The Borexino Nylon Film and the Third Counting Test Facility

Kevin B. McCarty

A dissertation presented to the faculty of Princeton University in candidacy for the degree of Doctor of Philosophy

Recommended for acceptance by the department of Physics

June 2006

© Copyright by Kevin B. McCarty, 2006. All rights reserved.

Abstract

Kevin B. McCarty

Advisor: Prof. Frank P. Calaprice

The Borexino solar neutrino detector should begin operations in late 2006. This scintillationbased detector will observe low-energy neutrinos, in real time, down to about 250 keV. The experiment should further tighten constraints on the neutrino oscillation parameters, and confirm the Standard Solar Model of solar neutrino production. It may also observe geoneutrinos; supernova neutrinos, should the timing of the experiment be fortunate; and perhaps other processes beyond the scope of the Standard Model of particle physics.

At the heart of Borexino lie 300 tons of organic scintillator fluid, contained by a spherical vessel composed of transparent nylon film. Roughly 300 tons of passive buffer fluid lie between this inner vessel and a second outer nylon vessel. Both vessels are located inside a steel sphere that also supports over 2000 inward-pointing photomultiplier tubes. The two most vital components of Borexino are these nylon vessels and the scintillator itself. Numerous measurements made at Princeton of the physical and radiochemical properties of the vessel film are reported in this thesis.

A 4-ton prototype of Borexino, the CTF, has been used to study scintillator radiopurity for over ten years. However, certain peculiarities of its design make determining the spatial positions of radioactive decays within the detector difficult. The development of a new position reconstruction code that takes these problems into account is reported herein.

Several studies of radiopurity in the latest version of CTF were made using this new code. These include a proposal for individually tagging decays of radon and four daughter isotopes; an attempt to detect convection using the radon daughters; a hypothesis to explain peculiar behavior of the crucial isotope ²¹⁰Po; and an analysis of ⁴⁰K contamination inside the detector based on models of the spatial distribution of external γ rays, leading to a somewhat concerning result. A final distillation test of the scintillator will take place shortly, and will match as closely as possible the procedure used to purify scintillator for the full Borexino detector. Analyses run on CTF data collected after this test should prove vital in understanding the detector sensitivity of Borexino.

Acknowledgments

No work is created in a vacuum, and this is especially true for anyone writing about aspects of a modern nuclear or particle physics experiment. This thesis builds on the work of hundreds of other people who have worked for the Borexino Collaboration, and that of tens of thousands who are responsible for the current state of the science of physics, the hardware and electronics used in the experiment, and the software used for data acquisition, reconstruction, and analysis. There is no space to thank all of them individually. I nevertheless want to acknowledge those people who have made the biggest difference to me personally and professionally.

Much of the analysis reported in this thesis was performed with Free Software. I therefore want to single out and mention the Free Software projects and people who have been the biggest help. Without the pioneering GNU tools, compilers, and General Public License provided by the Free Software Foundation, under the leadership of Richard M. Stallman, the world of physics analysis software would be a very different, and probably much poorer, place. The data acquisition system of the Borexino experiment is all based on the Linux kernel developed under the leadership of Linus Torvalds, used with the rich set of software distributed by the Debian Project. The analysis software I used mainly consisted of ROOT, authored by René Brun and Fons Rademakers, with some use of GEANT 4, distributed by the multinational GEANT 4 collaboration. Many of the figures in the text were created with XFig (by Brian Smith) or pre-processed with the GNU Image Manipulation Program (by Spencer Kimball and Peter Mattis). The text was prepared with teTFX, a version of Leslie Lamport's LATEX distributed by Thomas Esser and based on the original TEX by Donald E. Knuth. Particularly useful LATEX add-on packages included PSfrag by Craig Barratt, Michael Grant and David Carlisle; XMTEX by Shinsaku Fujita; and FeynMF by Thorsten Ohl.

Moving to the milieu of Borexino, I want to thank my Borexino Collaboration colleagues in the CTF data analysis group for all the stimulating ideas and discussions: Silvia Bonetti, Barbara Caccianiga, Davide d'Angelo, Davide Franco, Aldo Ianni, Emanuela Meroni, Maria Elena Monzani, Andrea Pocar, and Oleg Smirnov (with apologies to anyone whom I inadvertently missed). The collaboration members who have taught me the most about data acquisition hardware, software, and good coding practices are Cristian Galbiati, Marco Pallavicini, and Alessandro Razeto.

A number of other people at the Gran Sasso laboratory also made my several extended stays in Italy pleasant. The other Borexino folks based at LNGS who worked for the Princeton group at one time or another: Angelo Corsi, Ferenc Dalnoki-Veress, Steve Kidner, Antonio di Ludovico, Richard Ford, Augusto Goretti, Andrea Ianni, and Andrew Sonnenschein. Some other folks in the North American groups: Henning Back, Mark Chen, Bruce Vogelaar. The rest of the Borexino Collaboration. The always-friendly staff and researchers at the lab: Andrea Ford, Stefano Gazzana, Ersilia Giusti, Aldo Ianni, Matthias Laubenstein, Maria "del Bar" Mosca, and many others. I have memories of wonderful hikes in the nearby mountains at various times with Beth Harding, Mike Leung, Andrea Pocar, Domenic Schimizzi, Tom Shutt and Alex Wingler. I especially want to thank Richard and Andrea Ford, who were among my best friends in Italy, for many fine evenings and weekends. I'm happy to have met and befriended several of the people of the nearby city of L'Aquila, including Franca, Francesca, Marta, and Vincenzo. Finally, no one who knew him could forget Matt Pyle.

I next want to mention the following people who made a big difference in my time at Princeton. The inimitable Charles Sule, who was one of the first people I worked with in the lab at Princeton. Kamran Vakili and Walter Reisner, who helped me to survive qualifying exams. The Borexino vessel construction team: Julie Bert, Laura Cadonati, Richard Fernholz, Cristian Galbiati, Joel Greenberg, Beth Harding, Brian Kennedy, Ted Lewis, Allan Nelson, Andrea Pocar, Tom Shutt, and Charles Sule, who all managed to stay sane and even friendly despite a year of 12-hour days toiling in the clean room. Fred Loeser, who managed to escape the clean room by retiring. My other colleagues in the Princeton Borexino group: Jay Benziger, Frank Calaprice, Ernst de Haas, and Mike Leung. The staff of the Physics Department, who are all wonderful, with particular mention deserved by Claude Champagne, Paul LaMarche, Laurel Lerner, Kathy Patterson, Mike Peloso, Kathy Warren, John Washington, and especially the tireless Helen Ju, our group secretary. John Bahcall, a leader in the field of neutrino physics, whom I was fortunate enough to meet several times before his death in 2005. The National Science Foundation, which continued to support and believe in our experiment even when problems at the underground laboratory made the future look dark.

It would be remiss of me not to give special thanks to Cristian Galbiati, Beth Harding, and Andrea Pocar. In addition to being amazing and brilliant colleagues, they have been wonderful friends who supported me through good times and bad. My thesis advisor and mentor, Frank Calaprice, has always been compassionate, understanding, and concerned with my professional development.

Last but certainly not least, on a personal level I want to thank my parents, Marcia and Bryce, who have been supportive of my interests in science ever since their development at an early age; my brother Jason, who is heading into his own post-graduate long haul very soon; my extended family; my in-laws, Gene and Lee, who have been wonderful people to know since the first time we met; and my wife, best friend, and hiking partner Jennifer, who has been very patient with the long process of finishing this thesis. To them I am eternally grateful.

Introduction

This work discusses two important topics relating to the Borexino low-energy solar neutrino experiment. Before delving into the details, it gives a series of introductions to the prerequisite knowledge. Chapter 1 discusses the current state of neutrino physics. Chapter 2 describes the expected sensitivity of the Borexino detector to neutrinos originating from various sources, and Chapter 3 covers the physical design and components of the detector.

The first major topic is that of the thin nylon film making up the nested spherical scintillator containment vessels. Of the two vessels, the inner one contains 300 tons of organic scintillator fluid, the heart of the detector. The outer one encloses the inner one; between them is a region of buffer fluid. Beyond the outer vessel is another volume of buffer fluid, encased within a spherical steel structure that also supports over two thousand inwardpointing photomultiplier tubes. The design, construction, and installation of the vessels themselves are described briefly in a section of Chapter 3. Chapter 4 goes into detail about the material properties of the nylon film measured at Princeton University: its chemical properties, tensile strength, behavior when exposed to varying levels of humidity (both in air and immersed in scintillator), and innate level of radioactivity.

The scintillator fluid must, by its very nature, be the most radiopure component of the Borexino detector. A 4-ton prototype of Borexino has been operated for more than a decade, in three different incarnations, to study its radioactive contaminants. This Counting Test Facility (CTF) does not, unfortunately, have the sensitivity to prove once and for all that the Borexino scintillator is suitable, but only to point out potential areas of trouble with it. For reasons of practicality, the CTF incorporates some design elements that make data analysis difficult. Instead of an organic buffer fluid, for instance, water is used as a buffer in the CTF, leading to non-negligible refraction of light at the nylon vessel that separates the two liquids. The development of methods for accurately determining the positions of events in the CTF and the analysis of data from the third CTF incarnation, operated from 2001 to date, therefore comprise the last five chapters of this work.

Chapter 5 begins by discussing a theoretical likelihood-based approach to determining the expected spatial resolution of a scintillation-based detector. The ability of a detector to perform position reconstruction (determine the spatial position of an event) is important in separating events due to external radioactivity (on the surface of the detector, for instance) from those that actually represent the desired signal. Better spatial resolution implies a better separation efficiency. The results derived are general, and applicable to any large unsegmented scintillator-based detector. They may potentially be useful, therefore, in the design of new detectors intended to look for dark matter or to observe neutrinoless double β decay. The chapter concludes by presenting the expected radial distributions for "internal" events (those uniformly distributed throughout the scintillator) and "surface" events (those on the outer surface of the scintillator volume).

Chapter 6 describes the CTF itself. Topics covered include the physical design of the detector; the data acquisition system; the format of the data files (which is not, to my knowledge, documented anywhere else); and the history of the CTF and its previous results.

Chapters 5 and 6 provide the necessary background for Chapter 7, which discusses the many difficulties faced by any attempt to perform position reconstruction in the CTF. Some methods to overcome them are proposed, and these methods are compared using data from two sets of source calibration runs (data acquisition on a pointlike radioactive source moved to different known positions inside the CTF). A preferred algorithm to deal with the effects of refraction at the scintillator-water interface uses a numerical series expansion of a time-of-flight function for a light ray traveling from points inside the scintillator to points at the photomultiplier tube cathodes. The timing statistics of multiple photoelectrons detected at a single photomultiplier are also carefully considered. The performance of the preferred algorithm, developed by the author at Princeton, is compared to that of other position reconstruction codes that have been used with CTF data.

Finally, Chapters 8 and 9 present results for the radioactive contamination of the CTF scintillator and vessel, respectively, that were obtained using the reconstruction code described in Chapter 7. Particular attention is paid to the heavy-element decay chains ²³⁸U and ²³²Th. A likelihood-based method of tagging five isotopes in the ²³⁸U decay chain (²²²Rn, ²¹⁸Po, ²¹⁴Pb, ²¹⁴Bi, ²¹⁴Po) is proposed, and through Monte Carlo simulations and data analysis, found to be quite promising for use in Borexino. The distribution and behavior of the isotope ²¹⁰Po (an indicator of the more problematic ²¹⁰Bi) are found to be consistent with the presence of small particulate matter in the scintillator. In addition, the distribution of γ rays produced by the electron-capture decay of ⁴⁰K is simulated with three different models; two of the models yield a result suggesting significant ⁴⁰K contamination of the scintillator. A final distillation test of the scintillator, to take place within the next few months, will match as closely as possible the purification procedure used in the full Borexino detector. Analyses of CTF data collected after this test, performed in a similar way to those discussed in this work, should prove vital in determining the effectiveness of the distillation.

This work, despite its length, could not possibly discuss all aspects of the Borexino experiment and the CTF in an even-handed manner. For in-depth coverage of related topics of interest, the reader is referred to the literature on the Borexino experiment.

Contents

\mathbf{A}	Abstract ii					
Acknowledgments v						
In	trod	uction		viii		
1	Intr	oducir	ng the Neutrino	1		
	1.1	The S ⁻	tandard Model of particle physics	2		
		1.1.1	Quarks	2		
		1.1.2	Leptons	5		
	1.2	Neutri	ino mass and oscillations	6		
		1.2.1	Neutrino oscillations in vacuum	8		
		1.2.2	The Mikheyev-Smirnov-Wolfenstein effect in matter $\hdots\dots\dots\dots\dots$	13		
	1.3	Solar i	neutrinos	20		
		1.3.1	The pp cycle \ldots	21		
		1.3.2	The CNO cycle	25		
	1.4	Remai	ning questions	27		
		1.4.1	Absolute values and origin of neutrino masses	27		
		1.4.2	Sterile neutrinos and LSND	34		
		1.4.3	Neutrinoless double beta decay	35		
		1.4.4	New physics accessible to a low-energy solar neutrino detector $\ . \ .$	38		

2 The Borexino Neutrino Experiment

41

	2.1	Design parameters					
	2.2	Radio	active backgrounds	46			
		2.2.1	Types of radioactive decay	46			
		2.2.2	The isotopes present in Borexino	58			
		2.2.3	The heavy element decay chains	66			
		2.2.4	Categories of radioactive background	73			
	2.3	Prosp	ects for observing solar neutrinos	75			
		2.3.1	The ⁷ Be neutrinos	76			
		2.3.2	Annual variations in the $^7\mathrm{Be}$ signal	80			
		2.3.3	The <i>pep</i> and CNO cycle neutrinos	82			
	2.4	Prosp	ects for observing other neutrinos	85			
		2.4.1	Geoneutrinos and reactor antineutrinos	85			
		2.4.2	Supernova neutrinos	88			
3	\mathbf{Des}	ign an	d Hardware of Borexino	91			
	3.1	The so	cintillator and buffer fluids	93			
		3.1.1	The scintillator fluid	93			
		3.1.2	α particle quenching	99			
		3.1.3	The buffer fluid	101			
	3.2						
		Scintil	llator containment vessels	103			
		Scintil 3.2.1	Ilator containment vessels Requirements for the vessels	103 103			
		Scintil 3.2.1 3.2.2	Ilator containment vessels	103 103 106			
		Scintil 3.2.1 3.2.2 3.2.3	Ilator containment vessels	103 103 106 112			
		Scintil 3.2.1 3.2.2 3.2.3 3.2.4	Ilator containment vessels	 103 103 106 112 117 			
	3.3	Scintil 3.2.1 3.2.2 3.2.3 3.2.4 Photo	Ilator containment vessels	 103 103 106 112 117 124 			
	3.3	Scintil 3.2.1 3.2.2 3.2.3 3.2.4 Photo 3.3.1	Ilator containment vessels	 103 103 106 112 117 124 124 			
	3.3	Scintil 3.2.1 3.2.2 3.2.3 3.2.4 Photo 3.3.1 3.3.2	Ilator containment vessels	 103 103 106 112 117 124 124 126 			

4	Nyl	on Fil	m of the Borexino Vessels	133
	4.1	Chem	ical properties of nylon	135
		4.1.1	Nomenclature and structure	135
		4.1.2	Crystallinity	139
		4.1.3	Nylon hydrolysis	140
		4.1.4	Candidate materials for Borexino	140
	4.2	Polym	ner mechanical properties	143
		4.2.1	Glass transition in nylons	143
		4.2.2	Tensile strength, Young modulus, and creep	145
		4.2.3	Griffith model of brittle fracture	147
		4.2.4	Surface energy and critical thickness	149
	4.3	Testin	g material properties of nylon film	151
		4.3.1	The Controlled Relative Humidity Facility and nylon samples	152
		4.3.2	Tests of moisture content	154
		4.3.3	Tests of material strength	158
		4.3.4	Tests of creep	162
		4.3.5	Tests of brittleness	165
	4.4	Radio	active contamination of nylon surfaces	171
		4.4.1	Intrinsic activity in nylon film	171
		4.4.2	Radon diffusion through nylon	173
		4.4.3	Radon emanation from nylon	178
		4.4.4	Desorption of radioactive contaminants from nylon	181
	4.5	Solubi	ility of water in pseudocumene	185
		4.5.1	Water solubility in aromatic hydrocarbons	186
		4.5.2	Experimental tests of water solubility in pseudocumene	187
	4.6	Conclu	usions regarding the nylon films	189

