R_{\text{Kr}}R_{\text{Rb}}\Delta T$. However, we also need to multiply by the probability that two events, each distributed uniformly inside of a 3.5 meter sphere, will reconstruct within 150 cm of each other. To evaluate this quantity, a simple simulation was written that generated 10,000 points uniformly within such a sphere and calculated the distance between each pair of points — the fraction enclosed by a 150 cm cut is 6.13%. Thus, we would expect a total of 2.1 chance coincidences in this entire 592.4 day dataset.

As a cross check of the above result without using the geometrical efficiency correction we can consider that each candidate Kr event effectively opens a spherical volume for $5.6\mu\text{s}$ during which the occurrence of any Rb-like event would be identified as a coincidence. In other words, each ^{85}Kr event gives rise to a sensitive “4-volume” of $8.02 \times 10^{-10} \text{day} \cdot \text{ton}$. The total 4-volume opened by all Kr candidate events during the livetime under consideration is $0.341 \text{day} \cdot \text{ton}$. The rate of Rb-like events is just 8.02 per day per ton, so, the total number of expected accidental coincidences can be determined by multiplying this value with the total

4-volume: $8.02 \frac{1}{\text{day} \cdot \text{ton}} \cdot 0.341 \text{day} \cdot \text{ton} = 2.7$ — in good agreement with the 2.1 determined by the other method.

Likelihood Calculation

An additional diagnostic tool that is useful for identifying the presence of false coincidences in the data sample is to calculate, for each coincidence, its likelihood. Simply, the likelihood, \mathcal{L} is given by the product of the individual probability distributions for each variable used for identification — although typically one determines $-\log(\mathcal{L})$. The distributions of the individual variables used in the cuts are:

- The values for the Rb energy and Gatti parameter are assumed to be Gaussian, thus the $-\log \mathcal{L}$ value for a Gaussian distributed variable is:

$$-\log(\mathcal{L}_{\text{gaus}}) = \frac{(x - x')^2}{2\sigma_x^2} \quad (8.1)$$

For the Rb energy cut we use the mean and sigma reported in 8.12, and for the TTR parameter we use the fit values given in 8.13.

- The time between the candidate events should be distributed by the standard exponential decay equation and the $-\log$ of it is given by:

$$-\log(\mathcal{L}_t) = \frac{\Delta T}{\tau} + \log(\tau) \quad (8.2)$$

And the $1.464\mu\text{s}$ value for τ is used here.

- The spatial separation between the events is distributed according to a spherical Gaus-

sian (r^2 times a Gaussian), thus:

$$-\log(\mathcal{L}_d) = \frac{r^2}{2\sigma^2} - 2\log(r) \quad (8.3)$$

From a fit to the results of the simulation of the low-energy β and 514 keV γ described on page 311 we set the σ for the spatial separation histogram to be 30cm.

- Finally, the likelihood of the ^{85}Kr beta decay energy is:

$$-\log(\mathcal{L}_{85Kr}) = -\log\left(\frac{p(E)}{p_{max}}\right) \quad (8.4)$$

Where p is the beta decay distribution of ^{85}Kr convolved with the detector response function: the negative log of the probability of a given energy normalized to the maximum value that the energy distribution attains (such that for the most probable value the term is 0).

Normally, a likelihood-based analysis is used in situations where there are many candidate events in a given time interval and it is necessary to choose the events which best match the criteria for the coincidence — for instance, [57] uses it to attempt to tag all of the events of the ^{238}U chain from ^{222}Rn down to ^{214}Po . Fortunately, the short meanlife of ^{85}Kr means that we have very few situations in which there are more than one possible ^{85}Kr candidate to match to a ^{85m}Rb candidate (or vice versa), and thus the likelihood method is used to try to identify false coincidences. In the previous section, we estimated 2.1 – 2.7 chance coincidences in this time interval.

Results

Running the analysis over all 592.4 days of livetime (runs 5000-11304) that were available from *Mach4* cycle 97 at the time of this writing yields 16 candidate events. The energy distributions of the two species of events are presented in Figure 8.14 and 8.15. The energy of the Rb candidates is overlaid onto the data from the ^{85}Sr γ source since both decays produce the same γ — fortunately, the match is very good. Figures 8.16 and 8.17 show the coincidence-time spectrum and the correlation between reconstructed position and likelihood respectively. The fact that the point with the lowest likelihood reconstructs near the origin of the detector is important because many of the problematic events — such as electronics noise (which if it appears on all channels at the same time would be reconstructed at the origin) or events occurring in the IB due to the leak — reconstruct near this location as well.




Figure 8.14: Energy distribution of the 16 candidate events. While the statistics are too low to perform any sort of fit to the expected ^{85}Kr spectrum, it is an almost monotonically decreasing function over much of its range and the candidates seem to display approximately this trend. The event near the endpoint of the spectrum is the one with the smallest likelihood.




Figure 8.15: The energy distribution of the ^{85m}Rb candidates are displayed in red on top of the data obtained with the ^{85}Sr source at the center of the detector. The Rb data points have been scaled so that they are visible on the same plot. The agreement between the two data sets is worth noting.

Some relevant data on the candidate events is presented in Table 8.3. The sixteen coincidences observed during the data set indicate an average ^{85}Kr contamination of 28 ± 7 counts




Figure 8.16: The reconstructed time between the Kr and Rb events. Again, the statistics are too small to perform a fit. The point with the longest ΔT is also the one with the smallest likelihood, further suggesting that it might be a noise event.




Figure 8.17: This scatter plot shows the average reconstructed radii of the events versus the likelihood of the coincidence (multiplied by an arbitrary constant). The fact that the point with the lowest likelihood reconstructs so close to the origin is further evidence that it is probably noise as this is where many of the electronics noise and leak events tend to reconstruct.

per day per 100 tons.

The first calibration of the detector started on run 8793, so prior to the calibration we have 12 coincidences in 371 live days for a contamination level of 33.6 ± 9.7 counts per day per 100 tons.

In the 221 live days since calibrations first began, we have observed only four candidates — giving a central value for the contamination of 18.7 ± 9.38 counts per day per 100 tons. However, four events represent very low statistics, and, the 90% confidence level upper limit on the contamination since the calibrations have occurred is < 37.5 counts per day per 100 tons. While the statistics are simply too low to draw any firm conclusions, it certainly appears that the four calibrations did not introduce any additional ^{85}Kr .

Run	TrigId	E(Kr) (npe)	R(Kr) (m)	E(Rb) (npe)	R(Rb) (m)	ΔT (ns)	Δx (m)	Likelihood (a.u.)
5203	425320	45	2.0	206	2.2	992	0.7	64
5290	266820	100	3.4	229	3.1	2880	1.1	60
5494	158115	57	2.3	186	2.9	1152	0.8	63
5523	112277	45	1.1	209	1.6	4160	1.1	82
6324	353861	77	2.6	183	2.9	2400	0.8	53
7213	22087	50	1.9	191	2.6	4544	1.0	73
7295	530250	57	2.8	249	2.7	1504	0.4	53
7362	373531	55	1.0	214	0.8	3712	1.1	77
7887	273619	49	2.5	225	1.8	1312	0.8	64
8236	281709	48	3.5	220	3.4	1504	0.9	70
8245	562345	45	2.9	236	2.8	1920	0.1	58
8408	178089	114	0.1	255	0.1	5056	0.2	33
9332	650316	49	3.2	243	3.2	864	0.6	61
10534	90125	44	1.8	207	2.7	1152	1.2	86
10629	281227	46	1.8	193	1.1	448	1.5	99
11061	143408	47	2.8	203	2.5	1856	0.8	67

Table 8.3: The candidates from the ^{85}Kr analysis performed in *Mach4*. The information, from left to right is: the run in which the coincidence occurred, the trigger id of the coincidence, the energy and radius of the Kr event, the energy and radius of the Rb candidate, the time between the two events, the distance between them, and finally the likelihood.

8.2.4 Contamination In The ^7Be - ^{11}C Valley

In performing the ^7Be neutrino analysis for a forthcoming publication, it was noted that the count rate of events occurring in the valley between the ^7Be ν 's and the start of the ^{11}C beta decay spectrum (offset by the two 511 keV gammas from the capture of the positron) has been increasing. This valley is important as it sets the relative height of the $^{210}\text{Bi}+\text{CNO}$ component in the spectral fit, and can be used to determine the gamma light yield based on the starting point of the ^{11}C spectrum. The only spectral component in this region that is allowed to fluctuate is ^{210}Bi , and is the prime candidate for the source of this contamination.

^{210}Bi is produced from the decay of the long lived ($t_{1/2} = 22.3$ years) isotope ^{210}Pb as part

of the ^{238}U chain — once produced, it decays predominantly via emission of a β^- with a Q value of 1.162 MeV and a half life of 5.013 days[80]. Coincidentally, the neutrino emitted in the decay of ^{13}N in the CNO cycle is produced with a continuous spectrum up to 1.199 MeV — the end point is only 30 keV from that of ^{210}Bi and thus the two components cannot be fit separately and we can only extract the sum of the two. The ^{210}Bi acts as the “noise” on top of the CNO “signal” thus increasing the statistical error on the CNO flux determination in accordance with equation 3.35. In the region below about 700 keV, the spectrum is dominated by the ^7Be ν 's and the β decay spectrum from ^{85}Kr ; in the region above about 850 keV, the spectrum is dominated by the β^+ spectrum from the decay of ^{11}C — thus, the relative height of the $^{210}\text{Bi}+\text{CNO}$ contribution in the spectral fit is largely determined by the number of counts in the 150 keV wide valley region³.

To investigate the time evolution of events in this region, the data were divided into periods corresponding to various detector operations and the number of events in the energy region of [350, 425] npe was recorded. The raw data was processed with the ^7Be analysis code which utilizes a fiducial volume with a 3 meter radius and with the two poles removed by requiring $-1.8 < Z < 1.8$ meters. The data were also filtered to remove noise events (marked by a large fraction occurring in one particular crate, or a very asymmetric distribution of hits), vetoing all events occurring within an 80 cm radius sphere surrounding the mean position of any ^{214}Bi - ^{214}Po candidate (thus removing all components of the ^{238}U chain up to ^{222}Rn), and vetoing muons and all events occurring within 2 ms of a muon. Finally, statistical alpha subtraction was performed. A *simulated* energy spectrum is presented in Figure 8.18 where the red shaded region indicates the “valley” region where this analysis is performed.

Since the decay of ^{210}Bi is followed by decay of ^{210}Po , any increase in the rate of one should be accompanied by an increase in the rate of the other. To that end, the rate of alpha like

³The pep neutrinos also produce counts in this region, but due to its monoenergetic nature, the spectrum is basically flat in this region.




Figure 8.18: A simulated energy spectrum and spectral fit. It is important to note, this is not real Borexino data. The red rectangle extending from [350, 425] npe (or [700, 850] keV) shows the region which is defined to encompass the “valley” events.

events occurring inside of the fiducial volume in the energy range of [110, 300] npe was also tabulated and plotted together with the valley event rate. The results of the valley event evolution study are shown in Figure 8.19 — the ²¹⁰Po rate was scaled so that it fit on the same plot, the y axis thus only holds for the valley event rate.

A fit of the ²¹⁰Po rate in the time period of May 2007 up to August of 2008 yields a half life of 135 days — 1.9 σ away from the true half life of 138 days. The faster-than-expected decay rate is indicative of two groups of ²¹⁰Po — one distributed uniformly throughout the bulk and decaying with the expected lifetime, and the other diffusing outwards (in r) or falling out of the fiducial volume, thus giving the appearance of decay with a faster half life. Now, if we disturb the state of the detector by inducing convection currents, mixing should occur and the activity should change provided that the convection currents move at a different rate than the terminal velocity of the falling particulate (or diffusion speed). As Figure 8.19

shows, such an increase in ^{210}Po activity is not seen after the water loop began (orange line) — nor does the decay rate appear to change.




Figure 8.19: Evolution in time of ${}^{210}\text{Po}$ and events occurring in the ${}^7\text{Be}$ - ${}^{11}\text{C}$ valley; some relevant detector operations are also shown. The y-axis scale only holds for the valley events, as the ${}^{210}\text{Po}$ was scaled such that it would fit on the same plot.

However, starting with the water loop, the rate of events in the valley *does* start increasing, and increases at a fairly constant rate during the time before the calibrations and refillings were performed. Some general comments about Figure 8.19 follow:

- The valley event contamination does not fit particularly well to an exponential, nor a straight line. The first three points do fit well to a straight line, as do the fourth and fifth points (which is expected because there are only two points) — interestingly enough, the slopes of these two lines are identical to within the errors of the fits, which suggests that the rate of increase of the activity was constant over that time period.
- After the on-axis calibration and refillings in October of 2008, the valley event rate jumped upwards, but then continued to increase at the same rate as before these events. This seems to suggest that the refilling or calibrations added a parent of the decays occurring in the valley.
- Unfortunately, the proximity in time of the final refillings and last two calibrations do not permit us to see if the count rate continues to grow at the same rate that was observed in the other two time periods. However, it is very important to note that after the first off-axis calibration (which was the longest), the rate did not increase any more than was expected based on the rate increase seen in the first period. This seems to suggest that this calibration had no effect on the rate with which the contamination in the valley was increasing.
- The last two points are a bit curious as they seem to be rising at a steeper rate than before, although the large error on the point between the calibrations leaves some room for adjustment. At the time of this writing, the data for the time since October 2009 has not been processed. Being able to place another few data points on this plot would go a long way in determining the way in which the valley contamination is evolving.

Taken as a whole, the present data slightly disfavors any impact on the evolution of the contamination due to the calibration operations.

- The hypothesis of ^{210}Bi is also a bit suspect as it is hard to convince oneself that a similar increase in the ^{210}Po rate has occurred over the lifetime of the study. Unless the ^{210}Po atoms produced by the decay of ^{210}Bi are being swept out of the fiducial volume at the same rate in which they are produced, we would expect to see the ^{210}Po points start to bend upwards. Visual inspection and fits do not support this either.

Regardless of the actual source of the events residing in the ^7Be - ^{11}C valley, it is obvious from Figure 8.19 that the contamination has been increasing long before calibrations were ever attempted. There was only one calibration which occurred in a time where the detector was relatively stable, the data from this period do not indicate any increase in the rate or valley events. Data taken after the other two calibration campaigns do show “jumps” in the rates, however, these points were also taken in close proximity to refilling operations, thus making it impossible to deconvolve the source of the jump. With the present data, it seems unlikely that the calibrations played a significant role in the evolution of this unknown source of contamination.

8.3 Reconstruction of Calibration Data

Even at the time of this writing, the vast majority of the collaboration is still heavily involved in processing the abundance of data obtained during the calibration, tuning their analyses and Monte Carlos and working towards the reduction of the systematic errors — particularly concerning the energy and position reconstruction. An in-depth discussion of all of these results is beyond the scope of this dissertation and we report here some general results on

the position and energy reconstruction in the *Mach4* software.

8.3.1 Energy Reconstruction Performance

The energy reconstruction performance was known to be quite asymmetric in the z axis due to the asymmetric coverage of live PMTs as was discussed in Chapter 4. Probing the *relative* energy response throughout the detector is possible with any source, quenched or not, thus, we choose the ^{14}C - ^{222}Rn source since it was deployed in more locations than any other source. A scatter plot of the peak of the ^{214}Po energy is shown in Figure 8.20 where the energies have been normalized to the energy at the center of the detector. We can also exploit the calibration sources to determine the resolution of the energy reconstruction at the center of Borexino — performing a fit and taking the ratio $\sigma(E)/E$ gives a resolution of 6.6% at ~ 700 keV.

A more relevant source, particularly for the high-energy ^8B neutrino analysis is the AmBe neutron source due to the production of high energy neutron capture gamma rays with energies of: 2.2 MeV (^1H), 4.9 MeV (^{12}C), 7.6 MeV (^{56}Fe), 9.3 MeV (^{54}Fe), as well as numerous others at lower intensities. Deploying this source at several locations inside of the inner-vessel allows one to determine the radial bias as a function of energy (and thus, the fiducial mass enclosed by a hard cut on radius for a given energy), and also the error imposed by a hard cut on energy due to variations in the energy response as a function of position. Such information was very valuable in performing the ^8B and geo-neutrino analyses presented in [44] and [51] respectively.

To provide even more high energy gamma calibration points, studies were performed at VT to find a suitable “target” to place in the detector with the AmBe source; the leading candidate was chlorine. The ^{35}Cl isotope of chlorine is present at a 75% abundance, and




Figure 8.20: Relative energy response of the Borexino detector to a ^{214}Po emitting source deployed during the first off-axis calibration. The Y axis has been normalized to the peak of the ^{214}Po energy distribution obtained from the run with the source at the center (Run 9462). This figure reproduces Figure 4.3 in Chapter 4 where the relative response of ^{210}Po was plotted as a function of Z. The results of this figure show a variation in the energy response of $> 20\%$ in the whole inner-vessel, and $\sim 5\%$ inside of the fiducial volume ($-1.8 < Z < 1.8$ m).

has a thermal cross section for (n, γ) reactions of 44 barn. The capture can occur to a large number of excited states, however, the most common gammas produced have energies of 1.1, 6.1, 1.9, 7.4, 7.8, 6.7, 2.9, 5.7, 5.0, 3.1, and 8.6 MeV [120]. The chlorine is most easily introduced in the form of CTFE (chlorotrifluoroethylene, C_2ClF_3) chlorinated teflon (trade name: Kel-F) which is known to be chemically compatible with the Borexino scintillator. For a 1" diameter rod of this material (2" long), mounted $\sim 5\text{cm}$ below the AmBe source, we can expect approximately 500 capture gammas per hour. In order to provide sufficient statistics to resolve all of the peaks, such a source should be left in one position for at least 24 hours. While calibration with this target was planned for the calibrations in the summer of 2009, due to scheduling conflicts with the source it was necessary to cancel the plans for its use. Fortunately, the capture gammas on the Fe in the stainless steel insertion system

components was sufficient to obtain energy calibration points.

More information on the energy reconstruction performance will be presented in a forthcoming paper on the Borexino Monte Carlo simulation and calibration campaigns.

8.3.2 Position Reconstruction Performance

The studies of the position reconstruction were performed using the ^{14}C - ^{222}Rn source as it was a fully enclosed source in which the particles producing the scintillation light were confined to a small volume (except for some of the higher energy γ 's emitted by ^{214}Bi). Because the position reconstruction algorithm works under the assumption of a point like energy deposition, the gamma and neutron sources do not make good sources for studying the position reconstruction. Fortunately, impurity quenching of a scintillator causes some of the energy to be channeled away through reactions that either do not emit light, or absorb light already emitted — in either case, these processes do not affect the time profile of the emission, so, we can still use the quenched sources from the on-axis campaign to perform position reconstruction studies.

Within the ^{14}C - ^{222}Rn source we have a wide variety of energies to work with, each one likely to have a different reconstruction bias (as was discussed in Chapter 4). The ^{214}Bi - ^{214}Po events occurring in the source are relatively easy to tag via the method exploited in Section 8.2.1, and we choose to study the alpha events due to their limited range and the fact that ^{214}Bi emits many gamma rays in the majority of its decays. While it is true that the pulse shape from α events is slightly different than for β events (because of the larger slow component for the former: see parameterization in Table 3.4), this component is largely confined to the tail of the pulse and does not play a relevant role in the $[-5, 15]$ ns (relative to the peak) time interval in which the position reconstruction is performed.