5 Position Reconstruction in Scintillation Detectors

191

	5.1	The need for spatial reconstruction					
	5.2	Likelih	nood function derivation \ldots	193			
		5.2.1	Factoring the detector likelihood function	194			
		5.2.2	Scintillator dispersion time at the emission point	196			
		5.2.3	Photon attenuation	197			
		5.2.4	The PMTs not triggered	199			
		5.2.5	Specialization to a spherical detector	199			
	5.3	Analy	tical treatment of the likelihood function	200			
		5.3.1	Taylor expansion of the likelihood function	201			
		5.3.2	Likelihood function maximum and resolutions	203			
		5.3.3	Pattern matching	204			
		5.3.4	Comparison to observed resolutions	205			
	5.4	Multip	ble PMT occupancy	208			
		5.4.1	Correcting for timing bias	208			
		5.4.2	Effects on detector resolution	210			
	5.5	Expec	ted spatial distribution of reconstructed events	213			
		5.5.1	Calculating the expected observed event distribution	213			
		5.5.2	Radially symmetric event distributions	214			
		5.5.3	Previously used radial distribution functions	224			
		5.5.4	Cylindrical geometries	227			
6	The	Borez	kino Counting Test Facility	231			
	6.1	CTF p	physical characteristics	232			
		6.1.1	Details of the CTF design	234			
		6.1.2	The CTF photomultiplier tubes	237			
		6.1.3	The muon veto system	241			
	6.2	CTF d	lata acquisition	242			
		6.2.1	Group 1 and Group 2 events	242			

		6.2.2	Event energy: the ADC channels	243
		6.2.3	Event timing: the TDC channels	243
		6.2.4	Other data channels	245
	6.3	File st	ructure of events	248
		6.3.1	Note on different numerical formats used	249
		6.3.2	Structure of raw data files	251
		6.3.3	Reconstructed events	256
	6.4	Histor	y of the CTF	258
		6.4.1	CTF 1: The test of feasibility	259
		6.4.2	CTF 2: Tests with a new scintillator	263
		6.4.3	CTF 3: Purification and background tests	268
7	Wri	ting a	Position Reconstruction Code for the CTF 2	271
	7.1	Previo	bus approaches to CTF position reconstruction	273
		7.1.1	An effective index of refraction	273
		7.1.2	Monte Carlo techniques	274
	7.2	Outlin	e of the CTF reconstruction software	275
		7.2.1	Input of raw data	276
		7.2.2	Reconstruction code	277
		7.2.3	Output of reconstructed data	278
	7.3	Shared	d electronics channels	279
		7.3.1	Shared channels in Group 1 events	279
		7.3.2	Shared channels in Group 2 events	282
		7.3.3	Results using the approximate weight function	283
	7.4	Perfor	mance of spatial pattern recognition alone	284
	7.5	The so	cintillator PDF	287
	7.6	A sem	i-analytical ray tracing approach	290
		7.6.1	The path from event to PMT	291

		7.6.2	Enforcing Snell's law	293
		7.6.3	The problem of dark zones	298
	7.7	Testin	g algorithms with CTF 2 source runs	308
		7.7.1	The source calibration hardware	308
		7.7.2	Accuracy of reconstruction software with 214 Po events \ldots .	310
		7.7.3	Precision of reconstruction software with 214 Po events \ldots	316
		7.7.4	Data contamination by mis-reconstructed surface events	317
		7.7.5	Results with ²¹⁴ Bi events	320
	7.8	Testin	g algorithms with CTF 3 source runs	324
		7.8.1	The source calibration hardware	325
		7.8.2	The CTF 3 source run results	326
		7.8.3	Position resolution for CTF 3 source runs	329
8	Inte	ernal C	Contamination in the CTF	333
	8.1	Partic	le identification techniques	335
		8.1.1	Event position	336
		8.1.1 8.1.2	Event position Event energy	336 337
		8.1.18.1.28.1.3	Event position Event energy α/β discrimination	336 337 340
		8.1.18.1.28.1.38.1.4	Event position \ldots	336337340349
	8.2	8.1.18.1.28.1.38.1.4Heavy	Event position	 336 337 340 349 353
	8.2	 8.1.1 8.1.2 8.1.3 8.1.4 Heavy 8.2.1 	$ \begin{array}{l} \text{Event position} & \dots & $	 336 337 340 349 353 353
	8.2	 8.1.1 8.1.2 8.1.3 8.1.4 Heavy 8.2.1 8.2.2 	Event position	 336 337 340 349 353 353 359
	8.2	8.1.1 8.1.2 8.1.3 8.1.4 Heavy 8.2.1 8.2.2 8.2.3	Event position	 336 337 340 349 353 353 359 362
	8.2	 8.1.1 8.1.2 8.1.3 8.1.4 Heavy 8.2.1 8.2.2 8.2.3 The definition of the second secon	Event position	 336 337 340 349 353 353 359 362 365
	8.2	 8.1.1 8.1.2 8.1.3 8.1.4 Heavy 8.2.1 8.2.2 8.2.3 The de 8.3.1 	Event position	 336 337 340 349 353 353 359 362 365 366
	8.2	 8.1.1 8.1.2 8.1.3 8.1.4 Heavy 8.2.1 8.2.2 8.2.3 The de 8.3.1 8.3.2 	Event position	 336 337 340 349 353 353 359 362 365 366 370
	8.2	8.1.1 8.1.2 8.1.3 8.1.4 Heavy 8.2.1 8.2.2 8.2.3 The de 8.3.1 8.3.2 8.3.3	Event position	 336 337 340 349 353 353 359 362 365 366 370 374

Bi	ibliog	graphy		452
Co	onclu	ıding I	Remarks	450
		9.4.2	Determining the internal 40 K contamination	447
		9.4.1	The spatial distribution of external γ rays	442
	9.4	Extern	nal gamma rays	441
	9.3	210 Po	on the vessel	435
		9.2.2	212 BiPo coincidences: evidence for thorium $\ldots \ldots \ldots \ldots \ldots$	431
		9.2.1	$^{214}{\rm BiPo}$ coincidences and the implications for radon $~\ldots \ldots \ldots$.	422
	9.2	Rador	and thorium in or near the nylon film $\ldots \ldots \ldots \ldots \ldots \ldots$	422
		9.1.3	Attempts to observe the hold-down ropes	419
		9.1.2	Hot spots on the nylon film	417
		9.1.1	The north and south end regions	416
	9.1	Hot sp	pots on the CTF vessel	414
9	\mathbf{CT}	F Surf	ace Contamination and External Backgrounds	413
		8.5.2	Long-lived natural radioisotopes: $^{87}\mathrm{Rb}$ and $^{40}\mathrm{K}$	409
		8.5.1	The noble gases: 85 Kr and 39 Ar \ldots	407
	8.5	Lighte	er radioactive isotopes	405
		8.4.4	Looking for scintillator convection currents with tagged events	401
		8.4.3	Characteristics of the tagged events	396
		8.4.2	Monte Carlo simulation of the analysis	392
		8.4.1	Selection of candidate events	389
	8.4	Taggii	ng 222 Rn and its daughters $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots$	388
		8.3.5	Spatial distribution of 210 Po events	384
		8.3.4	Recent behavior of the 210 Po activity	381

Chapter 1

Introducing the Neutrino

The neutrino is perhaps the most negligible piece of matter that can be imagined. It is immune to effects of the electromagnetic and strong nuclear forces. Its mass, originally thought to be zero, is less than one-millionth that of the electron. The only one of the four fundamental forces that can affect it significantly is the weak nuclear force, the least powerful of the short-range interactions. When its existence was first postulated by Pauli in 1931, it was thought of as a sort of ghost particle whose offense against Occam's razor was deemed only slightly more palatable than the violation of matter conservation that would otherwise occur in some radioactive decays.

From this inauspicious beginning, the neutrino became one of the most-studied particles of the last quarter of the twentieth century. At the beginning of the twenty-first, the interest of the scientific community shows no sign of waning. Paradoxically, much of this interest results from the same reluctance to interact with other matter that made the neutrino so difficult to discover in the first place. Neutrinos produced at the center of the Earth, the Sun, or even in supernovae can tell us a lot about conditions in those otherwise inaccessible places—once we solve the problems of detecting them, that is!

The neutrino is even more fascinating because it is one of the first particles directly observed to contradict the predictions of the highly successful Standard Model of particle physics. It is not supposed to have a mass, but it does. It is supposed to exist as three separate species or "flavors," but these have been seen to convert back and forth between each other, or "oscillate." Recent (still controversial) experiments suggest that the neutrino may be its own antiparticle, violating the law of conservation of lepton number. These and other properties are the first direct observational evidence we have for physics beyond the Standard Model.

1.1 The Standard Model of particle physics

Before the recent revelations of non-zero neutrino mass and the consequent oscillations, the neutrino was thought to be well-understood. In the universe of the Standard Model, there are two fundamental types of fermions, the half-integer spin particles commonly thought of as matter (as opposed to bosons, whole-integer spin particles such as photons that mediate the exchange of forces). The fundamental fermions of the Standard Model are summarized in Table 1.1, and the force-carrying bosons are listed in Table 1.2.

1.1.1 Quarks

Fermions of one type, quarks, are never observed individually in nature or in any physics experiment to date since they have a property called "color charge." This charge has nothing to do with the spectrum of visible light. It is a convenient way of referring to the SU(3) "color" group embodied by the gluons that mediate quark-quark interactions in a way called the strong nuclear force.

The Standard Model requires that observable particles be colorless, because the potential energy between two particles with color charge is roughly a linearly *increasing* function of the distance between them. This property is called quark confinement. There are two ways of enforcing colorlessness. A combination of three quarks (qqq), each having one of the three different color charges, is colorless. As a pun on optics, the three color charges

		Quarks		Leptons		
Family	Particle	Charge $[e]$	Mass [MeV]	Particle	Charge $[e]$	Mass [MeV]
1	u	$+^{2}/_{3}$	~ 3	ν_e	0	$< 3 \times 10^{-6}$
	d	$-\frac{1}{3}$	~ 6	e^-	-1	0.5
2	c	$+^{2}/_{3}$	~ 1250	$ u_{\mu}$	0	< 0.2
2	s	$-\frac{1}{3}$	~ 100	μ^-	-1	105.7
2	t	$+^{2}/_{3}$	$1.73 imes 10^5$	$ u_{ au}$	0	< 18.2
5	b	$-\frac{1}{3}$	$\sim \! 4250$	$ au^-$	-1	1777.0

Table 1.1: The twelve fundamental fermions of the Standard Model. (Their antiparticles are not shown.) Each has a spin of $\frac{1}{2}$ in units of \hbar . Since the quarks are bound by color confinement, estimates of their masses are uncertain, especially for the lightest two. Estimates of the neutrino mass upper limits are actually for effective masses for the weak eigenstates, or equivalently, weighted averages over the mass eigenstates. Masses are taken from reference [1], except for the t mass [2].

Interaction	Particle	Symbol	Charge $[e]$	Spin $[\hbar]$	Mass [MeV]
Electromagnetic	Photon	γ	0	1	0
Weelt	W bosons	W^{\pm}	± 1	1	8.04×10^4
weak {	Z boson	Z^0	0	1	$9.12 imes 10^4$
Strong	Gluon (8 types)	g	0	1	0
Gravitational	Graviton		0	2	0
-	Higgs boson	H^0	0	0	$> 1.14 \times 10^5$

Table 1.2: The three fundamental interactions of the Standard Model, and their associated carrier bosons. Gravity is also listed, for completeness, although it is not in the scope of the Standard Model and gravitons have never been observed. The predicted scalar Higgs boson, not associated with an interaction, is listed as well. Masses are taken from reference [1]. The listed values of 0 are theoretical, although there are very strong upper limits for the photon ($< 6 \times 10^{-17} \,\mathrm{eV}$) and graviton ($< 7 \times 10^{-32} \,\mathrm{eV}$) based on astronomical observations.

are traditionally called red, green and blue. This arrangement forms a baryon, such as the familiar proton or neutron. Alternatively, a quark may pair with an antiquark having a complementary anti-color (red and anti-red, for instance). This arrangement $(q\bar{q})$ is called a meson. The exchange of π -mesons, or pions, between protons and neutrons in an atomic nucleus is one way in which they are held together against electromagnetic repulsion. Mesons and baryons are generically called hadrons.

Irrespective of their color charges, quarks come in six known types or "flavors," divided into three families: (u, d), (c, s), (t, b). Gluons can affect only color charges, not transform quarks between flavors. Each family consists of one quark with an electric charge of $+\frac{2}{3}$ (in units of e) and one with a charge of $-\frac{1}{3}$. The corresponding antiquarks have opposite charges. It can easily be seen that quarks can only combine in a way that produces an integer electric charge. For instance, the neutron (charge 0) comprises two d quarks ($-\frac{1}{3} \times 2$) and a uquark ($+\frac{2}{3}$), udd; the proton (charge +1) is uud. The π^+ is made up of $u\bar{d}$. Gluons have no electric charge.

As well as the strong nuclear force and electromagnetism, quarks are also susceptible to the weak nuclear force. The interaction of the charged weak force carriers, the W^{\pm} particles, with quarks may cause them to change flavor. For instance, one possible Feynman diagram vertex of the weak interaction is $u + W^- \rightarrow d$. One might expect that the flavor change would only occur within the same family $(d \leftrightarrow u, s \leftrightarrow c, b \leftrightarrow t)$. Curiously, however, the weak eigenstates of quarks are not precisely identical to their flavor eigenstates. The transformation between the two sets of eigenstates is the Cabibbo-Kobayashi-Maskawa (CKM) matrix:

$$\begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} |d\rangle \\ |s\rangle \\ |b\rangle \end{pmatrix} = \begin{pmatrix} |d'\rangle \\ |s'\rangle \\ |b'\rangle \end{pmatrix}$$
(1.1)

Primed states represent the weak eigenstates. The constraint of unitarity and the fact that complex phases can be absorbed into each eigenstate mean that there are really only four independent real parameters. Three are "mixing angles" θ_{us} , etc., and the fourth

is a complex phase δ that implies the existence of CP-violating phenomena. The CKM matrix is approximated by the identity matrix, with the largest off-diagonal terms being $V_{us} \approx V_{cd} \approx 0.22$.

In short, the existence of a non-trivial CKM matrix means that the Feynman diagram vertex $c+W^- \rightarrow d$ is possible. (Properly speaking this is really an interaction $c+W^- \rightarrow s'$, with subsequent observation of the flavor of $|s'\rangle$ yielding the flavor eigenstate $|d\rangle$ with probability $|V_{cd}|^2$.) As a direct result of such flavor-changing weak decays, particles containing the heavier quarks of the second and third families decay into those made only of the lighter u and d quarks. Hence in the universe at large, almost all the hadrons observed consist of combinations of u and d quarks and antiquarks, the members of the first quark family.

1.1.2 Leptons

Leptons, like quarks, are a fundamental type of fermion. The primary difference is that all leptons are colorless and unaffected by the strong nuclear force. As with quarks, there are also three families of leptons. Each family includes one particle with electric charge -1 and one particle with no charge. The charged leptons are named the electron, muon and tau $(e^-, \mu^- \text{ and } \tau^-)$, while the uncharged leptons are collectively called neutrinos. When it is necessary to distinguish between them, they are referred to as the electron neutrino, and so on (symbolized as the letter ν with the appropriate subscript). Writing the lepton families as pairs gives a familiar format (ν_e, e^-) , (ν_μ, μ^-) , (ν_τ, τ^-) similar to the quark families. This is no accident; each pair of particles is a weak isospin doublet. (Technically, this is only true for left-handed particles, as described in the next section.)

Like quarks, the leptons may be transformed into one another via the weak interaction. For instance, $\mu^- + W^+ \rightarrow \nu_{\mu}$. In fact (there being no free W^{\pm} particles, and none of the leptons having a mass as large as the W) two such vertices of the weak force must be combined to yield a real-world decay or interaction, *e. g.*, $\mu^- \rightarrow e^- + \nu_{\mu} + \bar{\nu}_e$. If the classic Standard Model held, neutrinos would have no mass. There would be only one possible measurement to perform on a neutrino that could distinguish between flavors: determining with which type of charged lepton it would interact.¹ In the classic SM, therefore, only the weak eigenstates of the neutrino are relevant, and there is no lepton equivalent of the CKM matrix.

1.2 Neutrino mass and oscillations

Until the last decade or so, it was believed that the masses of the neutrinos were exactly zero. The standard Hamiltonians for each flavor in a vacuum were therefore identical. As a result, a quantum state $|\nu\rangle$ consisting of a superposition $c_e |\nu_e\rangle + c_\mu |\nu_\mu\rangle + c_\tau |\nu_\tau\rangle$ of neutrino flavors would evolve over time in such a way as to preserve the amplitudes of the coefficients $c_{e,\mu,\tau}$; the only change in the state would be in the overall complex phase of the wave function. In particular, for a neutrino created in a weak eigenstate—one of the coefficients set to one and the others to zero, as should be the case with any standard weak interaction—the state would remain in that eigenstate forever. Therefore the individual lepton numbers L_e, L_μ, L_τ would be individually conserved.

If neutrinos have mass, contrary to the original formulation of the Standard Model, and if the mass eigenstates differ from the weak eigenstates, the situation changes. There exists a 3×3 matrix, which has been named the Maki-Nakagawa-Sakata (MNS) matrix, that describes the mixing between the mass eigenstates $|\nu_i\rangle$ (i = 1, 2, 3) and the weak eigenstates $|\nu_{\alpha}\rangle$ ($\alpha = e, \mu, \tau$):

$$\begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \\ |\nu_3\rangle \end{pmatrix} = \begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \\ |\nu_\tau\rangle \end{pmatrix}$$
(1.2)

¹Given a large number of neutrinos of the same flavor, one could also identify them in bulk by measuring their cross section for scattering on specific particles.

The neutrino mixing matrix is usually parametrized by the angles θ_{12} , θ_{13} , and θ_{23} and the phase δ . One common parametrization is given by [1]

$$U_{\rm MNS} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}$$
$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}, \quad (1.3)$$

with c_{ij} defined as $\cos \theta_{ij}$, and $s_{ij} \equiv \sin \theta_{ij}$. If non-zero, the phase δ would cause CPviolating physical effects. However, these have not yet been experimentally observed. All terms containing the CP-violating factors $e^{\pm i\delta}$ are proportional to the sine of θ_{13} , which is known experimentally to be small.

At first it may appear that the CKM and MNS matrices produce different effects. The CKM matrix results in the mixing of well-defined quark flavors during weak interactions, while the MNS matrix (as we will discuss momentarily) results in the mixing of well-defined neutrino flavors into different masses. This apparent difference is an illusion caused by psychology. We principally think about quarks in terms of the flavors that define hadron compositions; these are also mass eigenstates. The weak mixing that occurs via the CKM matrix is usually thought of as a secondary effect. On the other hand, the primary definition of neutrino flavors has always been in terms of the weak interaction, the only interaction that affects neutrinos: a neutrino that interacts with an electron to give a muon was by definition a μ -neutrino. In this case, the oscillation is thought of as the secondary effect— after all, we have no direct way yet to observe the neutrino mass eigenstates. But in both Equations (1.1) and (1.2), the vector on the left consists of the mass eigenstates, and is converted by the appropriate matrix to the vector on the right consisting of the weak eigenstates. Both mixings result from the non-correspondence between mass eigenstates (which in the case of quarks we think of as flavors) and weak eigenstates.

1.2.1 Neutrino oscillations in vacuum

Since a neutrino in a weak eigenstate (specific flavor) is a superposition of different mass eigenstates, as it travels through a vacuum (or any suitably diffuse material) the coefficients of the weak eigenstates will evolve. The wave function of the neutrino at any time after its formation may be described by expansion in terms of energy eigenstates:

$$|\nu(\mathbf{x},t)\rangle = \int dE \, g(E) \, e^{-iEt/\hbar} \sum_{i} c_i \, e^{i\mathbf{p}_i \cdot \mathbf{x}/\hbar} \, |\nu_i\rangle \,. \tag{1.4}$$

The function g(E) is arbitrary, and each energy eigenstate has three terms: one for each mass eigenstate $|\nu_i\rangle$. The source of neutrinos is the origin of the coordinate system, where we suppose they are created in weak eigenstate $|\nu_{\alpha}\rangle$. This boundary condition gives the requirement $c_i \equiv U_{\alpha i}$.

At a neutrino detector a distance L from the source, the imaginary phase of the component $|\nu_i\rangle$ with energy E is given by $\phi_i(L\hat{\mathbf{x}}, t; E) = (p_iL - Et)/\hbar$. The energy eigenstates (except for those with very small values of E, which are present in negligible quantities) are all relativistic. Hence $p_i \equiv \sqrt{E^2/c^2 - m_i^2c^2}$ is closely approximated by the first-order Taylor expansion $p_i \approx E/c - m_i^2 c^4/2cE$. The overall phase of $|\nu_i\rangle$ at a given energy becomes

$$\phi_i(L\hat{\mathbf{x}}, t; E) = \frac{E}{\hbar} \left(\frac{L}{c} - t\right) - \frac{Lm_i^2 c^3}{2E\hbar} + O\left[\left(\frac{m_i c^2}{E}\right)^4\right].$$
(1.5)

Notice that the first term is the same for all three mass eigenstate components. The expansion of the observed neutrinos in terms of the flavor eigenstates is then

$$|\nu(L\hat{\mathbf{x}},t)\rangle \approx \sum_{i,\beta} U_{\alpha i} U_{i\beta}^{\dagger} \int \mathrm{d}E \, g(E) \, e^{-i(L/c-t)E/\hbar} \, e^{-iLm_i^2 c^3/2E\hbar} \, |\nu_\beta\rangle \,. \tag{1.6}$$

For a monoenergetic neutrino beam, therefore, the probability of observing a flavor $|\nu_{\beta}\rangle$ given that the beam originated as $|\nu_{\alpha}\rangle$ is

$$\mathbf{P}(\nu_{\alpha} \to \nu_{\beta}|E) \approx \left| \sum_{i} U_{\alpha i} U_{i\beta}^{\dagger} e^{-iLm_{i}^{2}c^{3}/2E\hbar} \right|^{2}.$$
(1.7)

The fact that a neutrino may be observed in a flavor different from its original state is commonly called neutrino oscillation. This terminology is a bit misleading since a neutrino is hardly ever 100% in a specific flavor; it is actually the probabilities of observing it to be in a specific flavor eigenstate that oscillate. It should also be noted that expanding the initial neutrino source in terms of eigenstates of momentum rather than energy will lead to the same result. This is because the two formulations are completely equivalent, as detailed in, for instance, reference [3].

Electron neutrino oscillation

Plugging the values of U_{e1}, U_{e2}, U_{e3} into the equation above and making the additional assumption $\theta_{13} \approx 0$ results in finding the survival probability of an electron neutrino of energy E to be

$$P(\nu_e \to \nu_e | E) \approx 1 - \sin^2 2\theta_{12} \sin^2 \left(\frac{(m_2^2 - m_1^2)Lc^3}{4E\hbar}\right).$$
 (1.8)

The argument of the second sine function is often seen in the literature in the convenient form $1.27\Delta m_{12}^2 L/E$, where $\Delta m_{ij}^2 \equiv m_j^2 - m_i^2$ is in eV^2/c^4 , L is in meters, and E is in MeV. Put another way, the characteristic oscillation wavelength of an electron neutrino of energy E is

$$\lambda(E) = \frac{4\pi E\hbar}{\Delta m_{12}^2 c^3} \approx \frac{2.47 E/\text{MeV}}{\Delta m_{12}^2 c^4/\text{eV}^2} \text{ meters.}$$
(1.9)

The value $\sin^2 2\theta_{12}$ has been determined, through numerous experiments (Figure 1.1), to be 0.86 ± 0.04 . (More usually it is cited in terms of $\tan^2 \theta_{12}$, given in reference [4] as 0.45 ± 0.05 .) These parameters for historical reasons are named the Large Mixing Angle, or LMA, solution. Other proposed values for these parameters, which were ruled out by recent observational data from the SNO and KamLAND experiments, were the Small Mixing Angle (SMA) and vacuum oscillation (VAC) solutions. With the LMA parameters, the maximum probability to observe an electron neutrino is one, and the minimum probability is $\cos^2 2\theta_{12} \approx 0.14$. This is quite a variation!



Figure 1.1: Global fit of the parameters $\tan^2 \theta_{12}$ and Δm_{12}^2 for electron neutrino oscillation. Both solar neutrino experiments and experiments with reactor antineutrinos have contributed to the current small region in parameter space. Figure taken from reference [4].

Suppose a neutrino source is not pointlike, but instead has a radius R near or greater than the oscillation wavelength λ for a particular neutrino energy E. A detector situated at a slightly variable distance $L \pm \frac{1}{2} \delta L$ from the source, where $L \gg \delta L \gg R$, will not be able to discern individual peaks and troughs in the electron neutrino survival probability. The peaks and troughs generated by each point within the source cancel out, leaving only the average probability for electron neutrino survival, $\langle P(\nu_e \rightarrow \nu_e) \rangle = 1 - \frac{1}{2} \sin^2 2\theta_{12} \approx 0.57$. Thus the neutrino energies for which peaks and troughs due to vacuum oscillations may be observed fall in the range given by (approximately) $R < \lambda(E) < \delta L$:

$$\frac{0.4\Delta m_{12}^2 R}{\text{eV}^2 \text{ m}} \text{ MeV} < E < \frac{0.4\Delta m_{12}^2 \delta L}{\text{eV}^2 \text{ m}} \text{ MeV}.$$
(1.10)

To take a specific example, the core of the Sun where neutrinos are produced has a radius of $R \approx 0.2 R_{\odot} = 1.4 \times 10^5 \text{ km}$ [5]; refer to Figure 1.5. The difference between the Earth's annual minimum and maximum distances from the Sun is $\delta L \approx 5 \times 10^6 \text{ km}$. Given the value for

 $\Delta m_{12}^2 = (8.0 \pm 0.3) \times 10^{-5} \,\mathrm{eV}^2/c^4$ [4], the energy domain where vacuum oscillations of solar neutrinos might be directly observed on Earth is between 4.5 and 160 GeV. No neutrinos of such high energies are produced in the Sun; vacuum oscillation wavelengths for typical solar neutrinos are 27 km for an 862-keV electron neutrino, and 310 km for a 10-MeV ν_e . Hence we expect to detect solar electron neutrinos with constant probability ~ 57%. This expectation will be altered by mass interaction effects; see Section 1.2.2.

This analysis has not yet answered the question: if a neutrino that originally had ν_e flavor is not detected as a ν_e , what is it detected as? For solar neutrinos, the question is moot. The species ν_{μ} and ν_{τ} interact with "ordinary" matter with the same cross sections. Solar neutrinos carry far less energy than necessary to interact with electrons to produce muons $(m_{\mu}c^2 = 106 \text{ MeV})$ or tau particles $(m_{\tau}c^2 = 1.8 \text{ GeV})$. For all practical purposes, ν_{μ} and ν_{τ} in this energy domain are experimentally indistinguishable.

Muon neutrino oscillation

Consider now the case of μ -neutrinos. They will oscillate to electron neutrinos with the probability (using the same units as previously)

$$P(\nu_{\mu} \to \nu_{e}|E) \approx \cos^{2}\theta_{23} \sin^{2} 2\theta_{12} \sin^{2} \left(\frac{1.27\Delta m_{12}^{2}L}{E}\right), \qquad (1.11)$$

again under the assumption $\theta_{13} \approx 0$. In this limit, $P(\nu_{\mu} \rightarrow \nu_{e}|E) = P(\nu_{e} \rightarrow \nu_{\mu}|E)$ as well.

 μ -neutrinos are produced continually in the Earth's upper atmosphere as cosmic rays interact and decay in it. Typical energies for the resulting neutrinos range from 0.1–1000 GeV. The rough height at which the majority originate is $h \approx 15$ km, so a detector capable of observing μ -neutrinos will see them travel distances between 15 km and $2R_{\oplus} = 1.3 \times 10^4$ km, depending on whether they are created overhead or on the other side of the planet. Suppose we select only those neutrinos with E > 10 GeV. In this case, values of the argument of the second sine function in Equation (1.11) are no greater than 0.13. Regardless of the value of θ_{23} , at most 1% of these high-energy μ -neutrinos will be detected as electron neutrinos,

Parameter	Value	Units	Useful oscillations	(Anti-)neutrino sources
Δm_{12}^2	$(8.0 \pm 0.3) \times 10^{-5}$	eV^2/c^4	$ u_e \rightarrow u_{\mu, au}$	Sun
$\sin^2 2\theta_{12}$	0.86 ± 0.04	-	$\bar{ u}_e ightarrow \bar{ u}_{\mu, au}$	Reactors, geo- $\bar{\nu}$
$ \Delta m^2_{23} $	$(2.5 \pm 0.3) \times 10^{-3}$	eV^2/c^4	$ u_{\mu} ightarrow u_{ au}$	μ^{\pm} reacting with
$\sin^2 2\theta_{23}$	1.02 ± 0.04	-	$ar u_\mu o ar u_ au$	upper atmosphere
$\sin^2 2\theta_{13}$	< 0.12 (99% CL)	-	$\bar{\nu}_e ightarrow \bar{\nu}_e$	Reactors

Table 1.3: Known parameters of neutrino oscillations. Δm_{13}^2 is not listed; it is determined completely by the requirement $\Delta m_{13}^2 = \Delta m_{12}^2 + \Delta m_{23}^2$. The main questions are: what is the actual value of θ_{13} ? What is the sign of Δm_{23}^2 ? What are the actual values of m_1, m_2, m_3 ? Listed values taken from reference [4].

and the converse is also true. For high-energy neutrinos over baselines of planetary scale or less, little $\nu_e \leftrightarrow \nu_{\mu}$ oscillation will be seen. (Indeed, very little oscillation of high-energy atmospheric ν_e 's should be observed at all.)

On the other hand, consider the probability of μ -neutrino survival. If we now approximate $\Delta m_{12}^2 L/E \approx 0$ as well, this is given by

$$P(\nu_{\mu} \to \nu_{\mu} | L/E < 1.5 \,\mathrm{km/GeV}) \approx 1 - \sin^2 2\theta_{23} \,\sin^2\left(\frac{1.27\Delta m_{23}^2 L}{E}\right).$$
 (1.12)

This is precisely the form of Equation (1.8), with "12" replaced uniformly by "23." Since few high-energy μ -neutrinos are detected as electron neutrinos, the balance must become τ -neutrinos:

$$P(\nu_{\mu} \to \nu_{\tau} | L/E < 1.5 \,\mathrm{km/GeV}) \approx \sin^2 2\theta_{23} \,\sin^2 \left(\frac{1.27\Delta m_{23}^2 L}{E}\right). \tag{1.13}$$

Indeed, in most regimes neutrino oscillation can be approximated as occurring in a quasitwo-neutrino system. For solar neutrinos, that system consists of $(\nu_e, \sim (\nu_\mu + \nu_\tau)/\sqrt{2})$; for high-energy atmospheric neutrinos, (ν_μ, ν_τ) .

In principle, it is possible to assume that Δm_{23}^2 is similar to or smaller than Δm_{12}^2 . In this case, the μ -neutrino survival probability at high energies would always be roughly one. However, this is not what is observed. The large phase space in L and E available to detectors of atmospheric neutrinos has made it possible to determine the values of Δm_{23}^2 and θ_{23} experimentally. As a simple example, if we assume $R_{\oplus} \gg \lambda \gg h$, upgoing μ -neutrinos (coming from the other side of the Earth) survive with probability $1 - \frac{1}{2} \sin^2 2\theta_{23}$, while downgoing μ -neutrinos have undergone essentially no flavor conversion; that is, $\sin^2 2\theta_{23} \approx$ $2(1 - N_{\uparrow}/N_{\downarrow})$. The actual best-fit results for the parameters are given in Table 1.3. With $\theta_{23} \approx 45^{\circ}$, this incidentally answers the question above about solar electron neutrinos: the ratio that oscillate into μ -neutrinos versus τ -neutrinos is essentially 1:1.