Due to the large variation in the energy response for a given position within the detector, a strong cut on the alpha energy was not applied. Instead, the data were filtered such that only alpha events occurring within 1ms of another trigger were left, the energy spectrum of these events was fit to a Gaussian. The ^{214}Bi - ^{214}Po tag was then employed with an energy cut on the alpha defined to be the mean of the aforementioned Gaussian plus and minus two sigmas. Two other species of events were considered: the ^{214}Bi events, and the “ ^{14}C ” events which were defined to be any beta-like event in the energy range of [45, 80] npe. For each event in these species, their x, y, and z positions were binned into histograms and then fit to a Gaussian and entered into a table. After adding the “true” position of the source obtained from the cameras, we then compute the difference in each of the three coordinates and can use these data to gather information about the performance of our position reconstruction algorithms.

A very informative plot for examining position reconstruction biases can be made by plotting the camera determined coordinate on the x axis, and the deviation between the reconstructed position and the camera determined position on the y axis. This plot can exhibit several different trends:

- For an ideal position reconstruction, there will be no bias and we would expect simply a flat line colinear with the x axis.
- If the reconstruction is systematically biased in a particular coordinate, it will manifest itself on such a plot by having a y intercept other than zero.
- For a position reconstruction which systematically pushes events outwards, this will be indicated by data points which follow a line with a positive slope. A reconstruction which pulls events inwards would have data points on a line with a negative slope.

Figures 8.21 and 8.22 show the deviation from the camera position in X and Y respectively.

Both plots exhibit essentially the same features: an upward trend of the data points indicating an outward bias of the data and a very slight negative offset indicated by the y intercept of the fit line. The smallness of the bias in the negative coordinate direction is confirmed by fitting a Gaussian to the Y-axis projection of these plots.




Figure 8.21: Plot of the deviation from the true X coordinate for ^{214}Po events occurring inside of the ^{14}C - ^{222}Rn source during the first off-axis calibration campaign. The positive slope of the line indicates that the reconstruction is systematically biased outwards for ^{214}Po energies. Also present on the plot is a best-fit line with a y intercept of less than a centimeter.

The position reconstruction on the Z axis is unfortunately worse as can be seen in Figure 8.23. The data from the calibration suggest a systematic downward bias of 3-4 cm. A possible explanation is that the reconstruction algorithms themselves are flawed — if this is the case, Monte Carlo simulations of the sources deployed in the detector (using the exact detector conditions at the time of the run) would show the same defect when simulated at the “true” location determined by the cameras. The results of these studies seem to vindicate the reconstruction algorithms as the Monte Carlo data are reconstructed in the same location as the source data — however, this argument only holds if the simulation accurately reproduces




Figure 8.22: The same plot as 8.21 but for the Y coordinate. Again, we have a positive slope confirming the outward bias of these events, and again the best fit line shows a y offset of much less than a centimeter.

the detector, which it does, but only by using some non-physical factors.

It is also beneficial to look at the response of the position reconstruction algorithms for events with different energies and particle types. The same types of plots that were presented for ^{214}Po are shown for ^{214}Bi and ^{14}C calibration source events (only the Z axis deviations are shown, the X and Y are the same as for ^{214}Po) in Figures 8.24 and 8.25.

Indeed, the outward bias and the deviations are larger for the higher energy ^{214}Bi events as can be expected by the fact that higher energy events produce more photoelectrons per PMT. The arrival time spectrum for multi-pe events is biased towards earlier times (see Equation 3.31) — earlier times would mean an event occurring closer to the PMTs (i.e. at a larger radius). It is important to note here that the offset on the Z axis has remained the same as for the ^{214}Po data, thus, the systematic downward bias in the data does not appear to be energy dependent.




Figure 8.23: Plot of the deviation from the true z coordinate of the source. The Z axis displays a larger offset than the X or Y coordinates, indicating a systematic downward bias in the data of a few cm. Again, the upward trend of the points means that ^{214}Po events are pushed radially outwards.

The distribution for the low energy ^{14}C events is much less structured than the ^{214}Bi and ^{214}Po events considered thus far. It is clear that there is a universal reconstruction of the positions below the true source position indicating an approximate downwards offset of 3-4 cm, and it might be possible to convince oneself of a slight inward bias (negative slope) of the data.

By looking at the widths of the Gaussian fits to the three coordinates for the different particle types, we can determine the resolution of the position reconstruction for events of various energies. The results are given in Table 8.4. Plots of the position and energy distribution for a source approximately at the center are given in Figure 8.26.




Figure 8.24: Plot of the Z axis deviations for ^{214}Bi events in the January 2009 off-axis calibration. While the scale of this figure is larger than that of Figure 8.23, the general trend is the same. Filtering the data to use only those ^{214}Bi events with energies reconstructed in the lower energy, flatter, portion of the ^{214}Bi decay spectrum yields a plot with smaller vertical deviations.

8.3.3 Attempts to Resolve the Z offset

The results of the calibration have highlighted that the position reconstruction is very accurate in the X and Y directions, but there is a systematic discrepancy between the camera-determined positions and those determined from the position reconstruction. Several studies were performed to try to determine what the cause of the Z offset is and how to deal with it; unfortunately, none of the studies found a reason, but several potential explanations were eliminated from contention. A possibility not discussed in detail in the sections that follow is that the temperature gradient along the Z axis could affect the propagation of the light, however, this hypothesis has not yet been tested.




Figure 8.25: The Z axis deviations for ^{14}C data. It is difficult to extract any trend from these data, except that they are uniformly reconstructed downwards, perhaps with a weak inward bias.

Asymmetric PMT Distribution

One of the hypotheses for the Z offset is that it is a result of the strong asymmetry in the number of live PMTs in the Z coordinate. One way to test whether this is in fact an effect is to force the distribution of PMTs to be symmetric. This procedure is accomplished by looping over each PMT and finding the PMT on the opposite side of the detector from it by projecting a ray from the center of each PMT, through the origin, and finding which

Event Type	Energy (npe)	σ_x (cm)	σ_y (cm)	σ_z (cm)
^{14}C	45-80	31.8	31.1	34
^{214}Po	250 - 350	13.5	13.5	14.8
^{214}Bi	500-1600	12.1	11.6	12.2

Table 8.4: Position reconstruction resolution for a ^{14}C - ^{222}Rn source located at the center of the detector. As expected, the resolution gets better at higher energies and is slightly worse for the Z axis.

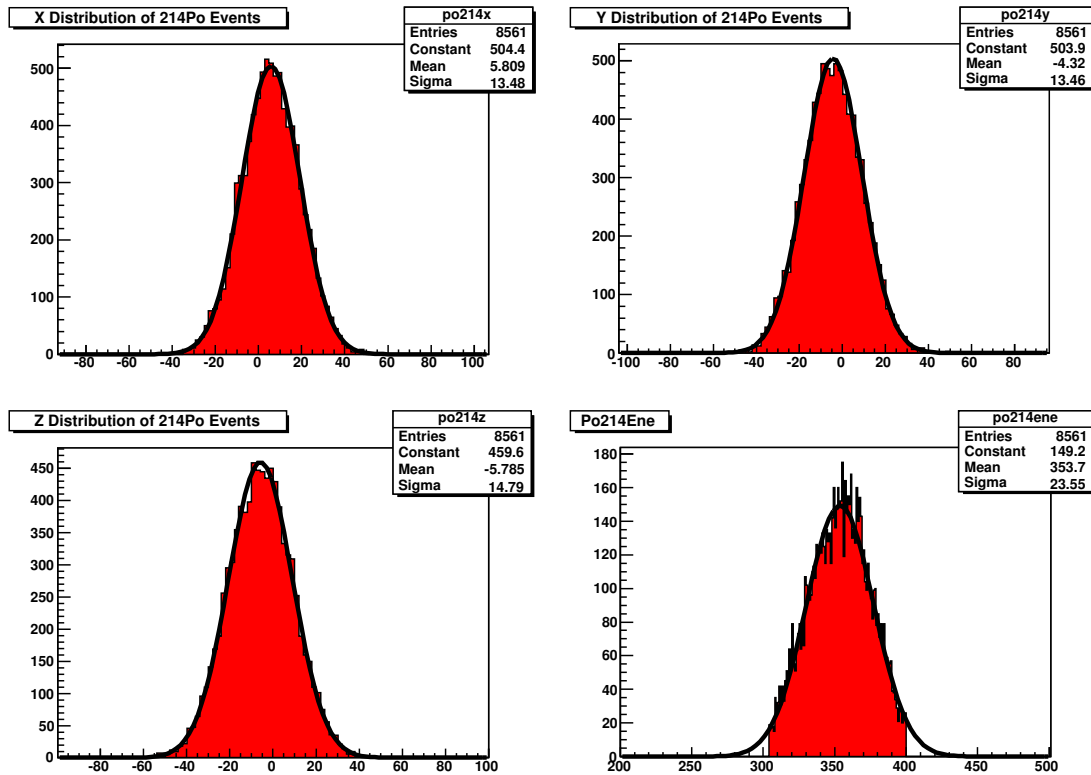


Figure 8.26: Plots and fits to the position distribution of ^{214}Po events in Run 9462 where the ^{14}C - ^{222}Rn source was deployed at $x = 6.8\text{cm}$, $y = -3.9\text{cm}$, $z = 0.1\text{cm}$. The bottom right panel shows the energy distribution (in npe) of the ^{214}Po events.

PMT it intersects on the opposite side of the detector. Due to the presence of the camera housings, scaffolding supports, sphere legs, and other structural members, there is not always a PMT on the diametrically opposite side of the detector. For each PMT or channel which is determined to be dead for timing purposes during pre-calibration, we then manually set a flag to tell the rest of the reconstruction modules to ignore the opposite PMT completely. Plots of the distribution of dead PMTs after this procedure are shown in Figure 8.27.

To test the efficacy of this procedure, several source runs were reconstructed⁴ after forcing the PMT distribution to be symmetric as discussed in the preceding paragraph. The results

⁴This study was performed with *Mach4* cycle 93 (read: 9.3); at the time of this writing, the production version of *Mach4* is cycle 96, however, the position reconstruction algorithms have not been modified since cycle 93, so, the results of the comparison are still valid.




Figure 8.27: Plots of the X, Y, and Z coordinate of PMTs not used in the reconstruction after manually forcing *Mach4* to ignore the PMT directly opposite to any legitimately dead PMT. The distributions are not completely uniform due to the fact that some PMTs do not have one directly opposite to it. This procedure was done to test if the asymmetric distribution of dead PMTs was the source of the systematic offset between the source position determined by the cameras and from the data.

of all comparisons between the Z distribution of calibration source events with and without a symmetric PMT distribution all conclusively show that the position reconstruction does not get better with a symmetric distribution, it in fact gets worse by another $\sim 2\text{cm}$. An example of the output of the comparison can be seen in Figure 8.28 — the fact that the z coordinate gets further from the true position is evidence that the asymmetric PMT distribution is not the leading cause of the observed Z offset.




Figure 8.28: Comparison of the Z coordinate distribution of data from calibration source Run 9460 using the normal position reconstruction from *Mach4* cycle 93 (in blue), and the reconstruction performed after forcing the algorithm to use a uniform distribution of dead PMTs (red line). Gaussian fits yield a mean of 96.5cm for the normal reconstruction and 94.8cm for the symmetric distribution (the true Z coordinate determined by the cameras is 99.7cm).

Camera System Origin

The z-offset could be explained quite easily if the detector and camera system had different coordinate systems — however, this is not the case. From the discussion in Chapter 5, we know that the cameras are in fact calibrated to the PMTs, using the same PMT location data as the reconstruction algorithms use. Moreover, the cameras were all installed with only a rough (by eye) angular alignment of the cameras along the camera housing axis, and thus the cameras do not “know anything” about the detector’s Z axis — there is nothing special about the Z axis for the camera system. Furthermore, the camera-reconstructed positions do not change appreciably when, for instance, cameras 1, 3 and 5 are used to reconstruct the image versus reconstructing it with cameras 2, 6, and 7 — it thus seems *highly* unlikely that an offset would happen to manifest itself in the same direction as the detector’s Z axis

(which it knows nothing of).

Vessel Shape

If the shape of the nylon vessel has any impact on the position reconstruction, it certainly could have caused an effect in the calibration data as the vessel shape before the first off-axis calibration was quite different from that before the second off-axis calibration (see Figure 8.29). As can be seen from Figure 8.29, the refilling operations caused the vessel deformation to have decreased by $\sim 15\text{cm}$ at the points of maximum deviation. Such a concern is feasible, as the main asymmetry in the Borexino vessel is in the Z (or theta) direction. When the plan for the summer 2009 calibration campaign was being formulated, we purposely planned the repetition of roughly fifteen points that had been done in the first off-axis calibration in January of 2009.

Comparing Run 9460 from the first off-axis calibration in which the source was placed at $(x, y, z) = (6.8, -3.9, 0.1)\text{cm}$ to Run 10309 from the second off-axis calibration when the source was placed at $(x, y, z) = (7.8, -9, 0.8)\text{cm}$. For the January run, the source reconstructs 5.8 cm lower than the camera suggests, and for the June run, it reconstructs 5.5 cm lower — the difference is well within the 2 cm error of the camera system. Similar comparisons for other source positions indicate the same level of difference, thus, it is highly unlikely that the source of the Z offset is the asymmetry in the vessel shape.

The shape of the Borexino inner-vessel has also been simulated in the Borexino Geant4 simulation program `g4bx`. The reconstructed position of events generated inside of a spherical vessel do not differ from those generated with a severely deformed inner-vessel shape[121]. This is further confirmation that the vessel shape is not the cause of the systematic offset in the Z direction.

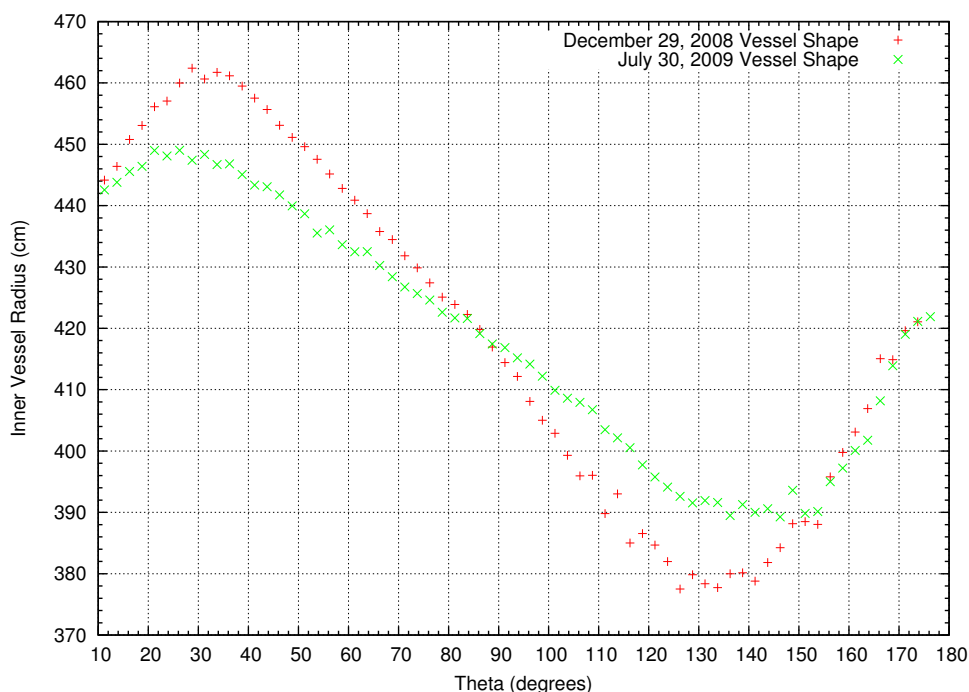


Figure 8.29: Vessel shape profiles before the first off-axis calibration (red “+”), and the shape just after the last off-axis calibration (green “x”). The refilling operations performed in November of 2008 and June of 2009 caused the shape to become more spherical. To see if this caused any shift in the reconstructed source positions, we re-mapped some of the same source positions in the second off-axis calibration campaign that were done in the first campaign.

Timing Calibration

During 2009, two short calibrations were performed with an optical source: a Teflon diffuser sphere emitting light from a 394 nm laser. The goal of this test was to deploy the source at the center of the detector and pulse the laser while triggering the DAQ on its synchronization pulse — since the source is deployed at the origin, all PMTs are equidistant from the source and, if they are properly calibrated in time, should be hit at the same time. This test revealed a small, albeit relevant, difference in the hit times of some groups of channels. Recall that the temporal alignment of the electronics channels are performed with the same laser and synchronization pulse and assume that the light from the laser will reach the ends of the fiber

optics at the same instant — this test indicated a small difference for some fiber bundles. The results of these tests were used to determine additional timing offsets to be applied during the timing calibration. After applying these ad hoc offsets, the Z offset decreased slightly, but was nevertheless still present (and in fact a slight offset was introduced in X and Y)[122]. At the time of this writing, these studies are still ongoing, but are very difficult as they involve determining corrections on the order of a ns or less.

A Position Dependent Scintillator PDF

The final attempt made at resolving the Z offset was to assume that the issue lies in the scintillator PDF itself; it is clear that there is a fault in the PDF because we see a large radial bias that requires us to use an effective index of refraction of 1.7 instead of the true index of refraction of 1.53. Recall from Chapter 3 that the detector simply supplies a list of the PMTs which were hit and at what time they were hit, the latter is given by (neglecting PMT time jitter): $t_{hit} = t_{event} + t_{TOF} + t_{PDF}$. t_{event} is the actual time that the energy was deposited into the scintillator, t_{TOF} represents the amount of time it takes for the light to get from the point where the energy was deposited to the PMT that was hit (given simply by $t_{TOF} = \frac{n}{c}d$), and finally, t_{PDF} is the emission time profile of the scintillator (the PMT jitter time is absorbed here for convenience). The task of the position reconstruction module is to find the position and time, t_{event} , such that $t_{hit} - t_{event} - t_{TOF}$ gives a spectrum that matches the scintillator PDF as closely as possible.

Effects such as scattering cause light to be delayed — this predominantly affects the light seen by the PMTs farthest from the event because the longer the photon spends in the detector, the more probable it is to undergo scattering. The “delay” of the arrival of light tricks the position reconstruction algorithm into choosing a position further out from the actual position — choosing a position further away accommodates the delayed light by pulling it

back into the peak area, thus providing a narrower spectrum for $t_{hit} - t_{event} - t_{TOF}$ than the true position does. This delay due to scattering, and its dependence on distance, can be seen in Figure 8.30 where the time-of-flight-subtracted arrival time of light from a calibration source is shown for PMTs grouped by their distance from the source. It is clear from Figure 8.30 that the PMTs which are further away from the true source position see a smaller component in the peak and more in the tail, indicating a larger component of light which has been delayed by scattering.




Figure 8.30: Time-of-flight-subtracted arrival times (normalized to unit area) for PMTs that have been grouped by their distance to the true event position: ^{14}C - ^{222}Rn calibration source run 9596 — $(x, y, z) = (170.4, 5.7, -36.9)\text{cm}$. For an ideal detector, all of the histograms would be identical. However, the presence of scattering causes a reduction in the prompt peak and an enhancement in the tail of the arrival time. It is clear that one single PDF cannot adequately accommodate all of the histograms shown. It should be noted that we only consider hits in which the PMT detected one photoelectron, thus, the differences are not due to the variation in the arrival time distribution due to multiple occupancy.