1.2.2 The Mikheyev-Smirnov-Wolfenstein effect in matter

The above discussion of oscillations is valid only for neutrinos passing through a vacuum, or at least a region with a relatively low number density of electrons. As they travel, neutrinos will interact to some extent to the surrounding matter. Virtually all normal matter exists in the first family of fundamental particles, so electron neutrinos are more likely to interact with it than μ - and τ -neutrinos. This fact causes two effects: first, in very dense material, neutrinos are slowed down by interactions; second, and more importantly, ν_e 's are slowed down more than $\nu_{\mu,\tau}$'s. The difference in velocities causes different components to become out of phase, and as we saw above, that inevitably leads to flavor oscillation. This type of oscillation is termed the Mikheyev-Smirnov-Wolfenstein (MSW) effect after its discoverers [6, 7].

When an electron neutrino is created in the core of the Sun, suppose that it has definite momentum (not energy) p. To first order, its vacuum Hamiltonian H_0 in the mass eigenstate basis is a diagonal matrix, with eigenvalues $\sqrt{p^2c^2 + m_i^2c^4} \approx cp + m_i^2c^3/2p$. However, there are now two additional terms to the Hamiltonian; one to describe neutral-current interactions with matter, and one to describe charged-current interactions.

In the *flavor* basis, where the charged-current term discriminates for electron neutrinos and the neutral-current term acts equitably, the Hamiltonian is given by

$$H = U_{\rm MNS} H_0 U_{\rm MNS}^{\dagger} + k_{\rm NC} I + k_{\rm CC} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$
(1.14)

(note the basis transformation of H_0 with $U_{\rm MNS}$), where $k_{\rm NC}$ and $k_{\rm CC}$ are the respective interaction energies for neutral-current and charged-current interactions for neutrinos traveling through matter. The three flavors of neutrino interact with matter via neutral currents equally—I above is the 3 × 3 identity matrix—but only the electron neutrino interacts with normal matter through charged currents. The interaction coefficient in the charged-current case is actually $G_F n_e \sqrt{2}$, where n_e is the number density of electrons in a material and G_F is the Fermi constant, but we will continue to label it simply $k_{\rm CC}$ for notational convenience.

If we now transform back to the basis of the mass eigenstates in vacuum, then under the assumption of vanishing θ_{13} , we have

$$H = \left(k_{\rm NC} + cp + \frac{m_1^2 c^3}{2p}\right) I + \left(\begin{array}{ccc} k_{\rm CC} c_{12}^2 & k_{\rm CC} c_{12} s_{12} & 0\\ k_{\rm CC} c_{12} s_{12} & k_{\rm CC} s_{12}^2 + \frac{\Delta m_{12}^2 c^3}{2p} & 0\\ 0 & 0 & \frac{\Delta m_{13}^2 c^3}{2p} \end{array}\right).$$
(1.15)

Note that $|\nu_3\rangle$ is still an eigenstate of the Hamiltonian in matter! Due to the smallness of θ_{13} , there will still be essentially no mixing of electron neutrinos into that state. The situation is a two-neutrino problem.

Let us therefore focus only on the upper-left 2×2 block of the matrix. We subtract $k_{\rm CC} \sin^2 \theta_{12}$ from the matrix main diagonal and compensate by adding it back to the term multiplying the identity matrix. Then (by using the trigonometric double-angle identities) this block may be further simplified as

$$\left(k_{\rm NC} + k_{\rm CC}\sin^2\theta_{12} + cp + \frac{m_1^2c^3}{2p}\right) I + \left(\begin{array}{cc}k_{\rm CC}\cos2\theta_{12} & \frac{1}{2}k_{\rm CC}\sin2\theta_{12}\\\frac{1}{2}k_{\rm CC}\sin2\theta_{12} & \frac{\Delta m_{12}^2c^3}{2p}\end{array}\right).$$
(1.16)

The eigenvalues of this matrix, ignoring the part proportional to I and dropping the "CC" suffix, are

$$\varepsilon_{\pm} = \frac{\Delta m_{12}^2 c^3}{4p} + \frac{k}{2} \cos 2\theta_{12} \pm \frac{1}{2} \sqrt{\left(\frac{\Delta m_{12}^2 c^3}{2p} \sin 2\theta_{12}\right)^2 + \left(\frac{\Delta m_{12}^2 c^3}{2p} \cos 2\theta_{12} - k\right)^2}, \quad (1.17)$$

and the difference between them is

$$\Delta \varepsilon = \sqrt{\left(\frac{\Delta m_{12}^2 c^3}{2p} \sin 2\theta_{12}\right)^2 + \left(\frac{\Delta m_{12}^2 c^3}{2p} \cos 2\theta_{12} - k\right)^2}.$$
 (1.18)

Compare this with the difference between energy eigenstates in a vacuum, which is $(\Delta \varepsilon)_0 = \Delta m_{12}^2 c^3/2p$. In the limit $kp \ll \Delta m_{12}^2 c^3$, it is clear that $\Delta \varepsilon \to (\Delta \varepsilon)_0$.

Suppose that the normalized eigenvectors of Equation (1.16) are $\binom{a_1}{b_1}$ and $\binom{a_2}{b_2}$. We may use them to construct the unitary matrix $U_m \equiv \begin{pmatrix} a_1 & a_2 \\ b_1 & b_2 \end{pmatrix}$. Since U_m is real and unitary, it must satisfy $a_1 = b_2 = \cos \theta'$, $a_2 = -b_1 = \sin \theta'$ for some real angle θ' . Hence $\tan \theta'$ equals $a_2/b_2 = -b_1/a_1$. We now have

$$U_m \begin{pmatrix} |\nu_1^m\rangle \\ |\nu_2^m\rangle \end{pmatrix} = \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix}, \tag{1.19}$$

where $|\nu_{1,2}^m\rangle$ are the neutrino energy eigenstates in matter—not the same as the mass eigenstates in vacuum. (However, $|\nu_3\rangle = |\nu_3^m\rangle$.) This gives us, still in the approximation $\theta_{13} \approx 0$,

$$\begin{pmatrix} |\nu_e\rangle\\ |\nu_{\mu}\rangle\\ |\nu_{\tau}\rangle \end{pmatrix} = U_{\rm MNS} \begin{pmatrix} |\nu_1\rangle\\ |\nu_2\rangle\\ |\nu_3\rangle \end{pmatrix} = U_{\rm MNS} \begin{pmatrix} U_m & 0\\ 0 & 1 \end{pmatrix} \begin{pmatrix} |\nu_1^m\rangle\\ |\nu_2^m\rangle\\ |\nu_3\rangle \end{pmatrix}.$$
(1.20)

We therefore define $U_{\text{MNS}}^m \equiv U_{\text{MNS}} \operatorname{diag}(U_m, 1)$ to be the neutrino mixing matrix in matter. For electron neutrinos, it has the entries

$$U_{e1}^{m} = c_{12}\cos\theta' - s_{12}\sin\theta' = \cos(\theta_{12} + \theta')$$
$$U_{e2}^{m} = s_{12}\sin\theta' + c_{12}\cos\theta' = \sin(\theta_{12} + \theta')$$
$$U_{e3}^{m} \approx 0.$$



Figure 1.2: Dependence of the mixing parameter in matter, $R \equiv \sin^2 2\theta_m$, on the product 2kE in units of $\Delta m_{12}^2 c^4$. R is shown on the vertical scale, and the ratio (labeled x) on the horizontal. Values of x < 0 refer to antineutrinos. Two curves are shown: the red one (peaking farther to the left) uses $\sin^2 2\theta = 0.825$ (near the currently accepted value for θ_{12}), and the green (narrower) curve has $\sin^2 2\theta = 0.08$ (about the upper limit on the θ_{13} parameter). The mixing in vacuum is given by the point where each curve intersects the y-axis. Figure taken from reference [8].

Through exactly the same argument leading to Equation (1.8), we obtain the equivalent of that equation, making the replacements

$$\Delta m_{12}^2 c^4 / 2E \quad \to \quad \Delta \varepsilon$$
$$\theta_{12} \quad \to \quad \theta_m \equiv \theta' + \theta_{12}$$

(we have also approximated p by E/c). That is,

$$P(\nu_e \to \nu_e | E) = 1 - \sin^2 2\theta_m \, \sin^2 \left(\frac{L\Delta\varepsilon}{2\hbar c}\right). \tag{1.21}$$

By explicitly finding the eigenvectors of U_m , we can obtain values for $\tan \theta'$ as defined above; these may be substituted into the trigonometric identity

$$\sin 2\theta_m = \sin(2\theta_{12} + 2\theta') \equiv \frac{(1 - \tan^2 \theta') \sin 2\theta_{12} + 2\tan \theta' \cos 2\theta_{12}}{1 + \tan^2 \theta'}$$
(1.22)

to obtain, finally, the result

$$\sin^2 2\theta_m = \frac{\sin^2 2\theta_{12}}{\sin^2 2\theta_{12} + \left(\cos 2\theta_{12} - \frac{2kE}{\Delta m_{12}^2 c^4}\right)^2}.$$
(1.23)

Therefore, when $k = k_r \equiv \cos 2\theta_{12} \Delta m_{12}^2 c^4/2E$, Equation (1.21) reduces to the simple formula $P = 1 - \sin^2(L\Delta\varepsilon/2\hbar c)$. In this condition we have resonance (Figure 1.2): it is possible for electron neutrinos in matter to be *completely converted* to a different flavor. At the resonant point, $\Delta\varepsilon = (\Delta m_{12}^2 c^4/2E) \sin 2\theta_{12}$. This allows us to calculate the resonant wavelength,

$$\lambda_r(E) = \lambda(E) \csc 2\theta_{12} \approx 1.08\lambda(E), \tag{1.24}$$

with $\lambda(E)$, the vacuum oscillation length, given by Equation (1.9). This value $\lambda_r(E)$ is the greatest possible matter oscillation wavelength for a given neutrino energy E. When kis much less than its value at resonance, $\sin 2\theta_m \approx \sin 2\theta_{12}$, and the oscillations approach their behavior in vacuum. When $k \gg k_r$, $\sin 2\theta_m \approx (k_r/k) \tan 2\theta_{12}$: oscillations are strongly suppressed, as the electron neutrino is nearly an eigenstate of the matter Hamiltonian. In particular, it is approximated by the higher-energy state $|\nu_2^m\rangle$, since with k > 0, the chargedcurrent interaction term raises the energy of the electron neutrino flavor. In this limit, the oscillation wavelength goes to $\lambda \approx 2\pi\hbar c/k$ and may become very short.

An electron neutrino created in a region where $k \ll k_r$ will always behave nearly as if it is in vacuum, and will oscillate accordingly; matter effects may be neglected. On the other hand, if it is created where $k > k_r$, the situation is more complicated. A proper discussion may be found in, for instance, reference [9]. Briefly, at the point of creation, $|\nu_e\rangle \approx |\nu_2^m\rangle$. As the density (and therefore the value of k) changes, the matter eigenstates evolve slightly as a function of k. However, the neutrino will remain mostly in the $|\nu_2^m\rangle$ eigenstate as long as the change is gradual (adiabatic); the energies of the two matter eigenstates never cross, but are continuous functions through the point at which k = 0, where they transition smoothly into the corresponding vacuum mass states. The neutrino thus remains in the state $|\nu_2\rangle$ until reaching a detector, at which point it is seen as an electron neutrino with probability $|U_{e2}|^2 = \sin^2 \theta_{12} \approx 31\%$. As the neutrino energy E increases, $k_r(E)$ becomes smaller. We now ask for the minimum neutrino energy E_r required such that the maximum value of k attained in the Sun is $\geq k_r(E)$. For neutrinos with lower energies than E_r , k in the Sun is never sufficiently high, and vacuum oscillations dominate. On the other hand, neutrinos with much higher energies behave as described in the previous paragraph. In the limit of low E, the survival probability observed at Earth will be $1 - \frac{1}{2}\sin^2 2\theta_{12} \approx 0.57$, as discussed in the section above. In the limit of high E (but not so high that $\lambda(E)$ is on the same order as R_{\odot} !) it will be simply $\sin^2 \theta_{12} \approx 0.31$. (More detailed calculations with three neutrinos reveal that the measured values 0.57 and 0.31 are actually the products of the explicit functions of θ_{12} with $\cos^4 \theta_{13}$, but we will continue to ignore this factor.)

The core density of the Sun is about 150 g/cm^3 , and the helium mass fraction at the core is about 63% [5]. The resonant electron density is

$$n_e = \frac{\cos 2\theta_{12} \Delta m_{12}^2 c^4}{2\sqrt{2}G_F E}.$$
(1.25)

In terms of matter density, $n_e = \langle Z/A \rangle \rho/m_p$, with m_p being the proton mass and $\langle Z/A \rangle$ the average ratio of atomic number to atomic mass in the Sun's core, roughly 0.7. Therefore the critical energy is

$$E_r = \frac{\cos 2\theta_{12} \Delta m_{12}^2 c^4 m_p}{2\sqrt{2}G_F \langle Z/A \rangle \rho} \approx 1.9 \,\mathrm{MeV}. \tag{1.26}$$

The behavior change that occurs near this energy is known as the vacuum-matter transition (Figure 1.3). The transition will have a slightly different behavior for neutrinos produced in each reaction described in Section 1.3, due to their differing probabilities of being created at different positions in the Sun's core (Figure 1.5). The observation of solar neutrinos at varying energies to measure this transition has not yet occurred. All currently running solar neutrino experiments either are only able to observe those with energies above a few MeV, or else are radiochemical experiments (such as Homestake and GALLEX) that do not measure neutrino energies, but only count the number detected above a particular threshold. Observations of a substantially different shape of curve would be challenging to existing theories of neutrino oscillation.



Figure 1.3: The vacuum-matter transition for solar neutrinos: averaged survival probability for electron neutrinos as a function of neutrino energy. The expected value of the transition energy is ~ 2 MeV. The higher-energy data point comes from SNO data, and the lower one from various gallium-based radiochemical experiments. This plot is only an approximation; in reality, the curve is different for neutrinos created in each type of reaction, since they are created with different radial distributions in the Sun's core (see Figure 1.5). Figure taken from reference [4].

In passing, we mention that the observed high-energy atmospheric neutrinos, which oscillate only between ν_{μ} and ν_{τ} flavors, must consist almost entirely of the ν_3 mass eigenstate. This eigenstate remains essentially unchanged between the matter and vacuum mass bases. As a result, this sample of neutrinos should be very little affected by passage through matter of any density.

It is also worth noting that the product $\rho \langle Z/A \rangle$ at the center of the Earth is less than that in the Sun's core by about an order of magnitude. Hence we do not expect to observe a "day-night asymmetry," caused by matter oscillations as solar neutrinos travel through the Earth to reach us at night, for any solar neutrinos except possibly those of greatest energy.
The ν_e spectrum end point of the Sun (excluding the rare *hep* neutrinos) is 14 MeV, rather less than what should be the transition energy for neutrinos passing through the Earth's core. Therefore the detection of a solar neutrino day-night effect would also be potential evidence for new physics.

1.3 Solar neutrinos

By far the largest source of neutrinos that can be detected on Earth with $E \ll 1 \text{ GeV}$ is our Sun. The nuclear fusion reactions that power its core convert four hydrogen nuclei (protons) to one ⁴He nucleus; in the process, neutrinos are produced. Two main cycles implement this transformation. The *pp* cycle, believed to be dominant, consists largely of simply putting protons together until a helium nucleus is formed. The CNO cycle, on the other hand, uses nuclei of heavier elements (carbon, nitrogen, and oxygen—hence the name) as catalysts.

The Sun is a complex object. The rates of the individual fusion processes at its core are determined by many variables, including (as functions of radius) temperature, pressure, density, opacity, convection, and concentrations of different elements. These may be simulated by a complex set of differential equations describing energy transport and hydrostatic equilibrium. Essentially the only way to solve the system is numerically. Computer models of the Sun are divided into radial shells, assuming constant values of the variables within each shell and patching the solutions together at the interfaces. The models are required to be approximately static since the Sun is observed to be stable. Outward pressure must balance inward gravitational forces; the observed amount of radiation must be balanced by the total energy produced in nuclear reactions.