The large differences between the arrival times as a function of distance shown in Figure 8.30

show that a single PDF does not fit well to all of the histograms. If we assume then that all of the differences in Figure 8.30 are due solely to the different distance from PMT to the energy deposition point, we can attempt to generate a PDF which depends on the distance traveled. The scintillator PDF used in *Mach4* is given by Equation 3.29, and depends on two exponential time constants $\tau_{1,2}$, two relative weights of the exponential decay components $q_{1,2}$, one time offset t_0 , and one width σ . If we fit each of the histograms shown in Figure 8.30 to the one-photoelectron PDF, we will be provided with one set of the six PDF parameters for each distance bin considered. If we then plot the parameters, one by one, as a function of distance and fit them to a second-order polynomial, the result will be a set of six equations (as a function of distance) which we can use in place of the single values for the PDF parameters.

The parameter fitting process was performed for the ^{214}Bi events recorded during calibration run 9596 for PMTs which are located between 4.5 and 8 meters from the source position (grouped into 0.5 meter bins). Results of the fits to the parameters are shown in Figures 8.31, 8.32 and 8.33.




Figure 8.31: Distribution and fit of q_1 and q_2 as a function of distance. As expected, the PDF for events which occur farther away is better described by a larger component of the long-lived exponential.




Figure 8.32: Distribution and fit of τ_1 and τ_2 as a function of distance. In both cases, we see the expected behavior of longer lifetimes for the two exponentials as a way to account for the larger fraction of delayed light for PMTs located further away from the source.

All six parameters are fit to a second-order polynomial $a_o + a_1d + a_2d^2$, and the results of the fits are given in Table 8.5. The position reconstruction module⁵ was then modified to use these new distance dependent PDF parameters — this new PDF is referred to as the “experimental” PDF. An important difference is that for the experimental PDF, we use the *true* scintillator index of refraction, $n = 1.53$, since the effective index of refraction is no longer necessary because we have absorbed the effects due to the scattering into the PDF parameters.

Several ^{14}C - ^{222}Rn calibration runs were reconstructed using the experimental PDF in *Mach4* cycle 93 and the ^{214}Bi - ^{214}Po tagging script was executed on these runs. The results for ^{214}Po events in run 9562 are shown in Figure 8.34 and in Table 8.6.

⁵As with the symmetric PMT distribution discussed earlier, these studies were also performed with *Mach4* cycle 93.




Figure 8.33: Distribution and fit of the σ of the Gaussian in the scintillator PDF and the time offset of the Gaussian t_0 . The σ parameter obviously does not fit very well to a second order polynomial and indeed appears discontinuous across a 6.5m separation — this behavior is not yet understood.

Variable	a_0	a_1	a_2
q_1	1.077	-0.121	0.003612
q_2	-0.07745	0.121	-0.003612
τ_1	4.574	0.3613	0.00004891
τ_2	51.42	-5.316	0.6563
σ	2.971	-0.2202	0.01898
t_0	-6.084	0.7924	-0.03546

Table 8.5: The three constants used in the polynomial representation of the six parameters of the scintillator PDF. Each of the six parameters (x_i) is represented as a function of distance via the equation: $x_i = a_0 + a_1d + a_2d^2$. The source data and fits from which these data were extracted are shown in Figures 8.31, 8.32 and 8.33.

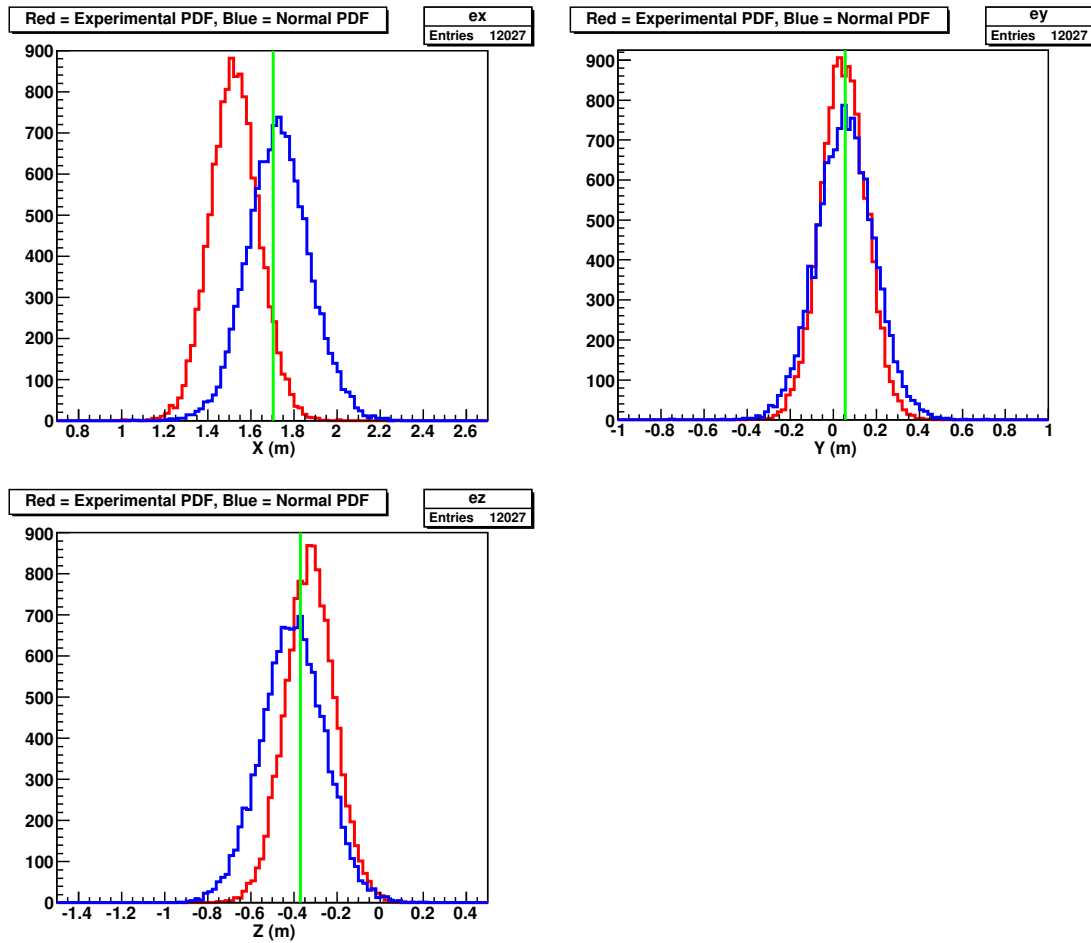


Figure 8.34: The results of a comparison between the normal reconstruction with $n = 1.7$ (in blue), and the “experimental” PDF (in red) generated with distance dependent position reconstruction parameters and an index of refraction of 1.53. The green lines on each panel show the position as determined by the calibration cameras.

As we can see from Figure 8.34 and Table 8.6, the new experimental PDF performs very well for the Y coordinate — however, this was to be expected because it is very close to the origin, and thus, the distance to every PMT is almost identical, so, very little correction is applied in this coordinate. The reconstruction in Z performs almost as well as for the normal PDF, except in the opposite direction. The experimental PDF unfortunately falls apart in the X coordinate with a distribution that is pulled inwards by about 15 cm. Interestingly, the experimental PDF does produce a reconstruction that has a resolution that is better by

	X (cm)	σ_X (cm)	Y (cm)	σ_Y (cm)	Z (cm)	σ_Z (cm)
Cameras	170.4		5.7		-36.9	
Normal PDF	172	14	5.5	13	-40.5	14.9
Difference ($x_{rec} - x_{Cam}$)	1.6		-0.2		-3.6	
Experimental PDF	153	11	4.8	10.8	-32.6	11.6
Difference ($x_{rec} - x_{Cam}$)	-17.4		-0.9		4.3	

Table 8.6: Reconstructed positions and widths for calibration source run 9562 using the normal cycle 93 reconstruction PDF with $n = 1.7$ as well as the experimental PDF with $n = 1.53$. While the experimental PDF does perform comparably well to the normal PDF in the Y and Z directions, the X coordinate is far from accurate.

about 2 – 3 cm in each direction. Tests on other calibration runs show similar results.

Ultimately, the experimental PDF generated in the method discussed in this section is unusable due to its performance in the X direction, nor did it resolve the problem of the Z offset (in fact, it reversed the direction and amplified it). The results point to the possibility that the parameters of the PDF need to be parameterized further, not just by the total distance from the PMTs, but perhaps the distance in Z to the PMT, the distance traveled in the buffer, etc. While these studies have not yet been completed, determining an accurate scintillator PDF which uses the true index of refraction of the scintillator is a primary goal of many collaboration members in the coming months.

8.4 Systematic Errors After Calibration

The previous sections have shown the performance of the reconstruction codes on calibration data, and the impact of the calibrations on the purity of the detector. However, the ultimate goal of the calibrations is the reduction of the systematic errors associated with our neutrino measurements. In what follows, some data and predictions about the systematic

error budgets for the ${}^7\text{Be}$, ${}^8\text{B}$, and geo-neutrino analyses are presented which make use of the results from the calibrations.

The biggest areas of improvement were the errors associated with the position and energy reconstruction. Some of the “second-order” improvements that came out of the calibration were: the timing calibration issue discussed in Section 8.3.3, α/β data for protons (from the AmBe source), acceptance regions for the event isotropy cut, multiple gamma source energies to exploit to determine the k_B value, and several low-energy sources for measuring the trigger efficiency of the detector.

8.4.1 Impact on the ${}^7\text{Be}$ Neutrino Analysis

In the first two publications of the ${}^7\text{Be}$ ν measurement ([83], [20]), the fiducial mass was one of the largest components of the overall systematic error — contributing roughly half of the total systematic error of 6%. If we wish to obtain a 5% total error measurement, this number must be reduced, by an even larger fraction if we aim for a 3% measurement. The origin of the 6% number is due to the method in which the fiducial volume was defined in the early ${}^7\text{Be}$ analyses.

The definition of the fiducial mass in the early analyses was performed by assuming that certain types of events were uniform inside of the entire inner-vessel. If we know that the mass of the entire inner-vessel is M_{IV} , and that the total number of counts of an isotropically distributed source is N_{tot} , we can determine the fiducial mass M_{FV} by finding the radius for which:

$$\frac{M_{FV}}{M_{IV}} = \frac{N(r)}{N_{tot}} \quad \Rightarrow \quad r = R_{FV} \quad (8.5)$$

Thus, the fiducial radius can be determined by counting the event rate inside of different spherical shells until the ratio $N(r)/N_{tot}$ is equal to M_{FV}/M_{IV} . The total mass of the inner-vessel was determined from integrating flowmeters (see Chapter 3) on the vessel fill lines and is known with an accuracy of 0.2%. However, the radial cut required to yield a fiducial mass of 100 tons varied by roughly 6% depending on which type events were used to calculate the ratio $N(r)/N_{tot}$ — this is due to the position dependent energy reconstruction.

The results of the calibration campaign have shown that the position reconstruction in x and y have $\sigma_{x,y} \sim 1.5$ cm, while z is a bit larger: $\sigma_z \sim 2.3$ cm — this translates to an error of about $\sigma_r = 3.1\% = \sigma_V/V$ (by equation 4.1) for the fiducial volume. Moreover, this high degree of accuracy for the position reconstruction means that it is possible to abandon the use of the mass normalization using ^{14}C and simply use the geometric volume. If we use just the ^{14}C in the northern hemisphere (because of the energy degradation in the southern hemisphere), we obtain a fiducial mass of 76.7 tons for a region enclosed by $R < 3.021$ m and $|z| < 1.7$ m. For comparison, the geometric mass enclosed by these cuts is 75.5 tons — a maximum difference of 1.6% [123].

While the 3.1% is a bit higher than one would hope, it is still a factor of two improvement over the error before the calibrations, and it does not preclude a 4% measurement of the ^7Be ν flux if the other systematic errors are kept below 2.5%. A 3% measurement will ultimately require changes to the position reconstruction modules in order to obtain more precise performance. The data with the experimental PDF in Section 8.3.3 did exhibit a smaller σ which implies there is room for improvement in this area.

Because the next iteration of the ^7Be ν analysis is still in its early stages at the time of this writing, the reduction in the detector response systematic has not yet been determined. The most recent ^7Be ν analysis attributed an error of 6% to the variation of the detector response function within the fiducial volume. Since this last measurement, the calibrations

have provided a fine mapping of the detector with sources covering the entire energy range of the ${}^7\text{Be}$ ν analysis. Following the approaches used in the ${}^8\text{B}$ and geo- $\bar{\nu}_e$ analyses, the calibration data will be used to create a position-dependent energy reconstruction function (which may also depend on the observed event energy). Based on results obtained with this procedure in other analyses, a factor of two reduction in the error is likely feasible with the current set of calibration data — further reductions may require a finer mapping of the detector in a future calibration.

8.4.2 Impact on the ${}^8\text{B}$ Analysis

For the higher energy ${}^8\text{B}$ neutrino analysis, the effect of the position reconstruction as a function of energy is much larger since this analysis extends to energies greater than 10 MeV. In the first ${}^8\text{B}$ analysis ([44]) which occurred before the calibrations, the error attributed to the knowledge of the fiducial mass was set at 6% in analogy with the ${}^7\text{Be}$ analyses.

During the calibration campaigns, the AmBe source was deployed in numerous locations throughout the inner-vessel (for exact locations, see the tables in Appendices D, E, and F), most important were several points on a three meter radius sphere along the principle axes. Events with different energies reconstructed at progressively higher radii and these data were used to define the fiducial volume as a function of position and energy. The combined error due to both the position reconstruction accuracy and the energy dependent bias was reduced to 3.8% in light of the calibration data[44].

The effect of the energy calibration on the second ${}^8\text{B}$ analysis is more difficult to quantify. Indeed, the most recent analysis has attributed an error due to the energy calibration of $(+3.6, -3.2)\%$ — due to the variation in the energy threshold. The first iteration of this

analysis quoted an error of 3% for this effect⁶, however, it has been regarded as an under-estimation. At the very least, the calibration provided a method to properly estimate the error associated with the variation in the energy response throughout the fiducial volume; few would argue that using the correct error is more important than using a lower one.

8.4.3 Impact on the Geo- $\bar{\nu}_e$ Analysis

The two-event signature for the detection of $\bar{\nu}_e$ greatly improves the discriminating power over background events. In theory, the tag is so efficient, that one may not even use a fiducial volume, and instead, use the entire active volume for the analysis. However, this tag can be mimicked effectively by fast neutrons. In the process of thermalization, the energy given to recoil protons can be observed as a prompt (emulating the positron annihilation) signal, while the capture of the thermalized neutron mimics the delayed signal. After muon-induced neutrons, the primary source of such fast neutrons is as the result of (α, n) reactions — most commonly due to ^{210}Po on carbon.

The earliest, unpublished, versions of the anti-neutrino analyses did not make use of the entire inner-vessel because of the possibility for (α, n) reactions initiated by α -decaying nuclei in the vessel material or in the inner buffer. Since the accuracy of the position reconstruction was not well known prior to the calibrations, a fiducial volume was utilized.

The results of the calibration combined with the ability to measure the shape profile of the inner-vessel — both products of the calibration system discussed in this dissertation — allowed the analysis to be extended to nearly the full inner-vessel. The only volumetric cut made was to veto any events occurring within 25 cm of the inner-vessel for a given θ . Without the ability to expand the target mass in this way, the statistical error on the measurement

⁶These numbers are for the analysis performed with a 2.8 MeV threshold, the errors for the analysis with a 5.0 MeV threshold are higher by a factor of about 1.5.

would have been much larger.

Analysis of the energy distribution of events in the geo-neutrino study required simulated geo-neutrino and reactor anti-neutrino spectra in order to compare with the experimental results (and to subtract the reactor anti-neutrino component). Since the energy spectrum for reactor anti-neutrinos extends up to approximately 8 MeV, the energy response of the detector must be known from the lower energy threshold (1.8 MeV) up to 8 MeV — and must be known over the entire analysis vessel.

The Geant4 Monte Carlo code used to generate the expected spectra was tuned to the data obtained with the AmBe source and many of the gamma sources to determine the energy response as a function of energy and particle type. After tuning the Monte Carlo to the calibration data, it was used to determine the efficiencies of the selection cuts and to obtain estimates on the residual backgrounds.

Chapter 9

Conclusion

In the first two and a half years of data-taking, the Borexino experiment has:

- Performed a 10% measurement on the ${}^7\text{Be}$ ν flux [20]
 - Combining this measurement with those from the radiochemical experiments led to the best upper limit on the contribution of the CNO neutrinos to the total solar flux: $< 5.4\%$ at 90% confidence level.
 - A similar calculation under the assumption of the luminosity constraint yields the best determination of the pp neutrino flux: $f_{pp} = 1.005^{+0.008}_{-0.020}$.
 - The shape of the low-energy portion of the α -subtracted spectrum used for the ${}^7\text{Be}$ analysis allowed the collaboration to place a limit on a non-zero neutrino magnetic moment: $\mu_\nu < 5.4 \times 10^{11} \mu_B$ at 90% confidence level — this is currently the best limit on the neutrino magnetic moment.
- Measured the ${}^8\text{B}$ solar neutrino flux with the lowest threshold to date[44]
 - Doing so has also made Borexino the first experiment to simultaneously measure

neutrinos which oscillate via the vacuum and matter enhanced mechanisms

- The survival probabilities extracted from the ^8B and ^7Be solar neutrino analysis differ by 1.9σ thus providing further evidence in favor of the MSW-LMA picture of solar neutrino oscillations.
- Established observation of geo-neutrinos at a 99.997% confidence level [51]
 - The anti-neutrino measurement necessarily includes a component due to nuclear reactors — the measurement performed by Borexino excludes the non-oscillation hypothesis of the reactor anti-neutrinos at a 99.6% confidence level.
 - The signal-to-noise ratio in the geo-neutrino analysis is 23:1 — the same ratio in the KamLAND result was 1:4. Thus, despite its smaller size, the unprecedented radiopurity and the lack of nearby reactors make Borexino a highly-competitive geo- $\bar{\nu}_e$ experiment.

The future of the experimental program is, in spite of the leak in the inner vessel, still quite promising. The collaboration is currently in the process of performing another ^7Be ν analysis (with lower errors and higher statistics), pep and CNO neutrino analyses, a possible measurement of the upper end of the pp neutrino spectrum, and further refinements on the geo-neutrino and ^8B neutrino results with larger data sets. The experiment has also joined the SNEWS (Supera Nova Early Warning System) network in 2009 and has increased its duty cycle to maximize the possibility of observing supernova neutrinos.

Without a doubt, the data collected from the calibrations described in this dissertation have been of vital importance to the reduction of the systematic errors; and future precision measurements of neutrino fluxes would be impossible without these data. Contamination analyses performed on the data after the calibrations yield results which are consistent with no introduction of any sources of long term contamination, however, more data is needed

in order to draw firm conclusions. The ability to monitor the status of the vessel with the calibration cameras has been invaluable, particularly in studying the evolution of the leak during refilling and the forthcoming purification campaigns.

Very recent measurements of the vessel volume indicate that the reduction of the DMP concentration has reduced the rate of the leak in the inner vessel to $0^{+0.5}\text{m}^3/\text{month}$. Moreover, a repurification campaign began in March, 2010 with the goal of reducing the ^{85}Kr and ^{7}Be - ^{11}C valley contamination. The combination of these two operations should enable the experiment to operate in perhaps an even more favorable condition than it has to date, thus ensuring its place as one of the world's premier neutrino observatories.