The currently recognized best model of this sort is the Standard Solar Model (SSM). The latest version of the SSM was until recently the one denoted BS05(OP) [10]. However, a more recent paper describes the unorthodox method of simulating 10^4 slightly different solar models in which the input parameters were varied according to Monte Carlo estimates of

statistical error distributions. The average output over all models is termed the BSB(GS98) model [5]. (The "GS98" refers to the assumed solar elemental abundances given in 1998 by Grevesse and Sauval [11]. More recent abundances were taken into account in [5] in the BSB(AGS05) model, but doing so yielded results that were incompatible with helioseismological data to > 3σ . Until this discrepancy can be explained, the BSB(GS98) model is to be preferred.)

1.3.1 The pp cycle

In stars with roughly the mass of the Sun or less, the pp cycle is dominant. Its overall input and output may be summarized as follows:

$$4p \rightarrow {}^{4}\text{He} + 2e^{+} + 2\nu_{e} + 26.7 \,\text{MeV}.$$
 (1.27)

Already we are provided with information, the "luminosity constraint," about the total neutrino flux that should reach us from the Sun. That is, if the energy released from the Sun in neutrinos is small relative to that released in photons, we have

$$\Phi_{\nu} \approx \frac{2\nu_e \text{'s}}{26.7 \,\text{MeV}} \frac{L_{\odot}}{4\pi a^2},\tag{1.28}$$

 L_{\odot} being the solar luminosity (energy produced per unit time) and *a* being the average Earth-Sun distance. This calculation yields $\Phi_{\nu} \approx 6.6 \times 10^{10} \,\mathrm{cm}^{-2} \mathrm{s}^{-1}$.

To begin the pp cycle, two protons interact, combining to form a deuteron:

$$p + p \rightarrow {}^{2}\mathrm{H} + e^{+} + \nu_{e} \tag{1.29}$$

As one of the protons must be converted to a neutron in the process, the reaction is governed by the weak interaction. It is therefore one of the slowest steps in the chain. It also produces a neutrino, the so-called "*pp* neutrino." The *pp* ν energy is a continuous spectrum with an end point at 420 keV.



Figure 1.4: Predicted spectra of solar neutrinos at Earth, based on the BS05(OP) solar model; figure taken from reference [10].

Source	E_{ν}	$E_{\rm max}$	Flux	Error
	[MeV]	[MeV]	$[\rm cm^{-2} s^{-1}]$	[%]
pp	≤ 0.42	0.26	5.99×10^{10}	0.9
pep	1.44	1.22	1.42×10^8	1.5
hep	≤ 18.8	18.5	7.93×10^3	15.5
$^{7}\mathrm{Be}$	0.862	0.667	4.34×10^9	10.5
"	0.384	0.231	$5.03 imes 10^8$	10.5
$^{8}\mathrm{B}$	≤ 14.06	13.81	$5.69 imes 10^6$	$^{+17.3}_{-14.7}$
^{13}N	≤ 1.20	0.99	3.05×10^8	$^{+36.6}_{-26.8}$
$^{15}\mathrm{O}$	≤ 1.73	1.51	2.31×10^8	$+\overline{37.4}$ -27.2
$^{17}\mathrm{F}$	≤ 1.74	1.52	$5.83 imes 10^6$	$+ \overline{7} 2. \overline{4} \\ -42.0$

Table 1.4: Tabulated predictions from the BSB(GS98) solar model [5] for solar neutrino fluxes at Earth. The E_{max} column gives the maximum electron recoil energy that can be observed in νe scattering (see Section 2.3).



Figure 1.5: Radial distributions of the probabilities that a neutrino from a particular reaction of the pp cycle will be created at a given distance from the center of the Sun. For ease of comparison, the distributions are displayed normalized such that the integral over each yields one. The dotted line represents the radial distribution for production of solar luminosity. (All curves go to zero at the center of the Sun because of the geometric effect in spherical coordinates. To get the neutrino creation probabilities per unit volume, which are maximal at R = 0, one would need to divide out a factor of $4\pi R^2$.) Figure taken from reference [5].

This reaction may be slightly modified by having an electron as input rather than a positron as output. Since there are three reactant particles, the occurrence is unlikely. The threebody *pep* reaction

$$p + p + e^- \rightarrow {}^{2}\mathrm{H} + \nu_e \tag{1.30}$$

is disfavored compared to the pp reaction by a factor of about 400. In this reaction there are only two products, so the energy of the resulting pep neutrino is completely determined $(E_{\nu} = 1.44 \text{ MeV}).$

However it was formed, the deuteron quickly captures another proton, forming ³He:

$${}^{2}\mathrm{H} + p \rightarrow {}^{3}\mathrm{He} + \gamma \tag{1.31}$$

The reaction is fast, as no flavor changes are involved. As a result, the concentration of deuterium in the Sun is never high; this prevents reactions such as ${}^{2}H + {}^{2}H \rightarrow {}^{4}He$ from being significant.

Since the most common species in the core of the Sun are p, ³He, and ⁴He, the next reaction to occur will most probably be one of the following:

$${}^{3}\text{He} + p \rightarrow {}^{4}\text{He} + e^{+} + \nu_{e}$$
 (1.32)

$${}^{3}\text{He} + {}^{3}\text{He} \rightarrow {}^{4}\text{He} + 2p$$
 (1.33)

$${}^{3}\text{He} + {}^{4}\text{He} \rightarrow {}^{7}\text{Be} + \gamma$$

$$(1.34)$$

The first of these is weak, while the other two that compete with it are not. It is therefore strongly disfavored. Still, it produces a particle, the rare *hep* neutrino, which could conceivably be observed by Earth-based detectors. The *hep* neutrino spectrum end point is at 18.8 MeV, higher than any other solar neutrino.

In the second case, the cycle is finished. Of six initial protons, four were converted to a helium nucleus and two are left over to participate in other reactions. This is the primary branch of the pp cycle. However, the resulting nucleus may still interact with ³He to produce ⁷Be. When it does, there are two possible fates for the new ⁷Be nucleus. It may capture an electron, followed by proton capture on ⁷Li generating two helium nuclei:

$$^{7}\text{Be} + e^{-} \rightarrow ^{7}\text{Li} + \nu_{e}$$
 (1.35)

$${}^{7}\mathrm{Li} + p \quad \rightarrow \quad {}^{4}\mathrm{He} + {}^{4}\mathrm{He} \tag{1.36}$$

In this case the monoenergetic $(E_{\nu} = 862 \text{ keV})$ ⁷Be neutrino is produced. In fact there are actually two possible energies, since the ⁷Li daughter may be produced in an excited state with Q = 478 keV. The branching ratio to the excited state, resulting in a lower-energy 384-keV neutrino, is only 10.4%, however.

Or, it may capture a proton, after which ⁸B and ⁸Be consecutively decay:

$${}^{7}\text{Be} + p \quad \rightarrow \quad {}^{8}\text{B} + \gamma \tag{1.37}$$

$${}^{8}\mathrm{B} \rightarrow {}^{8}\mathrm{Be} + e^{+} + \nu_{e} \tag{1.38}$$

$${}^{8}\text{Be} \rightarrow {}^{4}\text{He} + {}^{4}\text{He}$$
(1.39)

These ⁸B neutrinos are the primary variety that have been detected by most solar neutrino experiments to date. Though they are produced in relative numbers much less than most other solar neutrino types, they have a high energy spectrum end point (14.1 MeV). Most radioactive background in a neutrino detector (the uranium and thorium decay chains and a few other natural and anthropogenic isotopes) has an end point more on the order of 3-5 MeV. The ⁸B ν flux is also much greater than that of *hep* neutrinos.

It is unfortunate that, to date, no direct measurement of the vast majority of solar neutrinos—pp neutrinos—is possible. Still, as detectors with lower radioactive backgrounds and energy thresholds come into operation, we are nearing that goal. The Borexino experiment, described in Chapters 2–3, will be among the first detectors capable of making real-time observations of the ⁷Be neutrino, which makes up about ~ 10% of the total solar neutrino flux. It is even conceivable that under some circumstances, the Borexino experiment could measure the tail end of the pp neutrino signal [12].

1.3.2 The CNO cycle

The temperature and pressure at the core of the Sun are not sufficiently high to fuse helium nuclei into carbon. The Coulomb energy barrier is too great. Nevertheless, the Sun contains a non-negligible fraction of carbon and other "heavy" elements as a result of being a Population I (second-generation) star. The solar system was formed from gas and dust clouds enriched in heavy elements that were distributed by long-ago supernovae. Protons may be fused together into helium nuclei using these heavier elements as catalysts. This process is termed the CNO cycle. In stars slightly larger than the Sun, it is the dominant mode of fusion. The rate of CNO-cycle processes is strongly temperature-dependent, so the exact flux of neutrinos produced by it is quite uncertain, on the order of 30%.

In brief, the CNO cycle consists of the following six reactions. We begin with a nucleus of ¹⁴N, but that is of course an arbitrary starting point, as it is regenerated in the cycle.

$$^{14}N + p \rightarrow ^{15}O + \gamma$$
 (1.40)

$${}^{15}\text{O} \rightarrow {}^{15}\text{N} + e^+ + \nu_e \tag{1.41}$$

$$^{15}\mathrm{N} + p \rightarrow ^{12}\mathrm{C} + {}^{4}\mathrm{He}$$
 (1.42)

$$^{12}C + p \rightarrow ^{13}N + \gamma$$
 (1.43)

$$^{13}N \rightarrow ^{13}C + e^+ + \nu_e \tag{1.44}$$

$${}^{13}C + p \quad \rightarrow \quad {}^{14}N + \gamma \tag{1.45}$$

In the course of one cycle, four protons are converted to a helium nucleus, and two neutrinos are produced. These neutrinos are called the ¹³N and ¹⁵O neutrinos since they result from the β^+ decay of those isotopes. They have respective end points of 1.20 and 1.73 MeV.

Reactions (1.42–1.45) may be replaced by an alternative set of reactions to finish the cycle:

$$^{15}\mathrm{N} + p \rightarrow ^{16}\mathrm{O} + \gamma$$
 (1.46)

$$^{16}\mathrm{O} + p \rightarrow ^{17}\mathrm{F} + \gamma$$
 (1.47)

$${}^{17}\mathrm{F} \quad \rightarrow \quad {}^{17}\mathrm{O} + e^+ + \nu_e \tag{1.48}$$

$${}^{17}\mathrm{O} + p \quad \rightarrow \quad {}^{14}\mathrm{N} + {}^{4}\mathrm{He} \tag{1.49}$$

In this case, instead of the ¹³N neutrino, a ¹⁷F neutrino is emitted. This alternate cycle happens less frequently, and the ¹⁷F neutrino flux is much lower than that of the ¹³N and ¹⁵O neutrinos.

1.4 Remaining questions

The reader will have noticed that some seemingly important facts have been left out of the preceding introduction to neutrino physics. Now that we know, fairly accurately, the splittings Δm^2 between the three neutrino mass eigenstates, what are the actual values of the masses? Why is the neutrino mass so much less than that of any other fundamental particle? Why does it have mass at all? Can a neutrino oscillate into $\bar{\nu}$ as well? Do neutrinos interact by any mechanism other than the weak interaction? These are some of the questions that still must be answered at the beginning of the 21st century. Nor does our lack of understanding of these fundamental issues obviate the need to measure more precisely the semi-known neutrino parameters—particularly the question of the size of the mixing angle θ_{13} , which must be non-zero for CP-violating effects to be possible.

1.4.1 Absolute values and origin of neutrino masses

Experimental limits

One way to determine the actual value of the neutrino mass is to observe β decay spectra, looking for a distortion from the expected shape at the high-energy end point (Figure 1.6). Several experiments have done so using tritium, the lightest β -decaying isotope, where the effect would be most pronounced. Since the $|\bar{\nu}_e\rangle$ produced in the decay is a superposition of mass eigenstates, what is actually being measured is the quantity m_β , defined by

$$m_{\beta}^{2} = \sum_{i} \bar{m}_{i}^{2} |\langle \bar{\nu}_{e} | \bar{\nu}_{i} \rangle|^{2} = \sum_{i} m_{i}^{2} |U_{ei}|^{2}$$
(1.50)

(the last equality assumes, as is typically done, that the matter and antimatter masses and mixing angles are identical). Among the best such limits so far is that given by the Mainz experiment, $m_{\beta}c^2 < 2.3 \text{ eV}$ (95% CL) [14]. The KATRIN experiment, a next-generation tritium decay experiment, should be able to set limits at 200 meV (90% CL) or measure a 350 meV neutrino mass at 5 σ [13]. Much weaker analogous limits have been determined



Figure 1.6: Schematic of the β -decay spectrum of tritium near its endpoint. The spectral distortion is shown (lower, blue curve) in the hypothetical case of an electron antineutrino with $m_{\beta}c^2 = 1 \text{ eV}$. The theoretical prediction for a zero-mass neutrino is the upper, red curve. Note that the derivative of the upper curve goes to zero at the end point, but that of the lower becomes vertical. If the neutrino mass were zero, only a fraction 2×10^{-13} of all tritium decays would give the electron a kinetic energy within the shaded region. Figure taken from reference [13].

for the expectation masses of ν_{μ} (< 170 keV/ c^2 , 90% CL [15]) and ν_{τ} (< 18.2 MeV/ c^2 , 95% CL [16]). A *lower* bound for m_{β} may be estimated from solar and reactor neutrino oscillation data: we know that $m_2^2 \ge \Delta m_{12}^2$, giving $m_{\beta}c^2 \ge \sqrt{\Delta m_{12}^2} |U_{e2}|c^2 \approx \sqrt{\Delta m_{12}^2} |\sin \theta_{12}|c^2 \approx 5 \text{ meV}$. To reach this lower bound will require an experiment with a sensitivity 40 times better than that of KATRIN.

The neutrino mass may also be estimated from cosmological considerations. Statistical mechanics tells us [17] that the temperature of neutrinos left over from Big Bang nucleosynthesis (BBN) is about 70% that of the cosmic microwave background radiation, $T_{\rm CMB} =$ 2.728 K. This implies that those neutrinos now have energy $E_{\nu} = 0.7k_BT_{\rm CMB} \approx 0.2$ meV. Since the intermediate neutrino mass eigenstate is at least $\sqrt{\Delta m_{12}^2}c^2 \approx 9$ meV, for BBN neutrinos at least two of the three eigenstates are currently non-relativistic. The sum of the three eigenstate masses $\Sigma \equiv m_1 + m_2 + m_3$ is therefore related to the Hubble constant by the equation $\Omega_{\nu}h^2 = \Sigma c^2/(93.2 \text{ eV})$, with *h* the Hubble constant in units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$. (Ω_{ν} is the neutrino contribution to the curvature of space.) The Hubble constant is fairly tightly fixed at $h = 0.72 \pm 0.08$ from Hubble Space Telescope data [18].

If we assume that the universe is flat ($\Omega_{tot} = 1$), as all evidence suggests, then varying Ω_{ν} has two effects. First, it necessarily increases or decreases all other contributions to curvature. Second, if Σc^2 is greater than about 1 eV, at least the heaviest of the three neutrino mass eigenstates would have become non-relativistic even before the time that the falling average density of the universe caused neutrinos to "decouple" from (cease to interact significantly with) other matter. These two considerations mean that a significant neutrino mass would have had a measurable effect on large-scale cluster formation in the early universe, and could have been detected by experiments such as WMAP [19]. A discussion of these effects may be found in, for instance, [20].

No such cosmological signature of neutrino mass has been found. A global re-analysis of cosmological data performed in reference [21], for instance, gives an upper limit of $\Sigma c^2 < 470 \text{ meV} (2\sigma)$. Of course, a lower bound on Σ is given from atmospheric neutrino oscillation data by $\Sigma c^2 \ge \sqrt{|\Delta m_{23}^2|} c^2 \approx 50 \text{ meV}$. Additional constraints brought about by considering Big Bang nucleosynthesis, as well as the probability for the Z^0 boson to decay into "invisible" (*i. e.*, neutrino/antineutrino) channels, imply that the number of interacting neutrinos with $m < m_Z/2$ (about $45 \text{ GeV}/c^2$) is only three, the flavors we have observed.

The mass hierarchy problem

Though we know the values Δm_{12}^2 and $|\Delta m_{23}^2|$ (and that $|\Delta m_{23}^2| \gg \Delta m_{12}^2$), we do not know the sign of the larger quantity. This means that there are two possible orderings for the neutrino masses: $m_3 > m_2 > m_1$, the so-called normal or natural hierarchy, and $m_2 > m_1 > m_3$, the inverted hierarchy (Figure 1.7). It would be surprising if in fact the inverted hierarchy was the ordering found in nature. The known mass hierarchies



Figure 1.7: The two possible mass hierarchies with three neutrinos. The flavor compositions of each mass eigenstate were computed with current best values for the mixing angles. Figure taken from reference [22].

for the two sets $(+\frac{2}{3}, -\frac{1}{3})$ of quarks and for the charged leptons (Table 1.1) each have a lightest particle, a somewhat heavier intermediate particle, and a much heavier third particle. The inverted hierarchy instead has a lightest neutrino, and two much heavier neutrinos of approximately equal mass. On the other hand, we observe that, unlike the case of the nearly-diagonal CKM weak mixing matrix in the quark sector, the MNS mixing matrix has some off-diagonal terms that are at the same order of magnitude as those on the diagonal. To presume that the quark mass hierarchy applies equally in the neutrino sector might therefore be an invalid assumption.

It is also possible that the neutrino masses are nearly identical, a conclusion that would follow logically if they were found to be sufficiently large. If we suppose that the largest mass is given by about the upper limit (500 meV) of cosmological observations, the scale of absolute mass differences is then only $|\Delta m_{23}^2| \approx 2m_3 |m_3 - m_2|$, implying $|m_3 - m_2|c^2$ is only ~ 2.5 meV, or 0.5% of the actual mass. Clearly much better limits on the neutrino masses, from more accurate cosmological data or much better tritium decay shape measurements, are required to exclude this possibility.

If the inverted hierarchy prevails, then $m_{\beta}c^2 \approx m_{1,2}c^2 \geq \sqrt{|\Delta m_{23}^2|}c^2 \approx 50 \text{ meV}$. Hence a tritium decay experiment about four times more sensitive than KATRIN could, by being unable to detect a non-zero value of m_{β} , refute the inverted mass hierarchy at a 90% confidence level. However, a confirmed measurement of $m_{\beta}c^2 > 50 \text{ meV}$, while providing a highly significant and perhaps Nobel-prize-winning estimate for m_1 , would leave the hierarchy question open.

Dirac and Majorana masses

The Standard Model, as mentioned before, does not provide for neutrino mass in its Lagrangian. In analogy with the mass terms for the other fermions, one might consider introducing the term

$$-\mathcal{L}_D = \frac{1}{2} m_D (\bar{\nu}_L \nu_R + \bar{\nu}_R \nu_L) \tag{1.51}$$

where m_D is a constant, the Dirac mass. However, the presence of this term would at first seem to violate Occam's Razor by postulating an undetectable particle ν_R .

Particles may be considered either left-handed or right-handed depending on their helicity, a quantum number given by

$$\mathbf{h} = \frac{\mathbf{s} \cdot \mathbf{p}}{|\mathbf{p}|}.\tag{1.52}$$

The Standard Model treats fermions differently according to the sign of this quantum number. Left-handed particles exist in weak doublets, for instance (e_L^-, ν_{eL}) ; weak interactions may mediate at Feynman diagram vertices between them. Right-handed particles are singlets, and do not interact via the weak force. (The roles of L and R are reversed for antiparticles.) Thus, in the Standard Model assumption of massless neutrinos, processes such as β -decay must be maximally parity-violating: the neutrino is always emitted with left-handed chirality in β^+ decay (likewise, in β^- decay, the antineutrino is emitted with right-handed chirality).

With the assumption of massless neutrinos, there is no difficulty because, as they travel at the speed of light, it is impossible for them to reverse spin and become right-handed. Neutrinos interact only via the weak force, so even if right-handed neutrinos existed, they could not be detected in any way (except, in principle, through their cumulative gravitational force). This type of neutrino is termed *sterile*. When neutrinos have mass, however, they may convert between left- and right-handed states routinely. Though this is an odd side-effect (we should now expect to see neutrino oscillations into invisible states as well as between flavors), adding a Dirac term for neutrinos to the Standard Model Lagrangian might be an acceptable way to allow for neutrino mass. However, the Dirac term leaves unanswered the question of why neutrino masses are so many orders of magnitude smaller than those of any other fundamental particle.

A popular hypothesis that solves this problem is the see-saw mechanism, so named because it creates the observed neutrino masses by balancing a Dirac mass similar to that of the charged leptons with extremely heavy and extremely small "Majorana masses." Neutrinos may in principle be governed by a Majorana term in the Lagrangian,

$$-\mathcal{L}_M = \frac{1}{2} m_R^M \left(\overline{(\nu_R)^c} \,\nu_R + \text{h.c.} \right) + \frac{1}{2} m_L^M \left(\overline{(\nu_L)^c} \,\nu_L + \text{h.c.} \right), \qquad (1.53)$$

where $m_{L,R}^M$ are Majorana masses, not necessarily identical. Unlike the Dirac term, this term does not conserve lepton number L! Such a term is therefore only available to chargeless fermions (*i. e.*, neutrinos) since, if applicable to quarks or the electron family of leptons, it would additionally violate conservation of electric charge. Though neutrino oscillation clearly violates conservation of the individual lepton numbers $L_{e,\mu,\tau}$, breaking lepton number conservation completely is a bit much to swallow. Nevertheless, the see-saw mechanism requires it. The mechanism supposes that both Dirac and Majorana terms are present for neutrinos, and are mixed as follows:

$$-\mathcal{L}_{\text{mass}} = \frac{1}{2} \left(\bar{\nu}_L \ \overline{(\nu_R)^c} \right) \begin{pmatrix} m_L^M & m_D \\ m_D & m_R^M \end{pmatrix} \begin{pmatrix} (\nu_L)^c \\ \nu_R \end{pmatrix} + \text{h.c.}$$
(1.54)

In the most general case, each neutrino family will have a separate term of this nature, possibly with different values for the $m_{L,R}^M$ and m_D . (As a side effect, the MNS mixing matrix will contain an additional factor beyond those shown in Equation (1.3), a factor multiplying the value shown for $U_{\rm MNS}$ on the right by diag $(1, e^{i\phi_2}, e^{i\phi_3})$; the ϕ_i 's are CPviolation-inducing Majorana phases.)

For specificity, consider the m_1 neutrino mass state. When $m_R^M \gg m_D$ and $m_L^M \ll m_D^2/m_R^M$, the two eigenvalues of the above mass matrix are approximately m_R^M and m_D^2/m_R^M . Supposing that $m_D \sim m_e$, making the neutrino Dirac mass comparable to that of the charged leptons, we can let m_R^M be large enough to ensure $m_1 \sim m_D^2/m_R^M$. Then the two eigenstates of the mass matrix become approximately equal to $|\nu_L\rangle$ (light) and $|\nu_R\rangle$ (heavy). This would then explain why we only observe left-handed neutrinos; the probability to find a light right-handed neutrino is extremely small! The heavier particles, known as the N_i , may have masses up to the GUT scale. CP violation in their decay may conceivably have resulted in the current observed baryon asymmetry [23], the fact that antimatter is extremely rare in the universe today.

Mass-varying neutrinos

Recent observations have indicated that the relative densities of cold dark matter and dark energy in the universe are similar; in particular, $\Omega_{\rm CDM}/\Omega_{\Lambda} \approx \frac{1}{3}$. However, the ratio is changed as a function of the cosmic scale factor (the ratio between current and past comoving distances as the universe expands) as $1/a^3$. This coincidence seems rather too convenient for many physicists' taste. One possible, though slightly outré, explanation is that the mass of a neutrino is itself a variable quantity influenced by the local number density of neutrinos [24, 25]. The hypothesis of mass-varying neutrinos (MaVaNs) is modeled by their interaction with a field of "accelerons," so named because this supposition causes neutrinos to behave essentially as dark energy, accelerating the expansion of the universe. Constraints on naturalness of the model indicate that sterile neutrinos, a consequence of the model, have a mass on the order of 1 eV. This prediction could be tested by the MiniBooNE experiment; it has been suggested that it would also explain the puzzling LSND results [26] (see below). If confirmed, this model would have important implications for both cosmology and particle physics.

1.4.2 Sterile neutrinos and LSND

Though Big Bang nucleosynthesis puts a tight restriction on the number of light weakly interacting neutrinos, essentially constraining it to equal three if non-integer values are disallowed, it can say little about sterile neutrinos. Hypotheses such as the MaVaN model even question its ability to do this much. Sterile neutrino models are known as "3+n," n being the number of sterile neutrino flavors.

Fans of sterile neutrinos are therefore heartened by the result of the LSND experiment. LSND was an experiment performed from 1993–1998 with accelerator μ -antineutrinos up to 53 MeV; it attempted to detect the generation of electron antineutrinos in a baseline of only 30 m. The observed results suggested a neutrino oscillation caused by a squared mass difference $\Delta m_{\rm LSND}^2 c^4 > 0.2 \, {\rm eV}^2$ [27]. A simple look at Figure 1.7 makes it clear that one cannot have a value of Δm^2 so different from the known values of Δm_{12}^2 and $|\Delta m_{23}^2|$ without a fourth neutrino having a still different mass. As only three flavors of charged leptons are known, this fourth neutrino would have to be sterile. Oscillations from other flavors into it would look like the complete disappearance of the neutrino. However, other analyses have indicated that both possible types of four-neutrino mass hierarchies (2+2, with two pairs of similar masses; and 3+1, the situation of Figure 1.7 plus an additional mass eigenstate even more isolated than ν_3) exhibit very poor goodnessof-fit parameters [28]. The status of the LSND result is therefore quite puzzling. Some have suggested an indication of CPT violation [29], which would permit the probabilities of neutrino and antineutrino oscillation to differ. Others have proposed a 3+2 sterile neutrino model. Many of these questions could be resolved, or at least the peculiar observations confirmed, when the MiniBooNE experiment [30], an attempt to replicate the LSND result at a similar L/E ratio with much higher statistics, makes its experimental results public this year. A confirmation of the LSND observations would be an excellent indicator for new physics.

1.4.3 Neutrinoless double beta decay

If (as expected from the see-saw mechanism) the neutrino is a Majorana particle, then it can act as its own antiparticle. This hypothesis is testable: in particular, it would make the phenomenon of neutrinoless double β decay possible.

"Normal" double β decay $(2\nu 2\beta)$ is a process in which two electrons and two antineutrinos are emitted from a nucleus at once, that is,

$${}^{A}_{Z}X \to {}^{A}_{Z+2}Y + 2e^{-} + 2\bar{\nu}_{e}.$$
 (1.55)

There exist radioactive isotopes for which this process is energetically favorable, but a single β -decay is not; that is, the Q-value of the reaction ${}^{A}_{Z}X \rightarrow {}^{A}_{Z+1}Y' + e^{-} + \bar{\nu}_{e}$ is negative. Hence all four particles must be emitted within a time-frame short enough that the intermediate off-mass-shell state is permissible. Because satisfying this constraint is difficult, the half-lives of these isotopes are very long. For instance, that of 82 Se ($\rightarrow {}^{82}$ Kr $+2\beta^{-} + 2\bar{\nu}_{e}$) is 9.6 × 10¹⁹ yr [31]. Many isotopes that could be candidates for a double β decay (86 Kr, for example) are, however, observed to the best of our abilities to be stable.



Figure 1.8: Neutrinoless double β decay. A nucleus emits two virtual W^- bosons. The first decays into an electron and an electron antineutrino; the antineutrino interacts with the second boson to produce an electron. This diagram cannot exist unless (anti)neutrinos are their own antiparticles, *i. e.*, are Majorana particles, and total lepton number is violated.

Neutrinoless double β decay $(0\nu 2\beta)$ is a hypothetical process very similar in nature. Indeed, just as in normal $2\nu 2\beta$ decays, a nucleus emits two virtual W^- bosons. The crucial difference is that in $0\nu 2\beta$ decay (Figure 1.8), the antineutrino produced by the first boson interacts with the second to produce an electron. This reaction is not possible in the Standard Model since $W^- + \bar{\nu}_e \rightarrow e^-$ violates conservation of total lepton number. It may only occur if neutrinos are in fact Majorana particles and may act as their own antiparticles.

As with single-electron β decay, $2\nu 2\beta$ decay exhibits a continuous electron kinetic energy spectrum because much of the kinetic energy produced is carried away by antineutrinos, becoming basically invisible. However, in the case of $0\nu 2\beta$ decay, the total kinetic energy of the two electrons would be nearly monoenergetic (the recoil imparted to the daughter nucleus being negligible). Therefore, the signature of a $0\nu 2\beta$ decay event would be a peak in the observed energy spectrum at the Q-value of the decay. The $2\nu 2\beta$ continuous electron spectrum, on the other hand, extends from zero energy up to essentially the Q-value.

In a $0\nu 2\beta$ experiment, the continuous spectrum is unimportant. The relevant quantity is $m_{\beta\beta}$, the averaged value

$$m_{\beta\beta} = \left| \sum_{i} |U_{ei}|^2 e^{i\phi_i} m_i \right|. \tag{1.56}$$

(The ϕ_i 's are the same Majorana phases as those mentioned previously in connection with



Figure 1.9: The Heidelberg-Moscow claim for observation of neutrinoless double β decay of ⁷⁶Ge. At left, the region where the $0\nu 2\beta$ peak should appear given the calculated Q-value of 2.039 MeV. (The claimed peak is a bit right of center.) Other peaks are due to γ rays emitted in the radioactive decay of ²¹⁴Bi. At right, the full energy spectrum, with a fit (solid red line) to the theoretical $2\nu 2\beta$ continuous spectrum. Only events that matched a filter on pulse shape are shown in the spectra. Other experiments have been unable to reproduce this result. Figure taken from reference [32].

the see-saw mechanism; without loss of generality we may set $\phi_1 = 0$.) The quantity $m_{\beta\beta}$ may be determined because the partial half-life of the $0\nu 2\beta$ decay is inversely proportional to the effective neutrino mass squared:

$$\tau_{1/2}^{0\nu} = \left(G^{0\nu} | M^{0\nu} |^2 m_{\beta\beta}^2\right)^{-1},\tag{1.57}$$

where $G^{0\nu}$ is an accurately determined phase space factor, and $|M^{0\nu}|^2$ is an (in principle) independently observable nuclear matrix element. In more practical terms, $m_{\beta\beta}$ is proportional (through a more-or-less known proportionality constant) to the square root of the number of events observed in the $0\nu 2\beta$ energy spectrum peak (if observed) of a laboratory experiment. Hence, non-observation of a peak implies an upper limit on $m_{\beta\beta}$.

Clearly it is very important for a $0\nu 2\beta$ experiment to have very high energy resolution, in order for the $2\nu 2\beta$ continuum not to overwhelm the desired peak, and also so that other signals in the same energy range (mainly from γ rays produced in the decay of uranium and thorium chain isotopes) do not overlap. Generally, therefore, only calorimetric detectors are feasible. In total about nine different isotopes have been observed by various experiments in attempts to detect $0\nu 2\beta$ decay [31].

So far only one experiment, the Heidelberg-Moscow enriched Ge experiment, has claimed a positive result for observation of $0\nu 2\beta$ decay. They give an estimate of $m_{\beta\beta} c^2 = 440^{+140}_{-200}$ meV (the errors are given as 3σ values) [32]. This result is at the moment still hotly disputed in the physics community, particularly as it seems to be incompatible at $> 2\sigma$ with other limits on neutrino masses [21]. The best upper limit from any other experiment for $m_{\beta\beta} c^2$ is 0.2-1 eV [31]; the range is due to large uncertainties in the required nuclear matrix elements. A next generation of $0\nu 2\beta$ experiments will have sensitivities in the range ~ 80 meV, even taking the nuclear matrix element uncertainties into account, and will be able to conclusively confirm or refute the Heidelberg-Moscow result.