Bibliography

- [1] B. T. Cleveland *et al.*, “Measurement of the solar electron neutrino flux with the homestake chlorine detector,” *Astrophys. J.*, vol. 496, no. 1, pp. 505–526, 1998.
- [2] M. Altmann *et al.*, “Complete results for five years of GNO solar neutrino observations,” *Phys. Lett.*, vol. B616, pp. 174–190, 2005, hep-ex/0504037.
- [3] Y. Ashie *et al.*, “Evidence for an oscillatory signature in atmospheric neutrino oscillation,” *Phys. Rev. Lett.*, vol. 93, p. 101801, 2004, hep-ex/0404034.
- [4] B. Pontecorvo, “Inverse beta processes and nonconservation of lepton charge,” *Sov. Phys. JETP*, vol. 7, pp. 172–173, 1958.
- [5] B. Pontecorvo, “Neutrino experiments and the question of leptonic-charge conservation,” *Sov. Phys. JETP*, vol. 26, pp. 984–988, 1968.
- [6] A. B. McDonald, J. R. Klein, and D. L. Wark, “Solving the solar neutrino problem,” *Scientific American*, vol. 288, pp. 40–50, April 2003.
- [7] S. Abe *et al.*, “Precision measurement of neutrino oscillation parameters with kamland,” *Phys. Rev. Lett.*, vol. 100, p. 221803, Jun 2008.
- [8] J. N. Bahcall, *Neutrino Astrophysics*. Cambridge University Press, 1989.

- [9] M. C. Gonzalez-Garcia and M. Maltoni, “Phenomenology with Massive Neutrinos,” *Phys. Rept.*, vol. 460, pp. 1–129, 2008, hep-ph/0704.1800.
- [10] B. Aharmin *et al.*, “Low Energy Threshold Analysis of the Phase I and Phase II Data Sets of the Sudbury Neutrino Observatory,” *Submitted to Phys. Rev. C*, 2009, nucl-ex/0910.2984.
- [11] T. Schwetz, M. A. Tortola, and J. W. F. Valle, “Three-flavour neutrino oscillation update,” *New J. Phys.*, vol. 10, p. 113011, 2008, hep-ph/0808.2016.
- [12] S. Mikheyev and A. Y. Smirnov, “Resonant amplification of ν oscillations in matter and solar-neutrino spectroscopy,” *Il Nuovo Cimento*, vol. 9C, no. 1, pp. 17 – 26, 1986.
- [13] L. Wolfenstein, “Neutrino oscillations in matter,” *Phys. Rev. D*, vol. 17, pp. 2369–2374, May 1978.
- [14] A. Strumia and F. Vissani, “Neutrino masses and mixings and...,” 2006, hep-ph/0606054.
- [15] M. Pallavicini, “Solar neutrino results from Borexino and main future perspectives,” 2009, astro-ph.SR/0910.3367.
- [16] Q. R. Ahmad *et al.*, “Measurement of the charged current interactions produced by B-8 solar neutrinos at the Sudbury Neutrino Observatory,” *Phys. Rev. Lett.*, vol. 87, p. 071301, 2001, nucl-ex/0106015.
- [17] A. Y. Smirnov, “The MSW effect and solar neutrinos,” 2003, hep-ph/0305106.
- [18] H. O. Back, *Internal radioactive source calibration of the Borexino solar neutrino detector*. PhD thesis, Virginia Tech, 2004.

- [19] J. N. Bahcall, “The luminosity constraint on solar neutrino fluxes,” *Phys. Rev.*, vol. C65, p. 025801, 2002, hep-ph/0108148.
- [20] C. Arpesella *et al.*, “Direct Measurement of the Be-7 Solar Neutrino Flux with 192 Days of Borexino Data,” *Phys. Rev. Lett.*, vol. 101, p. 091302, 2008, astro-ph/0805.3843.
- [21] H. Back *et al.*, “Lens experiment proposal.” 2009.
- [22] H. A. Bethe, “Energy production in stars,” *Physical Review*, vol. 55, pp. 434–456, 1939.
- [23] J. N. Bahcall, A. M. Serenelli, and S. Basu, “New solar opacities, abundances, helioseismology, and neutrino fluxes,” *Astrophys. J.*, vol. 621, pp. L85–L88, 2005, astro-ph/0412440.
- [24] J. Bahcall, “Neutrinos from the sun.” URL http://nobelprize.org/nobel_prizes/physics/articles/fusion/index.html.
- [25] Wikipedia, “Cno cycle,” 2006. [Online; accessed 14-January-2010].
- [26] R. G. Lerner and G. L. Trigg, eds., *Encyclopedia of Physics*. Wiley-VCH, 2005. Available online at <http://www.sns.ias.edu/~jnb/Papers/Popular/Wiley/paper.pdf>.
- [27] M. Leung, *The Borexino Solar Neutrino Experiment: Scintillator Purification and Surface Contamination*. PhD thesis, Princeton University, 2006.
- [28] *Neutrino Mass*, vol. 190 of *Springer Tracts in Modern Physics*, ch. 1. Springer, 2003.
- [29] H. Klapdor-Kleingrothaus, “First evidence for lepton number violation and of the majorana character of neutrinos,” *Nucl. Phys. B (Proc. Supp.)*, vol. 145, pp. 219 – 224, 2005. NOW 2004.
- [30] O. Cremonesi, “Status of neutrinoless double beta decay searches,” *Hyperfine Interactions*, vol. 193, no. 1-3, pp. 261–267, 2009.

- [31] X. Guo *et al.*, “A precision measurement of the neutrino mixing angle $\theta(13)$ using reactor antineutrinos at Daya Bay,” 2007, hep-ex/0701029.
- [32] F. Ardellier *et al.*, “Double Chooz: A search for the neutrino mixing angle $\theta(13)$,” 2006, hep-ex/0606025.
- [33] M. G. Albrow *et al.*, “Physics at a Fermilab proton driver,” 2005, hep-ex/0509019.
- [34] N. Grevesse and A. Sauval, “Standard solar composition,” *Space Science Reviews*, vol. 85, no. 1, pp. 161–174, 1998.
- [35] S. Basu and H. M. Antia, “Constraining solar abundances using helioseismology,” *The Astrophysical Journal Letters*, vol. 606, no. 1, pp. L85–L88, 2004.
- [36] A. M. Serenelli, “New Results on Standard Solar Models,” 2009, astro-ph.SR/0910.3690.
- [37] M. Asplund, N. Grevesse, and A. Sauval, “Cosmic abundances as records of stellar evolution and nucleosynthesis,” vol. 336, p. 25, 2005.
- [38] M. Asplund, N. Grevesse, A. J. Sauval, and P. Scott, “The chemical composition of the sun,” *Annual Review of Astronomy and Astrophysics*, vol. 47, no. 1, pp. 481–522, 2009, astro-ph.SR/0909.0948.
- [39] T. Kovacs *et al.*, “Borex: Solar neutrino experiment via weak neutral and charged currents in boron-11,” *Solar Physics*, vol. 128, no. 1, pp. 38–938, 1990.
- [40] C. Pena-Garay and A. Serenelli, “Solar neutrinos and the solar composition problem,” 2008, astro-ph/0811.2424.
- [41] V. Barger, P. Huber, and D. Marfatia, “Solar mass-varying neutrino oscillations,” *Phys. Rev. Lett.*, vol. 95, p. 211802, 2005, hep-ph/0502196.

- [42] A. Friedland, C. Lunardini, and C. Pena-Garay, “Solar neutrinos as probes of neutrino - matter interactions,” *Phys. Lett.*, vol. B594, p. 347, 2004, hep-ph/0402266.
- [43] G. L. Fogli, E. Lisi, A. Marrone, A. Palazzo, and A. M. Rotunno, “Hints of $\theta_{13} > 0$ from global neutrino data analysis,” *Phys. Rev. Lett.*, vol. 101, p. 141801, 2008, hep-ph/0806.2649.
- [44] G. Bellini *et al.*, “Measurement of the solar 8B neutrino rate with a liquid scintillator target and 3 MeV energy threshold in the Borexino Detector,” 2008, astro-ph/0808.2868.
- [45] A. Ray, “Massive stars as thermonuclear reactors and their explosions following core collapse,” 2009, astro-ph.SR/0907.5407.
- [46] L. Cadonati, F. P. Calaprice, and M. C. Chen, “Supernova neutrino detection in Borexino,” *Astropart. Phys.*, vol. 16, pp. 361–372, 2002, hep-ph/0012082.
- [47] M. Wurm, *Cosmic Background Discrimination for the Rare Neutrino Event Search in BOREXINO and LENA*. PhD thesis, Technische Universität München, 2009.
- [48] T. Araki *et al.*, “Experimental investigation of geologically produced antineutrinos with KamLAND,” *Nature*, vol. 436, pp. 499–503, July 2005.
- [49] G. Fiorentini, F. Mantovani, and B. Ricci, “Neutrinos and energetics of the earth,” *Physics Letters B*, vol. 557, no. 3-4, pp. 139–146, 2003.
- [50] Y. Suvorov, *The Geo-neutrino signal in Borexino*. PhD thesis, Università Degli Studi di Milano & Université Paris Diderot Ecole Doctorale 381, 2009.
- [51] G. Bellini *et al.*, “Observation of Geo-Neutrinos,” *Submitted to Phys. Lett. B*, 2010, hep-ex/1003.0284.

- [52] K. Zuber, “Status of the SNO+ experiment.” Talk given at the MEDEX 09 conference, June 2009.
- [53] J. G. Learned, “Detection of geoneutrinos: Can we make the gnus work for us?,” in *The XXIII Conference on Neutrino Physics and Astrophysics*, vol. 136, Journal of Physics, IOP Publishing, 2008.
- [54] M. Wurm. Private communication, 2009.
- [55] D. D’Angelo, *Towards the detection of low energy solar neutrinos in BOREXino: Data readout, data reconstruction and background identification*. PhD thesis, Technische Universität München, 2006.
- [56] N. Ferrari, G. Fiorentini, and B. Ricci, “The ^{51}Cr neutrino source and Borexino: a desirable marriage,” *Physics Letters B*, vol. 387, pp. 427 – 431, October 1996.
- [57] K. B. McCarty, *The Borexino Nylon Film and the Third Counting Test Facility*. PhD thesis, Princeton University, 2006.
- [58] A. Pocar, *Low background techniques and experimental challenges for Borexino and its nylon vessels*. PhD thesis, Princeton University, 2003.
- [59] J. Benziger *et al.*, “The nylon scintillator containment vessels for the Borexino solar neutrino experiment,” *Nucl. Instrum. Meth. A*, vol. 582, p. 509, 2007.
- [60] “Pseudocumene MSDS.” URL <http://chemicaland21.com/industrialchem/organic/PSEUDOCUMENE.htm>.
- [61] OSHA, “Occupational safety and health guideline for dimethylphthalate.” URL <http://www.osha.gov/SLTC/healthguidelines/dimethylphthalate/recognition.html>.

- [62] “2,5-diphenyloxazole product description.” URL http://www.chemicalbook.com/ChemicalProductProperty_EN_CB3693821.htm.
- [63] “Mg500 diagram.” [Online; accessed 29-September-2009].
- [64] J. B. Birks, *The Theory and Practice of Scintillation Counting*. Pergamon Press, 1964.
- [65] M. C. Johnson, *Scintillator purification and study of light propagation in a large liquid scintillation detector*. PhD thesis, Princeton University, 1998.
- [66] M. Göger-Neff, *Development of a Liquid Scintillator and of Data Analysis Methods for BOREXINO*. PhD thesis, Technische Universität München, 2001.
- [67] G. Alimonti *et al.*, “Light propagation in a large volume liquid scintillator,” *Nucl. Instrum. Meth.*, vol. A440, pp. 360–371, 2000.
- [68] O. Smirnov, P. Lombardi, and G. Ranucci, “Precision measurements of time characteristics of ETL9351 photomultipliers,” *Instruments and Experimental Techniques*, vol. 47, no. 1, pp. 69 – 79, 2004.
- [69] J. Benziger *et al.*, “A scintillator purification system for the Borexino solar neutrino detector,” *Nucl.Instrum.Meth.A*, vol. 587, p. 277, 2008, physics.ins-det/0709.1503v2.
- [70] D. Manuzio, *Towards the detection of subMeV solar neutrinos in Borexino: data reconstruction and analysis tools*. PhD thesis, Università degli studi di Genova, 2004?
- [71] Mach4, “Mach4 code documentation.” URL <https://argus.princeton.edu/borexino/doxygen/>, 2009.
- [72] “Root: an object oriented data analysis framework.” URL <http://root.cern.ch/>.
- [73] R. N. Saldanha, “Mach4 clustering.” Presentation given at Mach4/Echidna analysis meeting, 2008.

- [74] O. Smirnov, P. Lombardi, and G. Ranucci, “Precision measurements of time characteristics of the 8” ETL9351 series photomultiplier,” *Instr. and Exper. Techn.*, vol. 47, no. 1, pp. 69–79, 2004.
- [75] E. Gatti and F. Martini, “A new linear method of discrimination between elementary particles in scintillation counters,” in *Proceedings of the Conference on Nuclear Electronics*, p. 265, IAEA, 1962.
- [76] M. Pallavicini, “Getting the first ${}^7\text{Be}$ ν detection: scintillator purification, detector response and data analysis in Borexino.” Talk at TAUP 2007, 2007.
- [77] Wikipedia, “Decay chain (4n+2, uranium series),” 2008. [Online; accessed 14-September-2009].
- [78] Wikipedia, “Decay chain (4n, thorium series),” 2008. [Online; accessed 14-September-2009].
- [79] G. Alimonti *et al.*, “Measurement of the ${}^{14}\text{C}$ abundance in a low-background liquid scintillator,” *Physics Letters B*, vol. 422, 1998.
- [80] C. Lederer and V. S. Shirley, “*Table of Isotopes*”. ”Wiley-Interscience”, 7 ed., 1980.
- [81] D. Franco, *The Borexino Experiment: Test of the purification systems and data analysis in the Counting Test Facility*. PhD thesis, Università degli studi di Milano & Ruperto-Carola University of Heidelberg, 2005.
- [82] M. Tokiwa *et al.*, “Atmospheric muon measurements at sea level iii: Muon flux,” in *28th International Cosmic Ray Conference Proceedings*, pp. 1175 – 1178, Universal Academy Press, Inc., 2003.
- [83] C. Arpesella *et al.*, “First real time detection of $\text{Be}7$ solar neutrinos by Borexino,” *Phys. Lett.*, vol. B658, pp. 101–108, 2008, astro-ph/0708.2251.

- [84] T. Hagner *et al.*, “Muon-induced production of radioactive isotopes in scintillation detectors,” *Astroparticle Physics*, vol. 14, no. 1, pp. 33–47, 2000.
- [85] C. Galbiati *et al.*, “Cosmogenic C-11 production and sensitivity of organic scintillator detectors to pep and CNO neutrinos,” *Phys. Rev.*, vol. C71, p. 055805, 2005, hep-ph/0411002.
- [86] R. N. Saldanha, “ ^{11}C subtraction in Borexino.” Presentation given at Borexino collaboration meeting, 2008.
- [87] S. Abe *et al.*, “Study of the Production of Radioactive Isotopes through Cosmic Muon Spallation in KamLAND,” 2009, hep-ex/0907.0066.
- [88] G. Alimonti *et al.*, “A large-scale low background liquid scintillator detector: the counting test facility at gran sasso,” *Nucl. Instrum. Meth. A*, vol. 406, p. 411, 1998.
- [89] H. O. Back *et al.*, “Phenylxylylethane (PXE): a high-density, high-flashpoint organic liquid scintillator for applications in low-energy particle and astrophysics experiments,” *submitted to Nucl. Instrum. Meth. A*, 2007, physics.ins-det/0408032.
- [90] G. Alimonti *et al.*, “Ultra-low background measurements in a large volume underground detector,” *Astroparticle Physics*, vol. 8, no. 3, pp. 141 – 157, 1998.
- [91] G. Bellini *et al.*, “Search for solar axions emitted in the M1-transition of $^7\text{Li}^*$ with Borexino ctf,” *The European Physical Journal C*, vol. 54, no. 1, pp. 1434–6044, 2008.
- [92] H. O. Back *et al.*, “Pulse-shape discrimination with the counting test facility,” *Nucl. Instrum, Meth. A*, vol. 584, no. 1, pp. 98–113, 2007.
- [93] H. O. Back *et al.*, “CNO and pep neutrino spectroscopy in Borexino: Measurement of the deep-underground production of cosmogenic ^{11}C in an organic liquid scintillator,” *Physical Review C*, vol. 74, no. 4, 2006.

- [94] M. Balata *et al.*, “Search for electron antineutrino interactions with the Borexino counting test facility at gran sasso,” *The European Physical Journal C*, vol. 47, pp. 21 – 30, 2006.
- [95] H. O. Back *et al.*, “New experimental limits on violations of the pauli exclusion principle obtained with the Borexino counting test facility,” *The European Physical Journal C*, vol. 37, p. 421, 2004.
- [96] H. O. Back *et al.*, “New experimental limits on heavy neutrino mixing in B-8 decay obtained with the prototype of the Borexino detector,” *JETP Letters*, vol. 78, no. 5, pp. 261 – 266, 2003.
- [97] H. O. Back *et al.*, “Study of the neutrino electromagnetic properties with the prototype of the Borexino detector,” *Physics Letters B*, vol. 563, p. 37, 2003.
- [98] H. O. Back *et al.*, “New limits on nucleon decays into invisible channels with the borexino counting test facility,” *Physics Letters B*, vol. 563, p. 23, 2003.
- [99] H. O. Back *et al.*, “Search for electron decay mode $e \rightarrow \gamma + \nu$ with prototype of Borexino detector,” *Physics Letters B*, vol. 525, p. 29, 2002.
- [100] G. Alimonti *et al.*, “Light propagation in a large volume liquid scintillator,” *Nucl. Instrum. Meth. A*, vol. 440, p. 360, 1998.
- [101] J. Benziger *et al.*, “Borexino internal memo: Update on bx backgrounds seen with mach4.” Mach4 Analysis Memo, November 2007.
- [102] J. C. Maneira, *Calibration and Monitoring for the Borexino Solar Neutrino Experiment*. PhD thesis, Faculdade de Ciências da Universidade de Lisboa, 2001.
- [103] F. Calaprice *et al.*, “Borexino internal memo: First data for the full boreixno detector, detailed analysis of april 2007 runs.” Mach4 Analysis Memo, May 2007.

- [104] H. Simgen, C. Buck, G. Heusser, M. Laubenstein, and W. Rau, “A new system for the ^{222}Rn and ^{226}Ra assay of water and results in the b project,” *NIM A*, vol. 497, no. 2-3, pp. 407 – 413, 2003.
- [105] W. Rau and G. Heusser, “ ^{222}Rn emanation measurements at extremely low activities,” *Applied Radiation and Isotopes*, vol. 53, pp. 371 – 375, 2000.
- [106] R. W. Kay, “Elemental abundances relevant to identification of magma sources,” *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 310, no. 1514, pp. 535–547, 1984.
- [107] B. Bajcsy, S. S. Beauchemin, R. Bajcsy, and G. Givaty, “A unified procedure for calibrating intrinsic parameters of spherical lenses,” in *In Vision Interface (VI 99)*, pp. 272–279, 1999.
- [108] A. Goretti. Private communication, 2007.
- [109] L. Cadonati, *The Borexino Solar Neutrino Experiment and its Scintillator Containment Vessel*. PhD thesis, Princeton University, 2001.
- [110] B. Berger *et al.*, “The kamland full-volume calibration system,” *JINST*, vol. 4, p. 04017, 2009, physics.Ins-det/0903.0441.
- [111] U. S. Government, “Military standard 1246c.”
- [112] A. Ianni. Private communication, 2009.
- [113] NIST, “X-ray mass attenuation coefficients,” 2004. [Online; accessed 5-February-2010].
- [114] M. Johnson *et al.*, “A Rn-222 source for low-background liquid scintillation detectors,” *NIM A*, vol. 414, p. 459, 1998.

- [115] Pylon Electronics, *RN-2015 / Th-1025 / RNC Flow-through Radon Source Datasheet*.
- [116] G. Alimonti *et al.*, “Measurement of the ^{14}C abundance in a low-background liquid scintillator,” *Phys. Lett. B*, vol. 422, p. 349, 1998.
- [117] “Midas daq.” URL <https://midas.psi.ch>.
- [118] G. F. Knoll, *Radiation Detection and Measurement*. John Wiley & Sons Inc., 3rd ed., 2000.
- [119] A. Chavarria. Private communication, 2010.
- [120] B. N. Laboratory, “Thermal neutron capture γ 's (capgam).” URL <http://www.nndc.bnl.gov/capgam/>, December 2008.
- [121] S. Perasso. Private communication, 2009.
- [122] A. Razeto, “Laser source calibration.” Presentation given at Borexino Collaboration Meeting, 2009.
- [123] G. Ranucci, “ ^7Be precision measurements.” Presentation given at Borexino Collaboration Meeting, 2009.
- [124] G. Fernandes, *Analisi della risposta del rivelatore Borexino con sorgente di calibrazione*. PhD thesis, Università di Genova, 2010.

Appendix A

Process Control for the Source Deployment System

The source deployment system relies heavily on maintaining an inert atmosphere of LAKN in several different volumes and precisely maintaining their pressures and flow rates. In addition to the gas flow requirements, we have several important safety instruments to alert shifters to the presence of oxygen in the glovebox, liquid in the cross, inadequate ventilation of instrumentation, etc. Due to the complexity of the operations involved, it was decided early on to have the operation of the system be as automatic as possible to allow the shifters to concentrate on the manipulation of the source. In the sections that follow, the hardware, electronics, and DAQ software of the system are presented.

A.1 Hardware

A piping and instrumentation diagram (PID, /P&ID) for the system is presented in Figure A.1 and will be a useful reference for the discussion that follows where repeated reference is

made to various valves, pressure gauges, etc.

Valves

The gas control system relies on a system of solenoid valves (V1-V8) which are operated by the computer and act to start or stop a supply of gas to a piping run, and a set of manual needle valves (NV1-NV5) which are used to control the flow rate through the system¹.

Valves V1-V8 operate on their own 24 VDC power supply housed in the electronics box and are cabled to individual solenoids. Since they control a gas stream directly in contact with the inner vessel (via the cross), they were chosen with cleanliness in mind and are all stainless steel with PTFE seals and discs. Since some of them also handle pseudocumene vapors, they are all NEMA rated as explosion proof.

The manual valves NV1-NV5 used to control flow rates through the system are all manually operated, severe service, needle valves. In addition to these needle valves, one metering valve, MV1 was installed on the line that feeds the glovebox with LAKN. If the box is not already pressurized, opening V4 to start pressurizing it, results in a very large flow due to the immense volume relative to the line feeding it; the flow, limited only by line impedance would very often overrun the flowmeter on this line (FM4), and cause it to output voltages high enough to cause cross talk with other channels. Limiting the flow on this line to a value that does not cause this to occur is the job of the metering valve MV1.

Three pressure regulators, R1-R3, step down the LAKN gas pressure provided to various areas of the system. All three regulators are from the Swagelok KLF series of high sensitivity diaphragm pressure regulators and have stainless steel bodies and PTFE seats. The regulator R1 sets the supply pressure for the entire gas flow system and R3 further steps this pressure

¹Regulating solenoid valves were extremely expensive and thus using one solenoid and one manual valve proved to be the more economical choice.

down to the pressure required in the cross to keep the liquid level below the gate valve — since these two pressures do not differ greatly, it was decided to use identical model regulators for both R1 and R3. The inlet pressure range for both R1 and R3 is 0—248 bar, with an outlet control range of 0–.68 bar. R2 on the other hand needs to produce a much lower pressure of just a few mbar for the glovebox, so, we chose a model with a more restricted outlet range of 0–0.13 bar and an inlet range of 0–1.0 bar.

The system also employs several check valves for ensuring that the pressures stay within range and also to provide a disposal path for any spilled scintillator. Check valve CV1, CV2a, and CV2b are all identical valves designed and fabricated at Virginia Tech and are simply cylinders with a delrin sphere which seals against a large o-ring, application of a pressure in excess of 4 mbar causes the valves to open. CV1 is located inside of the glovebox and protects the glovebox from under pressurization relative to the clean room. This can only occur if people withdraw their arms from the glovebox in a very rapid fashion. CV2a and CV2b are identical valves mounted outside of the glovebox and vent the excess pressure from the glovebox in case R2 fails, or shifters insert their arms too quickly — two valves are used for redundancy purposes in case one were to fail to open. A third check valve, CV3, exists in the pipe connecting the drain in the bottom panel of the glovebox to the waste tank — the weight of the scintillator in the tube section will open the valve and cause it to drain into the tank.

The exhaust of the entire system is routed to a large charcoal filter mounted outside of the clean-room which vents to the Hall C atmosphere. It is possible, although not terribly feasible, that the pressure in Hall C would increase to the point that it exceeds the pressure in the exhaust line, in which case exhaust gas could be blown back into the glovebox. To prevent this from occurring, the exhaust tube from the glovebox is immersed in a column of mineral oil contained in a trap — the trap contains another pipe (not immersed in mineral

oil) which routes to the exhaust filter. This setup would not permit exhaust gas to back-flow into the glovebox unless a very serious pressure difference were to occur (such a difference would likely cause structural damage to the clean-room before the gas was allowed to flow backwards).

Pumps

Obviously, the process of changing a source requires that the cross be opened in order to remove the old one and attach the new one, and it is inevitable that oxygen, and any radon in the clean-room air will also enter the cross during this time. In order to maintain the radiopurity, we need to not have any radon or krypton in the gas in the cross when the gate valve is open; to keep the optical properties in the desired regime, we should also remove any oxygen present. The solution is to purge the gas in the cross before opening it. However, purging a volume as large as that of the cross would take a very long time due to mixing. A much more efficient method is to first evacuate the cross, then backfill it with LAKN, evacuate it again and backfill it, etc. To that end, we perform at least seven iterations of evacuating to -800 mbar, and then backfilling to 5 mbar, this achieves a reduction by a factor of 10^{-5} in the concentration of the initial gas in the cross.

Additionally, we would like a diagnostic tool to know, at any given point in time, the level of oxygen in the atmosphere of the glovebox. The real concern is not oxygen itself, but radon, krypton, argon, and other such gases with radioactive isotopes — all of which come from the air in the surrounding environment — oxygen is the easiest of these to detect and industrial oxygen sensors are commonly available. The oxygen sensor chosen for the insertion system does not work with ambient pressure. Thus, a diaphragm pump is used to provide the specific pressure and flow rate required for the oxygen monitor.

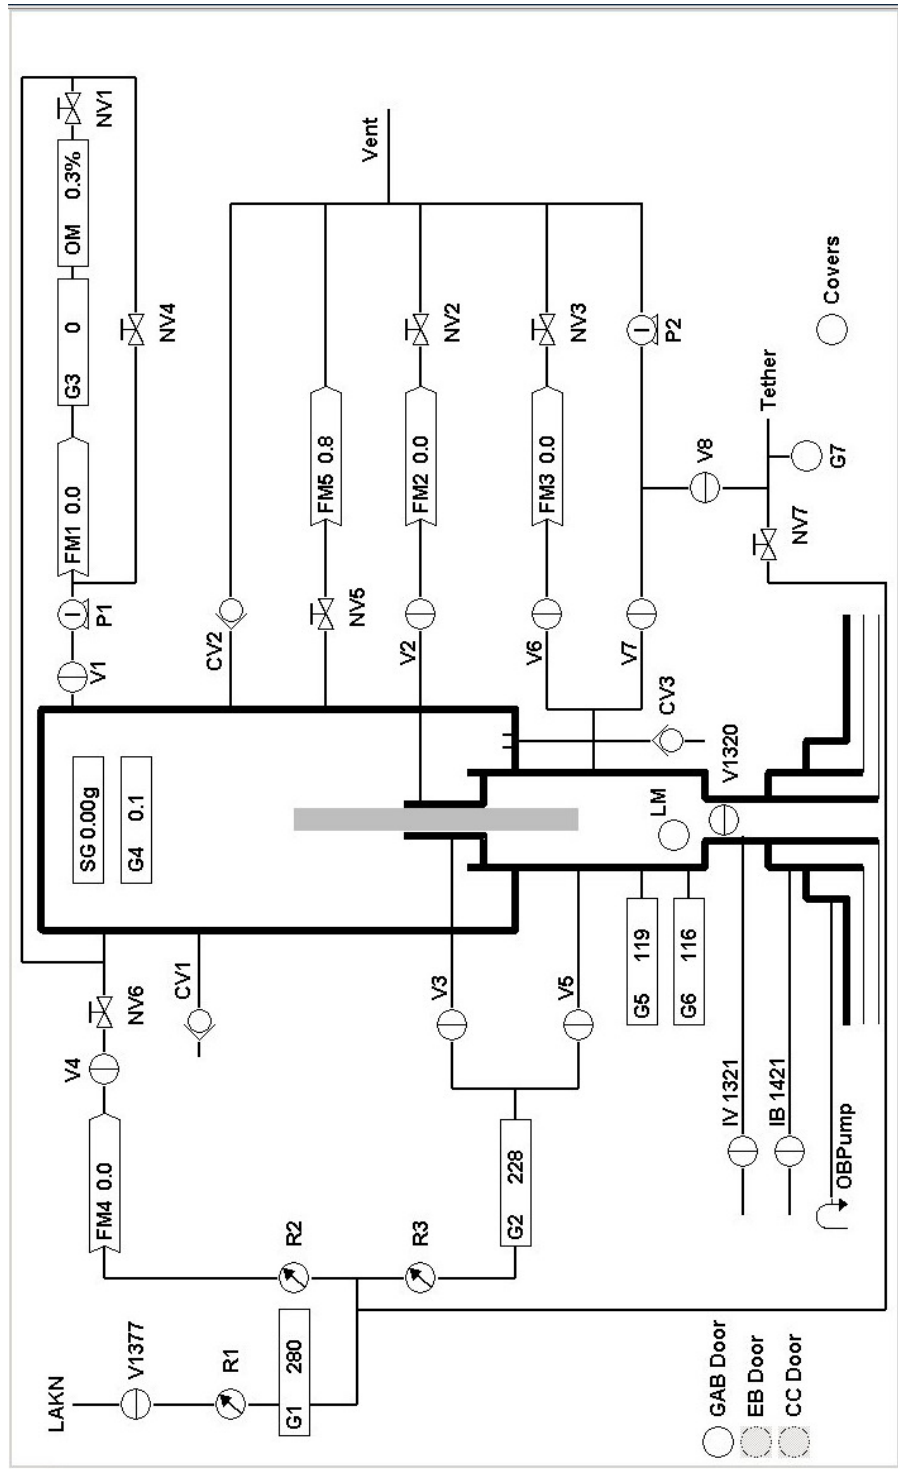


Figure A.1: Piping and instrumentation diagram of the gas control component of the source deployment system. In actuality, the photo is a screen shot of the main control panel for the control software.

In the interest of interchangeability and redundancy, it was decided to use vacuum pumps for both purposes because the oxygen monitor is not *absolutely* necessary, but being able to purge the cross in a timely manner is. The volume of gas and the pressures required are small enough that large pumps were not required and so we purchased two very small diaphragm pumps from Air Dimensions Inc. The pumps operate on 12VDC, draw two amps of current, and all wetted parts are stainless steel with PTFE diaphragms. They are dual headed pumps, but the heads are connected in series since we emphasize a lower pressure rather than a higher pumping speed.

Pressure Measurement

Six diaphragm pressure transducers, G1-G6, are installed throughout the system and give continuous feedback about the pressure inside of the conduits in which they are installed. Because many of them measure the pressure of a gas stream containing pseudocumene vapors, they are Wika E-10 stainless steel, explosion proof, transducers. The transducers operate on a 24 VDC supply and produce an output signal in the range 1–5 VDC which is linearly proportional to the pressure. Each gauge has been individually calibrated with a Bourdon tube and are guaranteed to be accurate to 0.25% of their full scale range.

For interchangeability purposes, we tried to keep the gauges as identical as possible, which meant choosing a range applicable to as many situations as possible. To that end, four of the gauges, G2,3,5, and 6 are all identical and are compound gauges with ranges of -30 inHg to $+30$ psi — G2 measures the pressure of the gas after the step down regulator R3, G3 measures the pressure of the gas entering the oxygen monitor, and G5/6 both measure the pressure in the cross — the presence of two is for redundancy reasons in the event that one should fail. Since the inlet pressure to the system can be more than two bars, G1 was specified with a range of 0-100 psi. Finally, the pressure in the glovebox is maintained at

2–3 mbar, or, 0.03 psi, this is a factor of two lower than 0.25% of the range of the compound gauges, so, using these gauges for this application would not be accurate, thus, G4 is specified with a range of 0 – 5 psi.

Flow Measurement

It is also important for us to know the flow rate of gas through various seals and areas to ensure that purging is indeed taking place, and also to perform a flow balance to look for the presence of any leaks. Five flow meters, FM1-FM5, installed throughout the system provide this information. The flow meters are all Sierra Instruments 824S thermal mass-flow meters calibrated for nitrogen service. As with the pressure transducers, these meters operate on 24 VDC and output a signal in the range 1–5 VDC which is linearly proportional to the flow rate in the range 0–4 SLPM. Four of the gauges, FM2-FM5 are specified with a maximum input pressure of 6 psi, and FM1, the one which measures the flow rate coming from P1, is specified with a maximum input pressure of 75 psi since the diaphragm pumps produce pressures over a wide range. Many of the flow meters do not contain the recommended² length of straight tubing at their inlet and outlet, nor do they contain flow straighteners, however, this does not affect the measurement of the flow at a magnitude to cause concern. As mentioned in the section on valves, a metering valve on the glovebox inlet line ensures that the flow of gas going through FM4 does not exceed 4 SLPM and avoids an output voltage in excess of the 10 VDC which has been seen to cause cross talk to other instruments.

Oxygen Monitoring

Measurement of the oxygen content of the glovebox (which is used as a proxy for measuring radon, krypton, etc.) is performed by a dedicated oxygen monitor produced by Advanced

²The usual rule of thumb is at least five times the diameter of the tubing

Micro Instruments. The monitor, which is an AMI model 65, is based on a Zirconium Oxide detection method which does not deplete the sensor or require replacement after several years. The sensor operates on 24 VDC and produces an output signal in the range of 0–2.5 VDC which is proportional to the oxygen content of the sampling stream in the range of 0–25%.

The feed for the oxygen monitor is provided by the output of P1, however, this pump produces a pressure in excess of that recommended for the sensor, and a flow that is too high. The bypass valve NV4 is used in tandem with flow controlling valve NV1 to set the pressure and flow rate through the oxygen monitor. The situation begins with a fully open NV4, and fully closed NV1, thus all of the gas is bypassed through NV4 — the operators then close NV4 while opening NV1 to establish the desired pressure of 1 bar, and the desired flow of 4 LPM.

Liquid Level Monitoring

During a calibration, the only proxy for knowing how high the liquid level is below the fill pipe is to monitor the liquid level in the outer vessel head tank (the valves to the inner vessel and inner buffer head tanks are closed during a calibration). Interpreting this value requires some level of experience and a more direct method was desired for the purposes of safety. Since *any* liquid in the cross is the sign of a severe problem in the configuration of the system, we simply require a method to see if the liquid level has reached the cross, not precisely how high it is. The liquid level monitor is constructed from a toroidal stainless steel float which contains a small magnet inside of it, the float fits around a piece of $\frac{1}{4}$ " tubing which contains a reed switch just slightly above the normal resting position of the float. If scintillator rises into the cross, the float will be lifted and the magnet closes the reed switch setting an alarm. Since the float and tube reside in the cross, they too underwent the same critical cleaning

steps as the rest of the cross, rods, etc. A photo of the float assembly installed in the cross can be seen in Figure 6.12 on page 224.

Rod Weight Measurement

As an additional precaution against continuing to use insertion rods that developed leak for any reason, a strain gauge was installed inside of the glovebox to monitor the weight of each rod. The basic idea is to measure the weight of each rod before it goes into the detector and immediately after removal, any increase in the weight must be due to additional mass acquired in the form of a leak³. The strain gauge chosen for the application was the LC601 series from Omega Engineering and has a maximum capacity of 5 pounds. The device produces an output voltage of $2\frac{\text{mV}}{\text{V}}$, thus, for a 12V excitation voltage, the output will be linear in the range of 0–24 mV over the range (0–5) pounds. The load cell is installed underneath of the rod holder in the glovebox such that placing a rod in the center hole in the front row will result in the weight of the rod being recorded.

Ventilation

All of the gauges, valves, pumps, etc. reside in a stainless steel box (the gas box, or GAB) which gets necessarily quite warm due to the heat generated by the components if left unattended. External to the clean-room is another enclosure (the electronics box, or EB) which houses the data acquisition system for all of the gauges and contains all of the terminations of the system — the components of this enclosure also generate heat. Finally, the computers which control the system reside in an enclosure (the computer cabinet, or CC) outside of the clean-room — the enclosure has a glass face which fits into a cutout in the clean-room

³Obviously, the rods will contain a thin film of scintillator after being withdrawn from the detector, and we would only be concerned about a leak on a scale larger than this.

wall. The computer cabinet also contains the control box for the camera system.

To provide ventilation to all of these systems, we have installed a take off from the clean-room air supply duct which is fed progressively through the gas box, the electronics box, and finally to the computer cabinet. With a series configuration like this, having the door to the gas box open will cease any ventilation to the electronics box or computer cabinet (likewise, opening the electronics box door ceases ventilation to the computer cabinet). To ensure that all doors remain closed, micro push button switches were installed on all three doors. The switches are fed with a constant +5 VDC signal, the presence of the logic 1 signal at the corresponding digital inputs of the data acquisition system ensures that the doors are closed. Should one of them be opened, an alarm occurs to alert the operators.

Enclosure

All of the aforementioned items and the vast majority of the plumbing is contained inside of a stainless steel box, known as the “gas box” (or GAB) mounted adjacent to the glovebox in the clean room. Two large conduits enter this box, one bringing fresh air from the air duct and the other exhausting it back into the electronics box (the second conduit also contains the vast majority of the cabling). An annotated photo of the inside of the gas box can be seen in Figure A.2.




Figure A.2: Photo of the “gas box” (GAB) which houses all of the sensing instruments that are part of the gas control system for the source insertion system. The positions of the eight solenoid valve, six pressure transducers, five flow meters, three regulators, two vacuum pumps, oxygen monitor and door switch are indicated in the figure. All of the wiring bundles enter and exit at the left of the photo — the terminal blocks along the left wall allow for easy troubleshooting of power connections.

A.1.1 Electronics

In order to allow the shifters to concentrate as much as possible on the manipulation of the source, the control of the system was designed to be as automatic as possible. The majority of the instrumentation described in the previous section produce analog or digital signals that must be acquired and decoded by a computer to produce the relevant information. Due to the geometry of the system within the clean room, cabling all of the devices to the computer cabinet would require passing cables on top of the clean room adjacent to high voltage cables for other clean room instrumentation. The potential for noise pickup for the analog signals was too high for comfort, thus, a data acquisition system was purchased that sends its results to the computer in the computer cabinet on the opposite side of the clean room.