As with measurements of the single β -decay neutrino mass m_{β} , a sufficiently small upper limit on $m_{\beta\beta}$ would (under the assumption that neutrinos are in fact Majorana particles) rule out the inverted mass hierarchy. However, the required limit on $m_{\beta\beta}$ is smaller than that of m_{β} . Since such a disproof would additionally be conditional on demonstrating the Majorana nature of the neutrino, it is likely that tritium decay experiments will be the first to conclusively refute (if it in fact does not occur in nature) the inverted mass hierarchy.

1.4.4 New physics accessible to a low-energy solar neutrino detector

None of the issues discussed above are really amenable to investigation by solar neutrino detectors, which have sensitivities in the energy region of about 0.25–10 MeV. Still, a low-energy solar neutrino detector could nevertheless explore unexpected physics. Observations that suggest variations in the neutrino flux [33] consistent with the solar rotation period of 28 days and the solar cycle of 11 yr might be explained by subdominant new physical processes such as resonant-spin-flavor-precession (RSFP). This process would require neutrinos to have a magnetic moment, giving them some probability to be "flipped" into a



Figure 1.10: Expected shape of the vacuum-matter transition curve (the survival probability of ν_e as a function of energy) for the standard LMA model "LMA-I" (higher, blue curves) and a model including non-standard interactions, LMA-0 (lower, violet curves). Solid lines represent the predicted survival probability functions for neutrinos produced in the Sun with a radial distribution following that of ⁸B neutrinos, and dashed lines are for neutrinos produced following the *pep* neutrino radial distribution function. (These radial distribution functions are shown in Figure 1.5.) Note the striking differences in the predicted survival probabilities at the *pep* and ⁷Be neutrino energies.

different flavor, into an antineutrino, or even into a sterile right-handed state, by the Sun's magnetic field [34]. Any hypothetical discovery of a day-night asymmetry in low-energy solar neutrinos, or of a seasonal variation in the neutrino flux beyond that expected from the shape of Earth's orbit, would imply neutrino oscillation wavelengths that can be explained by neither vacuum nor standard MSW oscillations. The most interesting, as-yet unexplored, and potentially easy-to-detect possibility, though, may be that neutrinos have non-standard interactions (NSIs) not permitted within the scope of the Standard Model.

NSIs, in their most general form, connect four fermions or antifermions involved in a reaction, of which at least two are (anti)neutrinos. An NSI may be either flavor-conserving or flavor-changing. For instance, an electron neutrino could scatter from a u quark and in the process be converted to a mixture of flavor eigenstates. As a result, instead of having a single non-zero entry, the matrix in Equation (1.14) may be completely filled:

$$H = U_{\rm MNS} H_0 U_{\rm MNS}^{\dagger} + k_{\rm NC} I + k_{\rm CC} \begin{pmatrix} 1 + \epsilon_{ee} & \epsilon_{e\mu}^* & \epsilon_{e\tau}^* \\ \epsilon_{e\mu} & \epsilon_{\mu\mu} & \epsilon_{\mu\tau}^* \\ \epsilon_{e\tau} & \epsilon_{\mu\tau} & \epsilon_{\tau\tau} \end{pmatrix}$$
(1.58)

The various dimensionless ϵ 's describe the probability for an NSI to occur relative to the charged-current interaction strength $k_{\rm CC} \equiv \sqrt{2}G_F n_e$. The definitions of the ϵ 's encompass all NSI effects in normal matter, both with electrons and with u and d quarks. Thus in general they are functions of chemical and isotopic compositions.

It should be clear from the discussion of Section 1.2.2 that the effect of these NSIs will go to zero in vacuum, and that in matter they will cause mixing between energy and flavor eigenstates which increases as n_e increases, similar to the standard MSW effect. The energy eigenvalues will be different than those caused by the standard MSW effect, as will the mixing angles. Since in principle all three neutrino flavors may mix in NSIs, the mixing matrix may also be forced to include a complex phase multiplying a non-negligible term.

What at first seems like a slight complication to the normal MSW effect could have real observable consequences. One region of values for the ϵ parameters, with magnitudes for the parameters no greater than 0.25, was found to permit an alternative solution to the observed solar neutrino and KamLAND data in the $(\Delta m_{12}^2, \theta_{12})$ parameter space. Specifically, $\Delta m_{12}^2 c^4 = 1.5 \times 10^{-5} \,\mathrm{eV}^2$, while $\tan \theta_{12} = 0.39$ [35], a solution dubbed "LMA-0." (Compare to the accepted standard LMA values of $8.0 \times 10^{-5} \,\mathrm{eV}^2$ and 0.45 [4].) With these parameters (and the assumed values for the ϵ coefficients), the shape of the solar electron neutrino survival probability as a function of E_{ν} —in particular, at the vacuum-matter transition energy—becomes quite different (Figure 1.10). It will be possible for a sub-MeV solar neutrino detector having the capability to observe scattering event energies to probe the shape of this curve at crucial points by studying the monoenergetic 862-keV ⁷Be and 1.44-MeV pep neutrinos. In so doing, we may very well get a glimpse of new physics!

Chapter 2

The Borexino Neutrino Experiment

Borexino is a neutrino detector located deep underground in central Italy, under a thickness of rock equivalent to 3800 meters of water. It is expected to begin operations in Fall 2006, complementing the similarly designed but already operational KamLAND neutrino detector in Japan and the Čerenkov light detecting Sudbury Neutrino Observatory in Canada. Both of those detectors are targeted at higher-energy neutrinos. Borexino should be the first detector capable of observing sub-MeV solar neutrinos in real time. Indeed, its principal goal is to measure the ⁷Be neutrino flux to high accuracy. The possibility to measure the rates of the slightly higher-energy *pep* and CNO-cycle neutrinos is also foreseen. The expected operational lifetime of the experiment is ten years, which will give it a reasonable chance to observe neutrinos from a supernova in our galaxy or nearby. Measurements of neutrinos from other sources (geoneutrinos, reactor neutrinos) may also be feasible.

The name of the detector comes from "Borex," since the fluid at the heart of the detector was originally planned to be trimethylborate-based [36], and "-ino," the Italian suffix meaning "little," as it was originally intended to be a prototype of a still-larger detector. Both parts of the name are now inaccurate, but it remains unchanged largely because of the well-known physics phenomenon of inertia. Various capitalizations are seen in the literature. The principal reaction by which Borexino will detect neutrinos is elastic νe scattering. The neutrino is not seen directly, but it imparts some of its kinetic energy to an electron. As the electron travels, it is slowed down by ionizing interactions with the surrounding fluid medium. The 2200 photomultiplier tubes in Borexino will mainly observe light emitted by molecules of the target mass that were electronically excited by the ionizing processes. The target mass is referred to as a "scintillator," and so Borexino is a scintillation detector.

Energy loss of a traveling charged particle also occurs via Čerenkov radiation. However, the number of Čerenkov photons detected will be at most a few percent of the total. Unlike Čerenkov light, scintillation light produced by excited molecules returning to their ground state is emitted isotropically. It will therefore not be possible to determine, even approximately, the initial direction of a neutrino that caused electron scattering in the detector.

In this chapter we discuss the radioactive backgrounds likely to be present in the Borexino detector (many of which have already been observed in its prototype Counting Test Facility) and the expected sensitivities of Borexino to neutrinos produced by various sources. A summary table is provided on the following page for quick reference.

2.1 Design parameters

Any radioactive decay inside the detector has some probability of being mistaken for a neutrino interaction. Many naturally occurring radioactive isotopes, in their decays, release energies in the sub-MeV range which is most important for the detector. Borexino is therefore designed around the principle of graded shielding, in which the central portion of the detector is the least contaminated. A detailed discussion of the detector scintillator and hardware can be found in Chapter 3. The following paragraphs, in tandem with Figure 3.1, should suffice to give a basic idea of the experimental design.

Signal	Source	Expected eve	References			
		$250800\mathrm{keV}$	$0.8 {-} 1.3\mathrm{MeV}$			
Internal backgrounds	(Individual isot	opes are each d	also described in	§2.2.2–2.2.3.)		
μ^{\pm}, n	Cosmic rays	$\sim 285~(0.4)$	$\sim 250~(0.3)$	[37, 38]		
$^{7}\mathrm{Be}$	Cosmogenic	0.4	-	[37, 39]		
$^{11}\mathrm{C}$	"	-	5.1 (1.0)	2.3.3, [22],		
				[40, 41, 42]		
$^{39}Ar + {}^{85}Kr$	LAK N_2	0.1	-	[42, 43]		
$^{14}\mathrm{C}$	Scintillator	1.5	-	[44]		
40 K, K _{nat} = 10^{-14} g/g	contam.	1.5	0.4	8.5.2, 9.4.2		
238 U chain 10^{-16} g/g	"	97.5(18.8)	13.1 (2.2)	8.2,		
232 Th chain 10^{-16} g/g	"	22.9(5.4)	4.0 (1.2)	8.4, [12]		
²²² Rn daughters	Nylon eman.	5.2(0.4)	1.2 (0.0)	4.4.3, [45]		
²¹⁰ Pb daughters	" wash-off	0.1 (0.0)	< 0.1	4.4.4, 8.3,		
"	Metal wo.	?	?	<i>[</i> 46, 47]		
External backgrounds						
$\gamma \text{ rays}$	PMTs/lt.conc.	1.2	0.8	[12, 48],		
"	Other	1.6	0.7	(49]		
Solar neutrinos						
$^{7}\mathrm{Be}$	Sun	27.7	-	2.3.1		
pep	"	1.3	0.9	2.3.3		
^{13}N	"	1.3	-	"		
^{15}O	"	1.5	0.5	"		
⁸ B	"	$0.1 \ [E >$	$2.8\mathrm{MeV}]$	[12, 50]		
Other neutrinos	(Expected events/year in 300 tons, full spectrum)					
$\bar{ u}_e$	Earth	10 - 24	Ň	2.4.1,		
$ u_e, ar{ u}_e$	Reactors	28 [8 wit	th $E < 3 \mathrm{MeV}$]	$\int [51, 52]$		
$ u_{e,\mu, au},ar{ u}_{e,\mu, au}$	$10\mathrm{kpc}~\mathrm{SN}$	186 events	in $\sim 10 \mathrm{s}$	2.4.2, [53, 54]		

Table 2.1: Summary table of expected signals and backgrounds in the Fiducial Volume (except where otherwise indicated) of Borexino. Note that as yet there is no accurate estimation for the internal scintillator contamination of ⁴⁰K and the heavy-element chains in Borexino; the values given are for the specified assumed contaminations. The amount of contamination due to ²¹⁰Pb wash-off from metal surfaces will depend upon their original exposure, and cannot be estimated in advance. Values in parentheses, where present, represent the reduced rates after individual event tagging methods are applied: searches for delayed coincidences, α/β discrimination, likelihood-based tagging (this largely removes the signal from ²¹⁴Pb; see Section 8.4), and so on. References are to sections in this thesis or to bibliographic entries; the list is at best representative, not exhaustive.

The central part of the detector consists of 300 tons of the organic liquid pseudocumene, in which is dissolved a wavelength-shifting compound that acts both to shift the scintillation light towards the wavelength at which the photomultiplier tubes are most efficient, and to increase the total production of scintillation light. Only the central 100 tons, which comprise the region best shielded against radiation from external sources, are considered to be the target mass. Since the data collected by the detector can be analyzed to determine the approximate position of any event, as will be discussed at length in Chapter 5, it will be possible to separate out the events that occur within this most pure region. This spherical "Fiducial Volume" may be redefined at any time, as desired, by simple changes in the Borexino analysis software.

The scintillator fluid is contained within two concentric spherical vessels of thin $(125 \,\mu\text{m}$ thick), transparent nylon film. The inner vessel has a radius of about 4.25 m, and the outer vessel, 5.5 m. Both vessels are intended to prevent radioactive and highly mobile atoms of radon gas from traveling into the scintillator fluid. In addition, the inner vessel serves to separate the scintillator fluid (within it) from a buffer fluid (between the two vessels, and outside the outer one). The buffer fluid is made of pseudocumene laced with a substance that prevents scintillation. The purpose of the two volumes of buffer fluid is to shield the target mass from external radiation. They are prevented from scintillating themselves in order to keep the trigger rate of the data acquisition system at a few tens of Hz.

A large steel sphere, 6.85 m in radius, surrounds and supports the two nylon vessels. The vessels are attached to the sphere at top and bottom by steel and bulk nylon tubing. The sphere also holds about 2200 inward-pointing photomultiplier tubes. These instruments observe photons produced within the scintillator. They provide the main means by which the energy of an event may be determined.

The energy resolution of the detector is expected to follow Poisson statistics, since it is based upon the number of scintillation photons emitted by excited molecules of scintillator, and upon the fraction of those photons which both strike a photomultiplier tube and are converted to a photoelectron. When the total number of observed photoelectrons is reasonably large (on the order of 20 or more), the energy spectrum of a monoenergetic source should closely approximate a Gaussian curve.

Borexino is expected to observe 400 photoelectrons (p.e.) for each MeV of kinetic energy released in a neutrino scattering event or radioactive β decay [55]. This ratio is nearly linear as long as the energy release is larger than 50–100 keV [56]. Therefore, if it were an ideal detector, the energy resolution for an event with true energy E would be given by¹

$$\frac{\sigma_E}{E} = \frac{\sigma_{N_{p.e.}}}{N_{p.e.}} = \frac{1}{\sqrt{N_{p.e.}}} = \frac{k_0}{\sqrt{E}},$$
(2.1)

with $N_{p.e.}$ being the number of photoelectrons detected in the event, and the proportionality constant $k_0 = \sqrt{(1 \text{ MeV})/(400 \text{ p.e.})} = 0.05 \text{ MeV}^{1/2}$. Hence the expected resolution σ_E for 250 keV events is 25 keV (10%), and for 1 MeV events, it is 50 keV (5%). In practice the value of the proportionality constant will be greater than k_0 due to light scattering, absorption, variances in photomultiplier efficiencies, and other effects. Denote the observed proportionality constant as k. For the third version of the Borexino prototype, the Counting Test Facility (CTF 3), k/k_0 is found in Section 8.1.2 of this work to be 1.34. We will suppose this also to be the case for Borexino; that is, we take $k = 0.067 \text{ MeV}^{1/2} = 2.11 \text{ keV}^{1/2}$.

Most likely an event in which a particular kinetic energy is released will result in the detector observing a number of observed photoelectrons that corresponds to a different energy. We will call the former value E_r ("r" for "real") and the latter value E_d ("d" for "detected"). The probability that a particle of kinetic energy E_r will be observed to produce a number of photoelectrons corresponding to kinetic energy E_d defines a resolution function $R(E_r, E_d) dE_d$. If we assume a Gaussian energy resolution, then $R(E_r, E_d) = (2\pi\sigma_E^2)^{-1/2} e^{-(E_d - E_r)^2/2\sigma_E^2}$. Substituting in Equation (2.1) gives us

$$R(E_r, E_d) = (2\pi k^2 E_r)^{-1/2} e^{-(E_d - E_r)^2/2k^2 E_r}.$$
(2.2)

¹Throughout the remainder of this thesis, the kinetic energy of a particle produced by a scattering or decay event will be designated E. When it is necessary to give the total energy of a particle, including the rest mass, that quantity will be designated U. In the case of neutrinos, since their masses are negligible compared to decay energies or the electron mass, the distinction will not be made.

Suppose that in some process, a particle is given a kinetic energy in the range [E, E + dE]with probability $\rho_r(E) dE$. That is, $\rho_r(E)$ is the energy spectrum of the process (again, the subscript "r" stands for "real"). It can be transformed into the form that will be observed in Borexino by a convolution with the resolution function. The convolution yields an observed energy spectrum

$$\rho_d(E) = \int_0^\infty R(E', E) \rho_r(E') \, \mathrm{d}E' = (2\pi k^2)^{-1/2} \int_0^\infty \frac{\rho_r(E')}{\sqrt{E'}} e^{-(E-E')^2/2k^2E'} \, \mathrm{d}E'. \tag{2.3}$$

When the process is monoenergetic, the real energy spectrum is a δ -function, and the observed energy spectrum becomes a Gaussian curve centered at the true event energy E_r , with width $\sigma = k\sqrt{E_r}$. For processes with a continuous spectrum, the convolution tends to smooth out the observed spectrum, as well as extend it slightly beyond the true high-energy end-point of the process.