A.2 Data Acquisition Hardware

The requirements for a system to read and control the system are rather simple:

- The system should be able to operate at least ten relays (eight solenoid valves, and two vacuum pumps), each capable of switching at least two amperes of current (the nominal current draw of the vacuum pumps P1 and P2). Additionally, the system should be able to read the state of the relays in order to recover from a program crash, or other such problem without having to reset each relay. The relays should also be non-latching for safety purposes — no harm can come from having all of the valves shut simultaneously. This by design isolates the entire system, whereas having a valve stick open could cause over pressurization or pressure loss.
- It should be able to acquire at least thirteen analog voltages (six pressure transducers,

five flow meters, an oxygen monitor, and a strain gauge) in the range 0–10 VDC. To minimize the possibility that noise contributes to the acquired signal, the analog signals should be acquired in differential mode — a mode in which a normal and inverted copy of the signal are transmitted along two different lines and are subtracted at the receiving end, providing common-mode rejection.

- The system should also be able to write to and read from several digital lines. The switches on the doors of the enclosures are constantly fed with a +5V signal and are acquired by the system; the lack of the +5V indicates that the door is open and the alarm is raised. A method of measuring the length of the tether inserted or removed from the detector based on a rotary encoder was developed (but not yet used), and acquiring the data from the encoder required being able to issue commands to the decoder chip (clear, read byte one, read byte two, etc.) as well as reading the output of the chip.

After searching through several commercially available options, the system adopted was provided by National Instruments and consists of the following components:

- A four-slot mini chassis (3U) which can accommodate a variety of different modules for performing various functions. The chassis was small enough to be mounted in the electronics box and affords an extra slot (we only use three) for expansion of the system in the future if required. The chassis is model number PXI-1031 from National Instruments.
- One 16 channel relay board for operating the valves and pumps. The relays in this board are all capable of switching a two ampere signal operating at 150 VDC or 125 VAC, and are all SPDT, form C, non-latching relays. Each channel can be switched up to 115 times per second. In addition to operating the solenoid valves and pumps,

two extra relays were envisioned to be used for operating a buzzer and warning strobes mounted in the clean room for alerting the shifters in the event of an alarm condition, however, these have yet to be installed. The module is model number PXI-2566.

- An I/O board for reading analog and digital lines as well as writing to the digital lines. The board chosen has 48 independent digital lines each capable of being placed in read, write, or read+write modes, as well as 32 analog inputs (16 analog inputs when in differential mode). The board also has several additional features which are not used, such as counters, digital triggering of I/O operations, etc. The module is model number PXI-6224
- For relaying the data to the computer, we still faced the problem of having to transmit the signals in a conduit on top of the clean room near high voltage alternating current power lines. To avoid picking up any noise during the transmission, the interface we employ is fiber optic. The DAQ chassis contains a fiber optic communications module (model number PXI-8336), and the gas control computer contains the PCI fiber optic interface card (PCI-8336). The signal travels over a 30 meter long fiber optic cable (model MXI-4) in a conduit which also carries the ventilation from the electronics box to the computer cabinet.

Cabling

Both the relay module and the I/O module provide connection boards with terminal blocks to facilitate tidier connections to the modules. The relay module has a terminal board which plugs into the front of the relay module, while the I/O module provides two break-out boards that are connected to the module via two 68 pin cables. Inside of the electronics box (EB), the cables from all of the instruments arrive and are terminated at a set of terminal blocks

which are then routed to either the I/O break-out boards, or the relay board. This type of setup does make the initial time investment in wiring longer, however, it pays off in that it is easier to debug problems and make changes to the wiring.

A photo of the electronics box can be found in Figure A.3. The DAQ chassis is mounted to an internal frame at the top of the box, leaving plenty of space at its front for the relay module's connection board and space for connections to the other two modules. Below the DAQ chassis are four DIN rails, each containing 80 DIN rail mounted terminal blocks — the rails are denoted by the letters A, B, C, and D, from top to bottom, and the numbering of the terminal blocks proceeds from the left to the right. The first rail, rail A, contains, from left to right:

- Two circuit breakers — one for power to the DAQ chassis, the other for power to the power supplies — both are 220 V breakers.
- Five power supplies. The $+V$ cable for each power supply was color coded in *all* subsequent connections to make identification easier when wiring — colors are indicated for each supply below. Since the wiring distance was not excessive, the supplies were left in “local” sense configuration (i.e. the sense terminal was simply connected to $+V$ to regulate the voltage at the supply connection — wiring the sense terminal to, for instance, a valve power connection would put the supply in “remote” sense configuration and the supply would self regulate such that $+V$ occurred at the valve power connection instead of the supply terminal in order to accommodate for voltage drop along the supply lines). The power supplies are (from left to right):
 - Power supply for the flow meters, pressure transducers, and strain gauge — 24V, 2A; color code: orange.
 - Power supply for the eight solenoid valves — 24V, 5A; color code: red.

- Power supply for the two vacuum pumps — 12V, 5A; color code: purple.
 - Power supply for the oxygen monitor — 12V, 2.5A; color code: yellow.
 - Power supply for the tether meter circuit — 5V, 3A; color code: grey.
- Fused circuit breakers for each of the power supplies.

The second row of terminal blocks, rail B, is actually not used for the gas control system, because we left open the possibility for switching the camera system control over from the separate control module in the computer cabinet, to being run by an additional National Instruments I/O module installed in the spare chassis slot. Thus, we capitalized on the opportunity to pre-wire much of the system for this purpose before wiring any of the gas control system. Two additional I/O module break-out boards were purchased and installed and were wired to the terminal blocks on rail B. There are still no concrete plans to perform this change in the future, but should the need arise, the majority of the work is already completed.

The third row of terminal blocks, on rail C, contains a few additional pre-wired camera connections in the first few blocks, followed by the terminals where one would connect the wiring for audible or strobe alarms. The middle of rail C is reserved as a ground bus and mains power entry point. The rest of rail C contains the digital I/O signals for the tether meter, door switches, etc., as well as the power and signal to gauge G4 since this gauge was added after the initial wiring was completed.

Finally, the last row of terminal blocks, rail D, contains the power and signal connections for all of the rest of the gauges. Below this last row of terminal blocks are four break out boards for the I/O modules — two of which are for the yet-to-be purchased module for controlling the camera system, the other two are for the module used for controlling the gas control system. The two boards on the left are for the former (camera control) and are labeled as

1A and 1B respectively, the two boards on the right are for the latter (gas control) and are labeled as 2A and 2B respectively.

In addition to rails A-D, there is a fifth rail (E) installed on the lower right side of the electronics box. The circuit for the tether meter was placed inside of a DIN rail mountable enclosure and rail E is where this circuit is mounted. In addition to the enclosure, a few terminal blocks are mounted next to it for easier connection to the rest of the system and debugging.

The actual mapping table showing which terminal block is connected to which device (via which break-out terminal) can be found in *Table B.1* in Appendix B.

A.3 Software

The acquisition of the data from the National Instruments system and the corresponding user interface are handled by a set of two programs — `gascontrol.exe` which controls the actual acquisition of the data, and `gascontrol.pl` which handles the user interface and sends requests for information to `gascontrol.exe` and decodes its response. We refer to `gascontrol.pl` as the user interface, or high-level code, and `gascontrol.exe` as the low-level code. LabVIEW was not utilized (despite all of the components being supplied by National Instruments), as, we deemed it wise to keep as much of the calibration software in the same language as was reasonably possible (recall that the camera software is written mostly in Perl/Tk but was done so years earlier). Instead, we simply wrote ad-hoc drivers to interface to the C++ libraries for the modules.




Figure A.3: The electronics box which houses all of the interconnections for the data acquisition of the gas control system instrumentation. See the body text for more details about the components.

Data Acquisition Operations

Communication between the two codes is via a named pipe (`\\.\pipe\GasControl` in the Windows vernacular) setup by the low-level code and connected to by the user interface code. The use of a named pipe is akin to a socket in network terminology and allows the two programs to communicate in a rapid manner that avoids time consuming disk read/write operations. The DAQ system can perform continuous data acquisition at a very high rate, however, this would swamp the computer and is really not necessary, so, the acquisition is configured to occur only when polled at a rate determined by the high-level code (currently set to once per second).

When the gas control software is started, by launching `gascontrol.pl`, it creates a few lock files and then launches the low-level code and passes an array of numbers as command line arguments which are used to tell `gascontrol.exe` which voltage ranges to assign to the various analog inputs. After creating the named pipe and connecting to the interface code, the low-level code configures the two DAQ modules — the relay module is addressed as `/PXI1Slot2/k0:15`⁴(i.e. the second module in the first, and only, PXI chassis; and we control all sixteen relays, labeled `k0`, `k1`, `k2`, ...), the digital I/O portion of the I/O module is addressed as `/PXI1Slot4/portX` (where `X` is 0, 1, or 2 for the three different sets of digital lines), and the analog input portion is addressed as `/PXI1Slot4/aiX` (where `X` is a number [0 – 7, 16 – 23] in differential mode).

After the setup of the devices has occurred, the high level code is then free to issue commands to the low-level code in order to acquire or perform various operations. The possible operations are:

⁴If the PCI card is switched to another computer (during an upgrade for instance), the new computer may identify the device as having a different name (`PXI2Slot1`) for instance, in this case it is necessary to change the device name in the software and recompile it.

- Kill and Exit — This command is executed by sending the string `0x` to the low-level code — the `x` marks the end of the command string. This command results in a graceful exit from the acquisition operation and closes the named pipe.
- Read relay positions — Polling the position of all of the relays occurs after receiving the command string `1x`. The status of all of the relays is then returned to the high level code as a bit field 16 characters in length.
- Operate a relay — When the user changes the state of a relay operated device in the high level interface, it sends a command of the form `2NzSx`. The digit `2` tells the low-level code that it will be changing the state of a relay, the digit `N` tells it which relay to operate — `N` can take on values in the range $[0, 15]$ and is followed by the character `z` which tells the low-level code that the one or two character string between `2` and `z` is the number of the relay to operate (to allow for relays 10–15), the value `S` denotes the state to switch the relay to either `0` (close) or `1` (open). Finally, `x` marks the end of the command string. After issuing the command to operate the relay, the low-level code waits for the relay to settle (a time interval defined by the National Instruments C Library), and then sends `1` back through the pipe if the operation was successful.
- Read analog inputs — The read analog function is executed by passing the command string `3(CHANNELS)x` where `(CHANNELS)` is a list of the analog inputs which we would like to be read — for channels > 10 , capital letters are used in analogy with hexadecimal notation. The analog voltages are sampled by the I/O module and then loaded into the output buffer, each value separated by an `x` and then sent to the high level code where they are split by the `x`'s. All of the analog inputs, with the exception of the one for the strain gauge are set for the range $(-10, 10)$ VDC — the input for the strain gauge is set to $(-1.0, 1.0)$ VDC. Moreover, the output of the strain gauge is volatile enough that a single voltage sample does not accurately allow one to extract the weight of the

rod, thus, its output is sampled 100 times and the average value is reported.

- Read digital inputs — This function is executed when the high level code sends command `4x`. The I/O module contains three digital ports, however, this section of code is hardwired to read only port 2, which contains 32 digital lines configured for input only. The low-level code reads the status of these lines and buffers them into the stream in binary and ends it with *x*.
- Write to digital output lines — The methods in this section are actually completely commented out at the time of this writing since they are not used, but, if uncommented, this section could be employed by sending the command string `5P(DATA)x` where the integer *P* (0,1,2) tells the program which port to write the values to, and the value (DATA) is the bit field that one would like to write, and as usual, *x* marks the end of the command.
- Read tether length — This function is only operational if one employs the automated tether meter, however, it has not yet been used during a calibration. The operation is performed when the command `6x` is received. The low-level code then proceeds to set the control lines to read the four bytes that comprise the contents of the internal scaler of the decoder chip, then the length is computed and sent back through the pipe followed by an *x*.
- Zero tether meter — Again, this function is only relevant when the tether meter is connected to the system, and is executed when the command `7x` is received. The low-level code then issues the commands to the decoder chip (on the digital lines) to reset the internal scaler.

User Interface

A screen shot of the user interface of the software (obtained by clicking on “Gas Control” from the main `vtcalib` launcher) is shown in Figure A.4. To help in debugging and for instructional purposes, the user interface code was written to also simulate values for the devices and operate completely without the DAQ system connected. This is accomplished by changing the value of the variable `$daqflag` from 1 to 0 in the file `gascontrol.pl`. The left hand side of the page allows the user to control the operation of the software; brief descriptions of each button follow:

- Start/Stop Auto Updates — These buttons control whether or not the software is polling the status of the instrumentation and updating it on the screen or not, the software defaults to off when first started. The colored box beneath these buttons flashes between blue and red when the system is polling the instruments.
- Unset Manual Checks — The procedure for the system requires that shifters mark certain manual instruments as set, open, closed, etc. For instance, after applying the light blocking covers, they should mark them as applied, or after setting the regulators, they should be marked as such. This button clears the state of all of the manual devices.
- All Auto Closed — This button simultaneously opens all of the relays in the system, in other words, it closes all of the solenoid valves V1-V8 and turns off both vacuum pumps.
- Save to XXXX — This button will save the state and value of every single instrument to a text file denoted by the value XXXX.
- Exit (+ Isolate Cross) — This button is the “nice” way of exiting the software and ensures that valves V5, V6, and V7 are all closed when the program exits, thus, isolating

the cross and ensuring that the pressure inside stays constant.

- **Load Settings For** — The individual buttons underneath this heading contain predefined configurations for various operations. For each button underneath this heading, there is another button labeled “Def” which allows the user to define what is allowed, not allowed, and what the acceptable ranges are for the instruments for that particular configuration.
 - All Enabled/False — Pressing this button allows the shifter to operate any device (hence the “All Enabled” part of the label), but “expects” all valves to be closed, and pumps to be off (hence the “False” part). Opening a valve, or turning on a pump, results in this device being displayed in red as it is in contradiction to what it expects (everything to be off).
 - All Enabled/True — This is analogous to the previous button, but it expects each item to be open or on.
 - All Disabled/T:F — This button disables the operation of everything regardless of what state it may be in.
 - All Enabled/T:F — Pressing this button causes the software to allow any operation regardless of its state — the opposite of the previous button.
 - Operations — Selecting this button loads the acceptable values and states for the system when a source is inside of the detector. It disables, or strongly discourages, the operation of several critical components and will advise the operator that performing certain operations is not advisable until other items have been set or configured (for instance, turning on P2 at all when this configuration is selected).
 - Ops Reg (G5 as ref) — This is the default configuration during calibration operations, the configuration is identical to the previous button, except that the

pressure in the cross is now regulated by opening and closing V5 and V6. Gauge G5 is used to set the upper and lower pressure allowed in the cross. When the pressure exceeds the upper limit of G5, the software automatically closes valve V5, and opens valve V6. The bleed valve NV3 is set so that the pressure drops very slowly ~ 10 mbar/minute. When the pressure measured by G5 drops below the lower limit, the software simultaneously closes V6 and opens V5 — the impedance provided by regulator R3 is high enough that it takes approximately a minute to bring the pressure back up by 10 mbar. This configuration is used for safety reasons — if V5 fails to close due to a stuck valve or broken relay, when the pressure exceeds the maximum set by G5, valve V6 will open and allow the cross to vent until the pressure drops below the lower set point — the same holds if V6 should stick open (V5 would open to maintain the pressure in the cross).

- Purge (G5 as ref) — This button is used to automate the rather monotonous process of purging the cross after one or more of the flanges were removed. The shifter first turns on P2 and then presses this button and the software takes over and opens V7 until the pressure on G5 reaches -800 mbar, then closes V7, and opens V5 until G5 reads 5 mbar — the process then repeats itself, writing to an output file `purge.log.dat`, until a shifter interrupts the process (procedurally this process can be stopped after at least seven iterations). After the purging process has finished, the operator would put the system into regulated operations.
- Alarms — The system has two levels or alarms, “Local” and “Remote”, the former is the less severe of the two and is intended just for the operators in the clean room, and the latter is intended as an alarm in the control room and DAQ room. The check box for each of the two alarms enables them, and the “Def” button brings up the configuration window (see Figure A.5) to define what triggers an alarm.

The lower section of the interface shows the current preset values for each of the devices and the ranges of the instruments. When an alarm occurs, the items which are out of range flash bright red to help identify the problem area.

The remainder of the window shows the status of the system and allows the shifters to operate the components. The layout of the system is such that everything on the left are inputs and the right hand side of the glovebox are all exhaust lines of the system. To change the status of a component in the system, the user has to first left-click on an item — if it is not advisable to operate the item in the current configuration, the conflicting items will be surrounded with a red box to alert the shifter that proceeding is not advisable. Actually changing the state of the system requires that the shifter then right-click on the component, then the system will perform the desired operation.

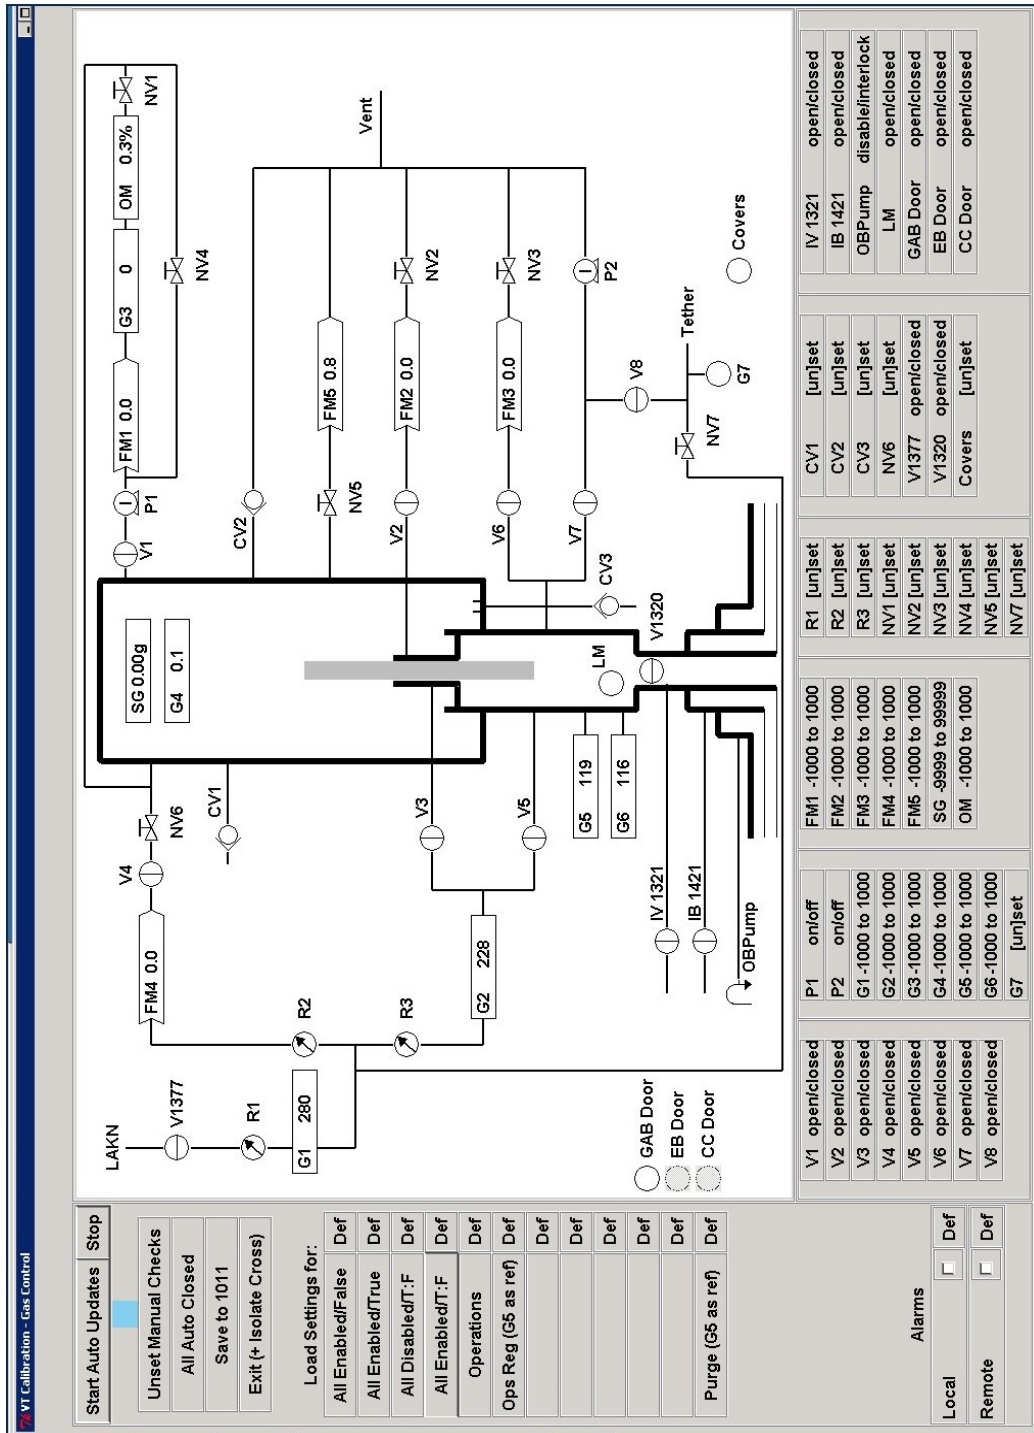


Figure A.4: User interface of the gas control portion of the source deployment system. The left panel controls the software and loads preset configurations, the bottom portion shows the set points for each instrument, and the main panel shows the present status of the system and allows the user to control the valves and pumps.

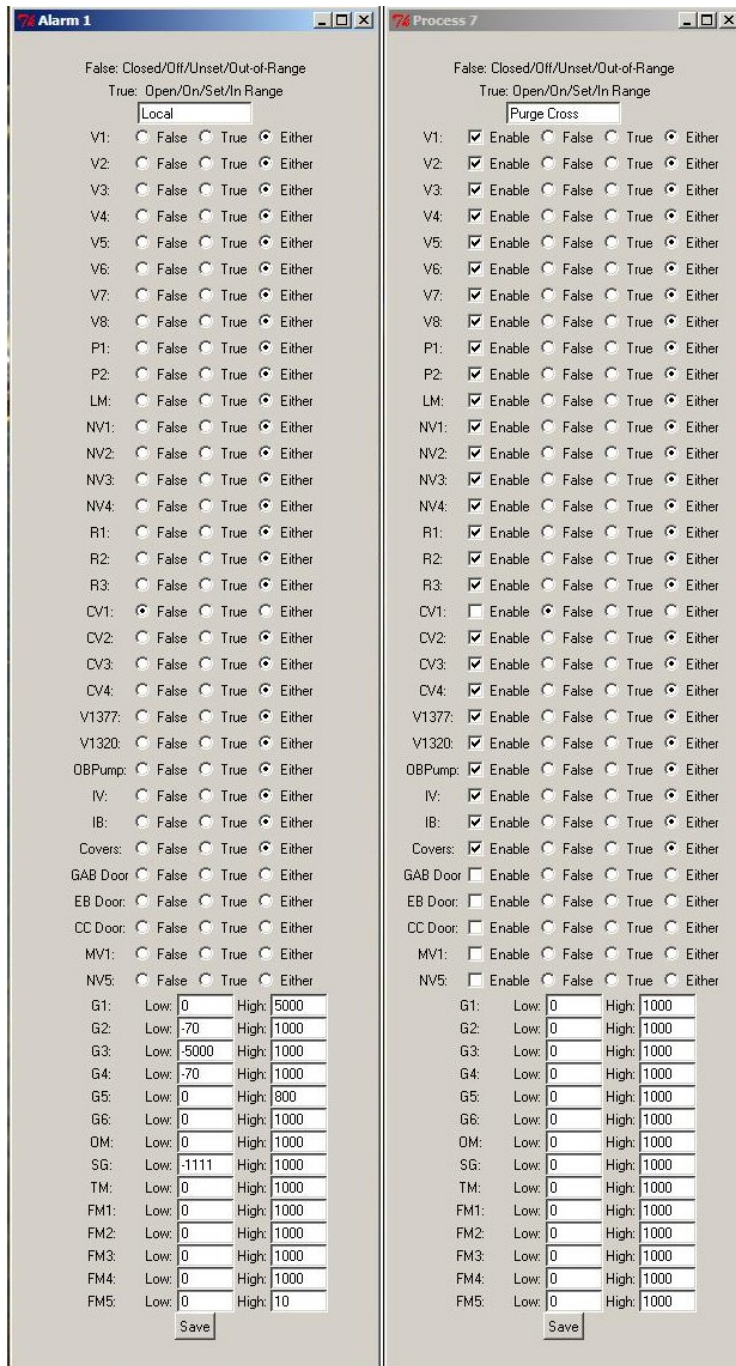


Figure A.5: Alarm and preset configuration portion of the user interface to the gas control system. The left pane allows the user to configure the parameters that will trigger an alarm. The right panel allows one to set up standard configurations for different system operations like evacuating the cross and backfilling it, purging the glovebox, etc.

Appendix B

Electronics Box Routing

Section A.2 explained the layout of the electronics box, which instruments are present, etc. Table B.1 shows the detailed routing table indicating which signals are routed to which terminal blocks in the event that any maintenance ever needs to be performed on the system.

Several abbreviations are used to minimize the width of the columns:

- Dig Gnd — Digital Ground
- P. — Power
- Sig — Signal (from a device)
- Com/NO — In reference to the relay devices: Common terminal, Normally Open terminal
- Alm — Alarm (meant to be connected to a buzzer, or strobe light)
- Gnd Bus — Ground bus

- kXY — The “k” is the National Instruments notation for a relay, the number following it (XY) is how each relay is identified by the software and on the connection board for the relays.
- (1/2)(A/B), YZ — Refers to which sets of I/O break-out boards the connection is to
 - set 1 is the unused, pre-wired, set for the eventual recabling of the camera system
 - set 2 is the one connected to the I/O module installed in the chassis. The letters A or B denote which of the two break-out boards within the set (1 or 2) is being referred to — the boards are all labeled, but A is always the leftmost of the two, and vice versa for the right. The number YZ determines which terminal the connection is to
- DCS — Digital Control System. This indicates that these items are connected to the main Borexino control system in the control room.
- TM — Tether Meter
- OM Current — An unused output of the oxygen monitor which would allow one to power a local LCD display for indicating the oxygen level.
- GAB/EB/CC Door — The door switches for the gas box, electronics box, and computer cabinet respectively.
- LM — Liquid monitor
- SW-X — Pre-wiring done for reading the status of any other switches we may wish to install.

Table B.1: Electronics Box Routing Table

Block	Purpose	Termination	Block	Purpose	Termination	Block	Purpose	Termination
B-1	Cam1, Lamps1	1A, 52	C-1	Cam P.	k10 Com	D-1	G3 P. +	
B-2	Cam1, Lamps2	1A, 17	C-2	Cam Coil	k10 NO	D-2	G3 P. -	
B-3	Dig Gnd	1A, 53	C-3	USB P.	k11 Com	D-3	G3 Sig +	2A, 68
B-4	Cam1, LED 1	1A, 29	C-4	USB Coil	k11 NO	D-4	G3 Sig -	2A, 34
B-5	Cam1, LED 2	1A, 27	C-5	RST P.	k12 Com	D-5	G2 P. +	
B-6	Dig Gnd	1A, 50	C-6	RST Coil	k12 NO	D-6	G2 P. -	
B-7	Cam1, Shutter	1A, 11	C-7		k13 Com	D-7	G2 Sig +	2A, 33
B-8	Dig Gnd	1A, 44	C-8		k13 NO	D-8	G2 Sig -	2A, 66
B-9	Cam2, Lamps1	1A, 19	C-9	Alm 1	k14 Com	D-9	G1 P. +	
B-10	Cam2, Lamps2	1A, 51	C-10	Alm 1	k14 NO	D-10	G1 P. -	
B-11	Dig Gnd	1A, 53	C-11	Alm 2	k15 Com	D-11	G1 Sig +	2A, 65
B-12	Cam2, LED 1	1A, 16	C-12	Alm 2	k15 NO	D-12	G1 Sig -	2A, 31
B-13	Cam2, LED 2	1A, 48	C-13			D-13	G5 P. +	
B-14	Dig Gnd	1A, 50	C-14	TM P. Bus		D-14	G5 P. -	
B-15	Cam2, Shutter	1A, 10	C-15			D-15	G5 Sig +	2A, 30
B-16	Dig Gnd	1A, 44	C-16	Gnd Bus		D-16	G5 Sig -	2A, 63
B-17	Cam3, Lamps1	1B, 52	C-17	Gnd Bus		D-17	G6 P. +	
B-18	Cam3, Lamps2	1B, 17	C-18	Gnd Bus		D-18	G6 P. -	
B-19	Dig Gnd	1A, 53	C-19	Gnd Bus		D-19	G6 Sig +	2A, 28
B-20	Cam3, LED 1	1B, 49	C-20	Gnd Bus		D-20	G6 Sig -	2A, 61
B-21	Cam3, LED 2	1B, 47	C-21	240V Gnd		D-21		
B-22	Dig Gnd	1A, 50	C-22	240V N		D-22		
B-23	Cam3, Shutter	1A, 43	C-23	240V L		D-23	F1 Sig +	2A, 60
B-24	Dig Gnd	1A, 44	C-24			D-24	F1 Sig -	2A, 26
B-25	Cam4, Lamps1	1B, 19	C-25			D-25		
B-26	Cam4, Lamps2	1B, 51	C-26			D-26		

Continued on next page

Table B.1: *continued*

Block	Purpose	Termination	Block	Purpose	Termination	Block	Purpose	Termination
B-27	Dig Gnd	1A, 53	C-27			D-27	F2 Sig +	2A, 25
B-28	Cam4, LED 1	1B, 16	C-28			D-28	F2 Sig -	2A, 58
B-29	Cam4, LED 2	1B, 48	C-29			D-29		
B-30	Dig Gnd	1A, 50	C-30			D-30		
B-31	Cam4, Shutter	1A, 42	C-31	G5 DCS	DCS	D-31	F3 Sig +	2A, 57
B-32	Dig Gnd	1A, 44	C-32	G5 DCS	DCS	D-32	F3 Sig -	2A, 23
B-33	Cam5, Lamps1	1B, 11	C-33			D-33		
B-34	Cam5, Lamps2	1B, 10	C-34			D-34		
B-35	Dig Gnd	1A, 15	C-35			D-35	F4 Sig +	2B, 68
B-36	Cam5, LED 1	1B, 43	C-36			D-36	F4 Sig -	2B, 34
B-37	Cam5, LED 2	1B, 42	C-37			D-37	F1-F5 P. +	N/A
B-38	Dig Gnd	1A, 18	C-38			D-38	F1-F5 P. -	N/A
B-39	Cam5, Shutter	1A, 41	C-39			D-39	F5 Sig +	2B, 33
B-40	Dig Gnd	1A, 13	C-40			D-40	F5 Sig -	2B, 66
B-41	Cam6, Lamps1	1B, 41	C-41			D-41	SG P. +	DR-30-12
B-42	Cam6, Lamps2	1B, 6	C-42			D-42	SG P. -	
B-43	Dig Gnd	1A, 15	C-43			D-43	SG Sig +	2B, 65
B-44	Cam6, LED 1	1B, 5	C-44			D-44	SG Sig -	2B, 31
B-45	Cam6, LED 2	1B, 38	C-45			D-45	OM P. +	
B-46	Dig Gnd	1A, 18	C-46			D-46	OM P. -	2B, 63
B-47	Cam6, Shutter	1A, 6	C-47	G4 P. +		D-47	OM Sig +	2B, 30
B-48	Dig Gnd	1A, 13	C-48	G4 P. -		D-48	OM Current	Not Used
B-49	Cam7, Lamps1	1B, 37	C-49	G4 Sig +	2B 26	D-49	V1 P.	k0 Com
B-50	Cam7, Lamps2	1B, 3	C-50	G4 Sig -	2B 61	D-50	V1 P.	k0 NO
B-51	Dig Gnd	1A, 15	C-51	TM D0	2A, 11	D-51	V1 Coil	Ground
B-52	Cam7, LED 1	1B, 45	C-52	TM D1	2A, 10	D-52	V2 P.	k1 Com
B-53	Cam7, LED 2	1B, 46	C-53	TM D2	2A, 43	D-53	V2 P.	k1 NO

Continued on next page

Table B.1: *continued*

Block	Purpose	Termination	Block	Purpose	Termination	Block	Purpose	Termination
B-54	Dig Gnd	1A, 18	C-54	TM D3	2A, 42	D-54	V2 Coil	Ground
B-55	Cam7, Shutter	1A, 5	C-55	TM D4	2A, 41	D-55	V3 P.	k2 Com
B-56	Dig Gnd	1A, 13	C-56	TM D5	2A, 6	D-56	V3 P.	k2 NO
B-57	Cam8, Lamps1	1B, 2	C-57	TM D6	2A, 5	D-57	V3 Coil	Ground
B-58	Cam8, Lamps2	1B, 40	C-58	TM D7	2A, 38	D-58	V4 P.	k3 Com
B-59	Dig Gnd	1A, 15	C-59	TM RST	2A, 52	D-59	V4 P.	k3 NO
B-60	Cam8, LED 1	1B, 1	C-60	TM OE	2A, 17	D-60	V4 Coil	Ground
B-61	Cam8, LED 2	1B, 39	C-61	TM SEL1	2A, 49	D-61	V5 P.	k4 Com
B-62	Dig Gnd	1A, 18	C-62	TM SEL2	2A, 47	D-62	V5 P.	k4 NO
B-63	Cam8, Shutter	1A, 38	C-63	GAB Door	2A, 37	D-63	V5 Coil	Ground
B-64	Dig Gnd	1A, 13	C-64	EB Door	2A, 3	D-64	V6 P.	k5 Com
B-65			C-65	CC Door	2A, 45	D-65	V6 P.	k5 NO
B-66			C-66	LM	2A, 46	D-66	V6 Coil	Ground
B-67			C-67	SW 4	2A, 2	D-67	V7 P.	k6 Com
B-68			C-68	SW 5	2A, 40	D-68	V7 P.	k6 NO
B-69			C-69	SW 6	2A, 1	D-69	V7 Coil	Ground
B-70			C-70	SW 7	2A, 39	D-70	V8 P.	k7 Com
B-71			C-71	+5V	2A, 14	D-71	V8 P.	k7 NO
B-72			C-72	+5V	2A, 8	D-72	V8 Coil	Ground
B-73			C-73			D-73	P1 P.	k8 Com
B-74			C-74			D-74	P1 P.	k8 NO
B-75			C-75			D-75	P1 Motor	Ground
B-76			C-76			D-76	P2 P.	k9 Com
E1	TM Gnd	Black Wire	C-77			D-77	P2 P.	k9 NO
E2	TM CHB	Blue Wire	C-78			D-78	P2 Motor	Ground
E3	TM CHA	White Wire	C-79			D-79		
E4	TM P.	Brown Wire	C-80			D-80		

Appendix C

On Axis Calibration Run List

Source	Run	X (cm)	Y (cm)	Z (cm)
^{14}C - ^{222}Rn	8793	0.2	-1.3	186.3
^{14}C - ^{222}Rn	8797/8	0.2	-2.5	285.6
^{14}C - ^{222}Rn	8799/8800	1.6	2.3	-299.4
^{14}C - ^{222}Rn	8801	1.8	0.4	-400.6
^{14}C - ^{222}Rn	8814/5	-1.7	0	392.7
^{14}C - ^{222}Rn	8817/9	1.9	-0.3	-4.4
^{14}C - ^{222}Rn	8822	-0.9	1.1	102.6
^{14}C - ^{222}Rn	8831/2	-1.7	2	-101.6
^{14}C - ^{222}Rn	8834-7	1.6	0.5	-172.1
^{54}Mn	8848	0	0	300
^{54}Mn	8849/50/1	0	0	0
^{54}Mn	8853/5	0	0	-300
^{85}Sr	8863/6/8/72/75	0	0	0
^{14}C	8881/3/7	0	0	0

Table C.1: Sources, DAQ run numbers, and positions of the sources during the on axis calibration from October 5 through October 10, 2008.

Appendix D

First Off Axis Calibration Run List

The list of sources, runs, and positions for the first off axis calibration are shown below in Table D.1.

Table D.1: January / February 2009 Off Axis Calibration

Source	Run	X (cm)	Y (cm)	Z (cm)
^{14}C - ^{222}Rn	9458	5.2	-1.9	300.5
^{14}C - ^{222}Rn	9459	7.6	-2.3	199.9
^{14}C - ^{222}Rn	9460	7.6	-5.2	99.3
^{14}C - ^{222}Rn	9461	7.5	-3.4	100.4
^{14}C - ^{222}Rn	9462	6.8	-3.9	0.1
^{14}C - ^{222}Rn	9464-7	5.4	-1.5	-99.9
^{14}C - ^{222}Rn	9468	6.7	-7.8	-199.9
^{14}C - ^{222}Rn	9469	96.3	18	-186.1
^{14}C - ^{222}Rn	9470	179.5	3.9	-152.6
^{14}C - ^{222}Rn	9471	232.9	6.8	-115.3
^{14}C - ^{222}Rn	9672-3	292.0	10.2	-54.3
^{14}C - ^{222}Rn	9474	352.9	13.2	63.7

Continued on next page

Table D.1: *continued*

Source	Run	X (cm)	Y (cm)	Z (cm)
$^{14}\text{C-}^{222}\text{Rn}$	9478	-21.5	51.2	-194.2
$^{14}\text{C-}^{222}\text{Rn}$	9479	-21.6	51.5	-194.8
$^{14}\text{C-}^{222}\text{Rn}$	9480/1/3/5	-21.3	50.3	-194.6
$^{14}\text{C-}^{222}\text{Rn}$	9486	-73.5	135.8	-162.5
$^{14}\text{C-}^{222}\text{Rn}$	9487	-103.1	182.9	-129.1
$^{14}\text{C-}^{222}\text{Rn}$	9488	-135.2	234	-74.6
$^{14}\text{C-}^{222}\text{Rn}$	9489	-174.6	293.7	44.7
$^{14}\text{C-}^{222}\text{Rn}$	9491	-29.8	-46.9	-194.4
$^{14}\text{C-}^{222}\text{Rn}$	9492-9501/3	-80.4	-128.3	-162.3
$^{14}\text{C-}^{222}\text{Rn}$	9505/6	-107.9	-175	-127.6
$^{14}\text{C-}^{222}\text{Rn}$	9507/8	-141.8	-226.2	-76
$^{14}\text{C-}^{222}\text{Rn}$	9510	-179.9	-286.5	42.3
$^{14}\text{C-}^{222}\text{Rn}$	9512	359.1	-30.8	130
$^{14}\text{C-}^{222}\text{Rn}$	9513	-80.1	346.8	127.3
$^{14}\text{C-}^{222}\text{Rn}$	9511	-255.3	-244	125.3
$^{14}\text{C-}^{222}\text{Rn}$	9514	-8.8	-5.4	-292.7
$^{14}\text{C-}^{222}\text{Rn}$	9515	60.6	-4.7	-292.8
$^{14}\text{C-}^{222}\text{Rn}$	9516	166.0	-13	-257.4
$^{14}\text{C-}^{222}\text{Rn}$	9517	249.9	-20.5	-195.8
$^{14}\text{C-}^{222}\text{Rn}$	9518	308.8	-24.5	-124.5
$^{14}\text{C-}^{222}\text{Rn}$	9519	348.6	-27.5	-38.3
$^{14}\text{C-}^{222}\text{Rn}$	9520	-27.4	66.3	-289.8
$^{14}\text{C-}^{222}\text{Rn}$	9521	-68.6	134.8	-261
$^{14}\text{C-}^{222}\text{Rn}$	9525	-112.7	211.9	-200.7
$^{14}\text{C-}^{222}\text{Rn}$	9529	-144.1	264.8	-128
$^{14}\text{C-}^{222}\text{Rn}$	9530	-165.7	301.3	-38.5
$^{14}\text{C-}^{222}\text{Rn}$	9531	-18.0	-31.5	-295.5
$^{14}\text{C-}^{222}\text{Rn}$	9532	-78.7	-122.4	-263.4
$^{14}\text{C-}^{222}\text{Rn}$	9565	-105.8	-209.9	-202.3
Continued on next page				

Table D.1: *continued*

Source	Run	X (cm)	Y (cm)	Z (cm)
$^{14}\text{C-}^{222}\text{Rn}$	9566	-140.4	-270.4	-117.9
$^{14}\text{C-}^{222}\text{Rn}$	9564	-292.0	-171.7	-41.6
$^{14}\text{C-}^{222}\text{Rn}$	9563	-161.2	-300	-40.6
$^{14}\text{C-}^{222}\text{Rn}$	9562	-39.2	-339.3	-40.9
$^{14}\text{C-}^{222}\text{Rn}$	9561	151.7	-309.9	-39.7
$^{14}\text{C-}^{222}\text{Rn}$	9559	313.9	-156.3	-38.2
$^{14}\text{C-}^{222}\text{Rn}$	9558	304.9	172.4	-36.1
$^{14}\text{C-}^{222}\text{Rn}$	9557	183.9	296.4	-37
$^{14}\text{C-}^{222}\text{Rn}$	9555	13.6	345.5	-39.6
$^{14}\text{C-}^{222}\text{Rn}$	9554	-296.0	167.2	-42.1
$^{14}\text{C-}^{222}\text{Rn}$	9550-9553	-338.8	-3.9	-42.6
$^{14}\text{C-}^{222}\text{Rn}$	9568	0.5	-1.7	-363.5
$^{14}\text{C-}^{222}\text{Rn}$	9571	80.2	-2	-353.8
$^{14}\text{C-}^{222}\text{Rn}$	9572	180.2	-3.6	-314.6
$^{14}\text{C-}^{222}\text{Rn}$	9573	-42.8	73.1	-351.2
$^{14}\text{C-}^{222}\text{Rn}$	9574	-91.1	142.1	-316.5
$^{14}\text{C-}^{222}\text{Rn}$	9575	-165.0	8.8	-316.7
$^{14}\text{C-}^{222}\text{Rn}$	9569	-49.1	-80.8	-347.8
$^{14}\text{C-}^{222}\text{Rn}$	9570	-88.1	-140.8	-317.2
$^{14}\text{C-}^{222}\text{Rn}$	9578	111.2	1.2	24.4
$^{14}\text{C-}^{222}\text{Rn}$	9579	174.6	2.2	65
$^{14}\text{C-}^{222}\text{Rn}$	9580	237.7	2	141
$^{14}\text{C-}^{222}\text{Rn}$	9581	265.4	1.8	211
$^{14}\text{C-}^{222}\text{Rn}$	9582	-60.1	106	32.2
$^{14}\text{C-}^{222}\text{Rn}$	9583	-90.2	155	72.3
$^{14}\text{C-}^{222}\text{Rn}$	9584	-119.6	204.3	147.1
$^{14}\text{C-}^{222}\text{Rn}$	9585	-132.6	225.4	213.9
$^{14}\text{C-}^{222}\text{Rn}$	9587	-53.5	-116.1	35.7
$^{14}\text{C-}^{222}\text{Rn}$	9589	-74.3	-157.8	69.7
Continued on next page				

Table D.1: *continued*

Source	Run	X (cm)	Y (cm)	Z (cm)
^{14}C - ^{222}Rn	9590	-99.2	-210.3	142.6
^{14}C - ^{222}Rn	9591	-111.0	-235.2	214.3
^{14}C - ^{222}Rn	9597/8	73.6	2.9	-88.3
^{14}C - ^{222}Rn	9596	170.4	5.7	-36.9
^{14}C - ^{222}Rn	9595	-32.7	57.9	-89.6
^{14}C - ^{222}Rn	9594	-78.8	141.8	-40.2
^{14}C - ^{222}Rn	9593	-35.4	-79.5	-82
^{14}C - ^{222}Rn	9592	-65.3	-140.3	-44.7
^{14}C - ^{222}Rn	9599	266.9	16.4	57.6
^{14}C - ^{222}Rn	9601	-131.2	225.7	56.7
^{14}C - ^{222}Rn	9602	-119.3	-230.2	56.1
^{14}C - ^{222}Rn	9603	78.1	0.6	117.8
^{14}C - ^{222}Rn	9604	134.0	2.9	164.1
^{14}C - ^{222}Rn	9605	169.1	2.90	235.6
^{14}C - ^{222}Rn	9506	-29	63	115.9
^{14}C - ^{222}Rn	9607	-59.3	117.4	166.4
^{14}C - ^{222}Rn	9608	-76.0	146.5	234.4
^{14}C - ^{222}Rn	9609	-34.7	-60.2	117.4
^{14}C - ^{222}Rn	9610	-66.0	-114.6	170.3
^{14}C - ^{222}Rn	9611	-82.1	-142.3	241.1
^{14}C - ^{222}Rn	9613	173.5	-4	300.3
^{14}C - ^{222}Rn	9615	-79.3	150.2	300.9
^{14}C - ^{222}Rn	9612	-83.3	-145.9	298.7
^{222}Rn In Vac	9630-35/7/8	5.7	6.7	0.6
^{222}Rn In Vac	9640-54	5.7	6.7	0.6
^{222}Rn In Vac	9655-65	169.1	10.6	6.8
^{14}C	9675-95	4.5	11.8	0.4
AmBe	9701/3-5/8	7.6	-7.0	0.9
AmBe	9710/1	7.6	-7.0	0.9

Appendix E

Second Off Axis Calibration Run List

Table E.1 below gives the list of sources, DAQ runs and source positions for all of the sources deployed during the second off axis calibration.

Table E.1: June 2009 Off Axis Calibration

Source	Run	X (cm)	Y (cm)	Z (cm)
^{14}C - ^{222}Rn	10301	6.3	- 1.4	342.7
^{14}C - ^{222}Rn	10302/4	6.1	-1.9	342.8
^{14}C - ^{222}Rn	10305	8.2	-1	292
^{14}C - ^{222}Rn	10306	13.2	-4.4	200.4
^{14}C - ^{222}Rn	10307/8	26.8	8.6	101.2
^{14}C - ^{222}Rn	10309	7.8	-9	0.8
^{14}C - ^{222}Rn	10310	10.3	-10.9	-98.6
^{14}C - ^{222}Rn	10311	136.6	15.5	-171.8
^{14}C - ^{222}Rn	10312	-60.7	91.7	-179.7
^{14}C - ^{222}Rn	10313	-43.5	-101.1	-179.2
^{14}C - ^{222}Rn	10314	-75.7	-192.6	-129.8
^{14}C - ^{222}Rn	10315	-120.8	-316.4	40.4
Continued on next page				

Table E.1: *continued*

Source	Run	X (cm)	Y (cm)	Z (cm)
$^{14}\text{C-}^{222}\text{Rn}$	10316	214.4	-3.2	-129.2
$^{14}\text{C-}^{222}\text{Rn}$	10317	350.5	2.1	52.8
$^{14}\text{C-}^{222}\text{Rn}$	10318	-89.1	199.8	-122.4
$^{14}\text{C-}^{222}\text{Rn}$	10319	-151.8	310.1	53.4
$^{14}\text{C-}^{222}\text{Rn}$	10320	10	8.6	-198.5
$^{14}\text{C-}^{222}\text{Rn}$	10321	6.4	8.4	-299.3
$^{14}\text{C-}^{222}\text{Rn}$	10322	-129.9	276.7	-123.2
$^{14}\text{C-}^{222}\text{Rn}$	10323	312	-3.1	-120.3
$^{14}\text{C-}^{222}\text{Rn}$	10324	-145.1	-267.6	-121.3
$^{14}\text{C-}^{222}\text{Rn}$	10349	-127	-214.1	-250.6
$^{14}\text{C-}^{222}\text{Rn}$	10350	257.3	-7.3	-248
$^{14}\text{C-}^{222}\text{Rn}$	10351	-128.5	-209.1	-250.7
$^{14}\text{C-}^{222}\text{Rn}$	10352	-76.4	-118.8	-339.1
$^{14}\text{C-}^{222}\text{Rn}$	10353	151.4	-11.5	-338.2
$^{14}\text{C-}^{222}\text{Rn}$	10354	-80.2	132	-334.1
$^{14}\text{C-}^{222}\text{Rn}$	10346	-43.4	3.9	-296.8
$^{14}\text{C-}^{222}\text{Rn}$	10345	1.2	49.8	-296.5
$^{14}\text{C-}^{222}\text{Rn}$	10344	52.3	-10.1	-296.5
$^{14}\text{C-}^{222}\text{Rn}$	10343	-7.7	-47.6	-296.4
$^{14}\text{C-}^{222}\text{Rn}$	10325	6.1	-148.3	-264.6
$^{14}\text{C-}^{222}\text{Rn}$	10326	156.5	-1.5	-263.7
$^{14}\text{C-}^{222}\text{Rn}$	10327	16.1	150.7	-264.5
$^{14}\text{C-}^{222}\text{Rn}$	10330	-142.6	4.8	-265.3
$^{14}\text{C-}^{222}\text{Rn}$	10331	-193.6	1	-224.2
$^{14}\text{C-}^{222}\text{Rn}$	10332	17.4	200.4	-223.3
$^{14}\text{C-}^{222}\text{Rn}$	10333	204.1	-12	-222.2
$^{14}\text{C-}^{222}\text{Rn}$	10334	1.6	-197.6	-223.9
$^{14}\text{C-}^{222}\text{Rn}$	10335	-13.1	-241.7	-168.1
$^{14}\text{C-}^{222}\text{Rn}$	10336	247.8	-22.8	-166.5

Continued on next page

Table E.1: *continued*

Source	Run	X (cm)	Y (cm)	Z (cm)
$^{14}\text{C-}^{222}\text{Rn}$	10337	34.4	243.5	-168.1
$^{14}\text{C-}^{222}\text{Rn}$	10338	-237.7	17.8	-168.8
$^{14}\text{C-}^{222}\text{Rn}$	10339	-265.8	48.6	-97.8
$^{14}\text{C-}^{222}\text{Rn}$	10340	46.2	273.6	-96.8
$^{14}\text{C-}^{222}\text{Rn}$	10341	275.9	-50.9	-95.6
$^{14}\text{C-}^{222}\text{Rn}$	10342	-14.2	-272.6	-97.9
$^{14}\text{C-}^{222}\text{Rn}$	10355	-103.2	146	72.5
$^{14}\text{C-}^{222}\text{Rn}$	10358	183.3	29.5	74.3
$^{14}\text{C-}^{222}\text{Rn}$	10359	-60.7	-167.2	72.6
$^{14}\text{C-}^{222}\text{Rn}$	10360	51.9	-251.5	141
$^{14}\text{C-}^{222}\text{Rn}$	10362	257.6	43.7	141.8
$^{14}\text{C-}^{222}\text{Rn}$	10363	-33.2	254.7	140.3
$^{14}\text{C-}^{222}\text{Rn}$	10364	-248.1	-44.7	138.7
$^{14}\text{C-}^{222}\text{Rn}$	10366	164.4	30.5	227.4
$^{14}\text{C-}^{222}\text{Rn}$	10367	5.1	163.7	225.9
$^{14}\text{C-}^{222}\text{Rn}$	10368	-159.3	-2	224.8
$^{14}\text{C-}^{222}\text{Rn}$	10369	9.9	-162.2	225.3
$^{14}\text{C-}^{222}\text{Rn}$	10376	74.3	-1	339
$^{14}\text{C-}^{222}\text{Rn}$	10377	-37.9	58.3	337.2
$^{14}\text{C-}^{222}\text{Rn}$	10378	-34.6	-60.6	338
$^{14}\text{C-}^{222}\text{Rn}$	10379	71.1	-0.5	290.8
$^{14}\text{C-}^{222}\text{Rn}$	10380	-1.2	97.5	283.1
$^{14}\text{C-}^{222}\text{Rn}$	10383	18	-67.7	289.1
$^{14}\text{C-}^{222}\text{Rn}$	10382	-64.2	2.9	289.6
$^{14}\text{C-}^{222}\text{Rn}$	10386	-39.7	55.9	216.2
$^{14}\text{C-}^{222}\text{Rn}$	10387	73.9	7.9	216.7
$^{14}\text{C-}^{222}\text{Rn}$	10388	-26.5	-63.1	216
$^{14}\text{C-}^{222}\text{Rn}$	10389	6.2	-5	192.3
$^{14}\text{C-}^{222}\text{Rn}$	10390	0.9	-2.8	292.9
Continued on next page				

Table E.1: *continued*

Source	Run	X (cm)	Y (cm)	Z (cm)
^{14}C - ^{222}Rn	10391	-0.2	-3.1	341.5
$^{85}\text{Sr} + ^{65}\text{Zn} + ^{60}\text{Co}$	10396-7	7.5	-9.0	1.6
$^{85}\text{Sr} + ^{65}\text{Zn} + ^{60}\text{Co}$	10399	-4.4	-1.9	-299.4
$^{85}\text{Sr} + ^{65}\text{Zn} + ^{60}\text{Co}$	10400	6.6	-3.3	300.4
^{203}Hg	10402	5	-2.3	282
^{203}Hg	10403	1.9	-6.7	147.2
^{203}Hg	10404-5	6.6	-10.8	0
^{203}Hg	10407	99.4	-3.8	-4.2
^{203}Hg	10408	135.4	-8.4	-9.7
^{203}Hg	10409	222.1	-58.5	-9.2
^{203}Hg	10410	-1.5	-8.5	-152.8
^{203}Hg	10411	-1	-7.1	-300.8
$^{54}\text{Mn} + ^{40}\text{K}$	10414	4.5	-4.5	299.4
$^{54}\text{Mn} + ^{40}\text{K}$	10415-6	8.6	-8.8	0.2
AmBe	10434/6/9	7.5	-3.3	292.5
AmBe	10440/1/3-5	0.3	-6.2	142.8
AmBe	10446/7/9	6.6	-7.1	-6.4
AmBe	10450-3	2.2	-13.1	-156.6
AmBe	10454/5/7/8	293.7	12.5	-11.4
AmBe	10460-3	260.1	13.4	-137.7
AmBe	10464-7	5.3	-9	-301.2
AmBe	10468/9/70/1	145.7	3.8	-346
AmBe	10473-6	310	20.6	-200
AmBe	10478/9/81/2	271.1	-8.7	216.1
AmBe	10483-7	130	-66	-14

Appendix F

Third Off Axis Calibration Run List

The list of runs and source positions for the third off axis calibration are presented in Table F.1. In Table F.1, the source labeled as “ ^{14}C - ^{222}Rn Inv” was a ^{14}C - ^{222}Rn source which was mounted upside down — with the neck pointing away from the source coupler, instead of at it. The idea was that perhaps the neck of the source was channeling a lot of light and pointing it at the source coupler where it was essentially lost. By mounting the source upside down, and deploying it at some of the same locations where normal ^{222}Rn sources were deployed, we can check for effects due to the orientation of the source. The data analysis of this source is discussed in detail in [124].

Table F.1: July 2009 Off Axis Calibration

Source	Run	X (cm)	Y (cm)	Z (cm)
^{57}Co	10546/7	5.9	-4	298.8
^{57}Co	10548	4.1	-7.3	150.3
^{57}Co	10550-2	1.9	-9.2	0.6
^{57}Co	10554	90.5	29.1	0
Continued on next page				

Table F.1: *continued*

Source	Run	X (cm)	Y (cm)	Z (cm)
^{57}Co	10557	145.5	49.1	1.4
^{57}Co	10558	217.7	74.6	-6.6
^{57}Co	10559	2.5	-7.7	-151.7
^{57}Co	10560	5.6	-4.7	-301.5
^{139}Ce	10562	6	-3.4	299.4
^{139}Ce	10563	4.1	-5	153
^{139}Ce	10564-6	9.8	-6.3	-0.5
^{139}Ce	10567	95.8	30.4	3.1
^{139}Ce	10568	138.4	59.1	-30
^{139}Ce	10569	243.6	105	86.6
^{139}Ce	10571	6.6	-4.9	-150.9
^{139}Ce	10572-3	5.9	-5.7	-300.9
^{85}Sr	10574-5	4.5	-4.2	299.1
^{85}Sr	10576	3.6	-5.9	148.4
^{85}Sr	10577/9	6	-9.8	-0.3
^{85}Sr	10580	104.3	30.4	6.2
^{85}Sr	10581	140.3	50	-0.2
^{85}Sr	10582	210.1	79.6	-17
^{85}Sr	10583	2.4	-7.2	-151.1
^{85}Sr	10584	8.8	-5.1	-299.8
^{14}C	10585-7	9.1	-5.3	-0.1
^{14}C - ^{222}Rn Inv	10606-7	4.2	-3.4	347.6
^{14}C - ^{222}Rn Inv	10608	4.8	-7	-1.8
^{14}C - ^{222}Rn Inv	10609	240.6	18.3	-102.7
^{14}C - ^{222}Rn Inv	10611	241.8	18.9	-103.4
^{14}C - ^{222}Rn Inv	10610	347.4	27.8	61.7
^{14}C - ^{222}Rn Inv	10612	314.5	21.4	-107.5
^{14}C - ^{222}Rn Inv	10613	115.6	5	-282.8
^{14}C - ^{222}Rn Inv	10615	2.9	-9.6	-353.7

Continued on next page

Table F.1: *continued*

Source	Run	X (cm)	Y (cm)	Z (cm)
AmBe	10631-2	8.5	-6.1	1.9
AmBe	10633-8	11.1	300	-13.1
AmBe	10639-44	-286	-32.9	-22.6
AmBe	10645-50	55.2	-288.7	-20
AmBe	10651	133.5	-334.4	-7.4
AmBe	10653	-332.3	-117.5	-10.9
AmBe	10654	-8.4	358.8	-6.8
AmBe	10655	362.6	-3.5	-5.6
AmBe	10656	189.6	179.8	-261.3
AmBe	10657	-242.2	16.7	-264.2
AmBe	10658	73.6	-245.2	-263.3
AmBe	10661/2/9/70-4/6-8/80/2-7	-2.2	5.6	-401.1
AmBe	10689/90/9	1.9	2.3	400.8
AmBe	10701-9/10-4	1.9	2.3	400.8
AmBe	10716	182.5	199.1	278.3
AmBe	10718	-257.7	60	274.3
AmBe	10719	107.2	-245.9	276.4
Laser	10747	4.1	-7.1	0.8
Laser	10744	Measured	Via	Source
Laser	10736	Measured	Via	Source
Laser	10738	Measured	Via	Source
Laser	10740	Measured	Via	Source
Laser	10742	Measured	Via	Source
Laser	10730	Measured	Via	Source
Laser	10733	Measured	Via	Source