2.2 Radioactive backgrounds

In this section we will first recall the basic types of background in Borexino, then break them down into individual isotopes and note the origin and potential danger posed by each. The section will be rather detailed since the information provided here will be used again in the final three chapters of this work.

2.2.1 Types of radioactive decay

A radioactive nuclide may decay in several different ways. It may emit an α , β^- , or β^+ particle (these are all often accompanied by one or more γ rays); capture an electron from the K shell; decay by pure γ ray emission to a lower energy state of the same nuclide; or break apart into roughly equal-sized pieces by spontaneous fission. All of these decay methods except the last are important contributors to the background event rate in Borexino. Additionally, in Borexino, neutrons may be produced by the reaction of incoming cosmic rays or α particles (themselves resulting from radioactive decay) with carbon nuclei. Since a free neutron will be captured by a proton (forming deuterium) with emission of a characteristic 2.2 MeV γ ray after a typical time of 250 μ s, they are not usually problematic.

Using the standard notation, consider an unspecified isotope of some element X, which has Z protons and N neutrons for a total of A = Z + N nucleons. Depending on whether Z and N are each even or odd, the isotope is classified as "even-even," "odd-odd," etc. Nucleons, as spin- $\frac{1}{2}$ fermions, find it energetically favorable to pair off within the nucleus. In general, therefore, even-even isotopes are the most stable, and odd-odd isotopes the least: only four odd-odd isotopes (²H, ⁶Li, ¹⁰B, ¹⁴N) are stable, ²H only barely.²

A nucleus is subject to three of the four fundamental forces (the effect of gravity is negligible), one struggling to hold it together and the other two to tear it apart. The binding force is the residual strong interaction. Nucleons, each being composed of three quarks, are colorless objects. Yet there is still some residual attraction between them, akin to the force seen between electric dipoles that have no overall charge. This attraction is mediated by the exchange of pions. There are three types, having charges of 0 and ± 1 . (Their quark constituents are, for the π^+ , $u\bar{d}$; for π^- , $d\bar{u}$; and for π^0 , $(u\bar{u} + d\bar{d})/\sqrt{2}$.) Thus they can mediate interactions between all three possible combinations pp, nn, and pn. The residual strong interaction is short-range and acts to create a potential energy well within a nucleus.

The other forces, tending to cause radioactive decay, are electrical repulsion between the positively-charged protons, and the weak interaction that permits conversions between protons and neutrons. On one hand, as the number of neutrons is increased, the isotopes of an element become less susceptible to α decay or fission caused by electrical repulsion. Adding neutrons increases the total nucleon binding energy while keeping the electrical repulsion constant. So the ratio N/Z is usually greater than one for heavy stable isotopes (reaching a maximum of about 1.5). On the other hand, if N and Z are too different, the Fermi energy level of the neutrons will be enough higher than that of the protons to make conversion of

 $^{^{2}}$ The 2 H nucleus, or deuteron, is an interesting case. It has only one energy level, the ground state, which if excited by more than 2.22 MeV will dissociate into its component proton and neutron.

a neutron into a proton via the weak process of β decay energetically favorable. The two stability constraints (large N on one hand and $N \approx Z$ on the other hand) become more and more difficult for heavy nuclei to satisfy, so there are no stable isotopes with Z > 83 or $N > 126.^3$

The difference between the mass of a radioactive isotope and the total mass of the products into which it decays is Q, the amount of energy released in the decay. This is called the Q-value; it must be greater than zero for the decay to occur. Energy released may appear as kinetic energy of the decay products, as γ -ray photons, or both. If more than one possible decay branch has a positive Q-value, the isotope may decay in any of these ways, with a specific probability for each.

α decay

Emission of an α particle, also known as the nucleus of ⁴He, is common among very heavy isotopes ($A \sim 200$ or more). The α particle, consisting of a pair of protons and pair of neutrons, is a very stable configuration. This decay mode is described by the Gamow-Condon-Gurney theory as the escape of a pre-formed α particle from the nuclear potential energy well by quantum tunnelling:

$${}^{A}_{Z} \mathbf{X} \to {}^{A-4}_{Z-2} \mathbf{Y} + \alpha. \tag{2.4}$$

The radius of the potential well, for the heavy elements that emit α particles, is approximately $R \approx 9.5$ fm for even-even isotopes [57].

Because only the α particle is emitted, the decay is monoenergetic, with the α having a kinetic energy almost equal to the decay's Q-value. The α energy is typically in the range 4–9 MeV for heavy nuclei. The remainder of the energy goes into recoil of the nucleus, which receives a kinetic energy on the order of 100 keV in order to conserve momentum of

³Technically, even the single common isotope of element 83, ²⁰⁹Bi, is slightly radioactive, though its half-life is many times the age of the universe.

the system. An important property of α decays for ultra-low background detectors is that the release of energy from them is very localized; α particles will typically travel less than 0.1 mm in solid or liquid materials. Therefore one needs to worry about α -decaying isotopes only in the scintillator or the vessel containing it, not in any of the other materials of the detector.

The Gamow-Condon-Gurney model can be used to show that the mean life τ of the isotope⁴ depends exponentially upon the Q-value of its decay, and also upon the atomic number Z' = Z - 2 of the decay product:

$$\log \frac{\tau}{\tau_0} \approx Z' \sqrt{\frac{E_0}{Q}} - C\sqrt{Z'}.$$
(2.5)

Here, $\tau_0 \equiv 2R/v_{\alpha}$ is the time it takes for the α particle to travel across the original nucleus, given roughly by $(2.74 \times 10^{-21} \text{ s})\sqrt{(1 \text{ MeV})/Q}$. The other constants are defined by $E_0 \equiv 8\pi^2 \alpha^2 m_{\alpha} c^2 \approx 15.67 \text{ MeV}$ and $C \equiv 8\sqrt{\alpha (Rm_{\alpha}c^2)/(\hbar c)} \approx 9.16$. This equation, the Geiger-Nutall relation, implies that isotopes with long half-lives emit low-energy α particles ($Q \approx 4 \text{ MeV}$), while very short-lived isotopes have higher-energy decays ($Q \approx 8 \text{ MeV}$).

Decays by α emission must of course conserve angular momentum. The α particle itself has spin 0; therefore the orbital angular momentum ℓ_{α} of the α must add to the spin of the daughter nucleus to yield the spin of the parent nucleus. If we write the spins of parent and daughter nuclides as J_i and J_f respectively, then the quantum mechanical rules for spin imply the inequality $J_f + \ell_{\alpha} \geq J_i \geq |J_f - \ell_{\alpha}|$. For an α particle emitted with no angular momentum, the parent and daughter will have equal spins. In principle, however, an α decay may connect any two spin states whose spins differ by an integer if the angular momentum of the α particle has the requisite value.

Often connected with angular momentum is the idea of parity. This is a quantum number π whose value is +1 for a state whose wave function is invariant under reversal of coordinate signs ($\psi(-\mathbf{x}) = \psi(\mathbf{x})$), and -1 for a state whose wave function changes sign under such a

⁴After one mean lifetime, a fraction 1/e of an initial sample of radioactive atoms will remain. For this reason, the symbol "log" will always be used in this thesis to represent the *natural* logarithm.

reversal $(\psi(-\mathbf{x}) = -\psi(\mathbf{x}))$. These values are often abbreviated simply + and -, or even and odd, respectively. Individual particles and nuclei have their own intrinsic parities: that of the α particle itself is +. Spin and parity are usually written together in the format J^{π} , for instance, $\frac{1}{2}^+$. The important point to note is that the spherical harmonic functions $Y_{\ell m}$, which describe the wave function of an α particle with angular momentum ℓ , have parity $(-1)^{\ell}$. Hence, for parity to be conserved in α decay, the parity of the daughter nucleus must be $(-1)^{\ell}$ times that of the parent. In particular, a nucleus with a ground state of 0^+ may decay by α emission only to a daughter in one of the states $0^+, 1^-, 2^+$, etc.

An α decay can in principle result in the daughter nucleus being in an excited state, meaning that it shortly afterwards emits one or more γ rays. However, all even-even isotopes have a ground state in which the nucleons are paired off with antiparallel spins. In this state, all the spins cancel: the spin and parity of the nuclide is 0^+ . Since α decay preserves the evenness of N and Z, the product also has a ground state 0^+ . All other factors being equal, the α particle itself is most likely to exit the parent nucleus with no angular momentum, because there is then no additional centrifugal energy barrier helping to prevent its escape. Recall also that a higher-energy α decay implies a shorter half-life: the α decay with the highest possible energy, that to the ground state of the daughter isotope, is thus strongly favored. As a result of these two factors, the most likely α decay mode of an even-even isotope is to the ground state of the daughter, with no γ rays produced. This situation holds for all α -decaying isotopes in the ²³²Th and ²³⁸U decay chains, with the single exception of odd-odd ²¹²Bi. Except for decays of that isotope, only a small fraction of α events in Borexino will produce γ rays. Since γ rays may travel many cm in scintillator, this fact makes α -decaying isotopes useful in evaluating the position reconstruction capabilities of scintillation detectors.

β decay

 β decay, unlike α decay, is a weak phenomenon. When a nucleus has too many neutrons, one neutron in the nucleus may become a proton. Odd-odd isotopes are particularly prone to this and related processes because they are converted to even-even isotopes. In β decay, both an electron (also called a β^- particle) and an electron antineutrino $\bar{\nu}_e$ are emitted:

$${}^{A}_{Z}X \rightarrow {}^{A}_{Z+1}Y + e^{-} + \bar{\nu}_{e} \tag{2.6}$$

The electron loses all of its kinetic energy to collisions with scintillator molecules within a few cm. β decay is therefore observed as an essentially point-like event (the position resolution of Borexino will be on the order of 10–20 cm). Like α decays, purely β -decaying isotopes—those which always decay to the ground state of the daughter isotope—only contribute to detector background when present in the scintillator itself or the vessel containing it. (In principle β decays in the volume immediately outside the vessel, if it is thin, may also contribute.)

At the quark level, a d quark emits a virtual W^- , becoming a u quark. The W^- immediately decays into an electron. A $\bar{\nu}_e$ is also emitted in order to conserve lepton number:



Since two additional particles are produced in the decay, each of them has a continuous kinetic energy spectrum ranging from 0 up to approximately the Q-value of the decay. Their total kinetic energy is almost monoenergetic. However, the antineutrino can essentially never be observed; only the electron will be seen by the detector. If the Q-value is sufficiently high, the electron kinetic energies observed from decays of a particular isotope will be spread

across the entire energy range where one hopes to see solar neutrinos. It is ironic that this most troublesome type of background in a neutrino detector owes its intractability to the near-invisibility of neutrinos!

Often the end product of the β decay is an excited state of the final nuclide. It will emit one or more γ rays to reach the ground state. As with α decay followed by an internal transition, the mean lifetime of the excited state is too short to distinguish the β and γ events.

A positron (also known as β^+) may be emitted instead of an electron, with conversion of one proton in the nucleus into a neutron:

$${}^{A}_{Z}X \rightarrow {}^{A}_{Z-1}Y + e^{+} + \nu_{e} \tag{2.7}$$

This process is very similar to "normal" β decay, but it cannot occur unless the Q-value of the decay would be greater than 1.022 MeV. This accounts for the γ rays that are produced by electron-positron annihilation immediately afterward. The maximum possible kinetic energy of the positron is thus about one MeV less than the Q-value.

The energy spectrum of the electron or positron produced in a typical β decay may be approximated by a phase space argument, and hence is called the statistical spectrum. The argument yields a formula for the rate of events as a function of the β particle momentum p. The electron or positron produced is usually relativistic. Its total energy is $U = \gamma m_e c^2$, and its kinetic energy is by definition $E \equiv U - m_e c^2 = (\gamma - 1)m_e c^2$, giving $p^2 c^2 \equiv U^2 - m_e^2 c^4 = E^2 + 2Em_e c^2$. The rate of events as a function of β kinetic energy E becomes

$$\frac{\mathrm{d}N}{\mathrm{d}E} = C\sqrt{E^2 + 2Em_ec^2} \left(E + m_ec^2\right) (E_{\mathrm{max}} - E)^2 F(Z', E).$$
(2.8)

Here, C is a normalization constant, Z' is the atomic number of the decay product (Z + 1)for electron emission and Z - 1 for positron emission), and E_{max} equals Q for an electron or $Q - 2m_ec^2$ for a positron. F is the Fermi function of atomic number and β kinetic energy. The Fermi function, a correction factor that accounts for Coulomb attraction between the



Figure 2.1: The idealized β particle kinetic energy spectra for two hypothetical cases in which the decay product is presumed to be an isotope of lead with an effective nuclear radius R = 9.5 fm. At left, the maximum kinetic energy of the β is supposed to be 100 keV; at right, 1 MeV. The dotted lines are the curves obtained purely from phase space considerations without incorporating the Coulomb correction factor, Equation (2.9). The solid lines show the energy spectra for emitted electrons. At left, the shape is representative of low-energy β emitters such as ²¹⁰Pb; at right, of higher-energy β emitters such as ²¹⁰Pb; at right, of higher-energy β emitters such as ²¹⁰Pb; at right, of the energies of the emitted electrons tend to be lower than those of the positrons due to the differences in the Coulomb force from the nucleus. These curves have been normalized such that the area under each is one, so the vertical axes are labeled in fractional probability per MeV.

nucleus and electron (or repulsion between nucleus and positron), is approximated by [58]

$$F(Z', E) = 2e^{\pi\eta} (s+1) \left[4\rho^2 (E^2 + 2Em_e c^2) \right]^{s-1} \left| \frac{\Gamma(s+i\eta)}{\Gamma(2s+1)} \right|^2,$$
(2.9)

$$\rho \equiv \frac{m_e c^2}{\hbar c} R$$

$$s \equiv \sqrt{1 - (\alpha Z')^2}$$

$$\eta \equiv \pm \alpha Z' \frac{E + m_e c^2}{\sqrt{E^2 + 2Em_e c^2}} \text{ for } e^{(\mp)}.$$

The Γ function is the continuous generalization of the factorial over the complex plane; for integer n, $\Gamma(n) \equiv (n-1)!$ Figure 2.1 shows the energy spectra predicted by Equations (2.8) and (2.9), for hypothetical β decays by electron and positron emission to a product with atomic number Z' = 82 (lead) and effective nuclear radius R = 9.5 fm. Most commonly the β decay of a nuclide with specific spin and parity J^{π} will yield another nuclide with the same parity, whose spin differs by 0 or ± 1 . This is an "allowed decay." The reason for this is the requirement to conserve angular momentum. Suppose for the moment that the total orbital angular momentum of the three products (the two leptons and the daughter nucleus) is zero, as is the most probable situation. The electron and antineutrino each have a spin of $\frac{1}{2}$, and therefore have a combined angular momentum of either 0 or 1. In the former case, a "Fermi transition," the parent and daughter nucleus must have the same spin. In the latter case, a "Gamow-Teller transition," the spin of the daughter may differ from that of the parent by $\Delta J = \pm 1$; or, if the spin of the parent is not zero, may be the same as that of the parent. Note that if the parent and daughter have the same spin, and that spin is non-zero, the decay may be a linear combination of Fermi and Gamow-Teller transitions.

For some isotopes, the nuclear structure of the possible β decay products is such that neither a Fermi nor Gamow-Teller transition may occur. In this case the parent nucleus, in order to decay, must emit an electron and antineutrino pair with a non-zero total orbital angular momentum. This situation is kinematically disfavored; such decays are suppressed, usually by at least two orders of magnitude, relative to allowed decays. The transition is called n^{th} -order forbidden if $\Delta J = \pm n$ or $\pm (n+1)$ and the product of parent and daughter nuclide parities is $\pi_i \pi_f = (-1)^n$. The possible second-forbidden transitions, for instance, have $\Delta J = \pm 2, \pm 3$ with no parity change. First-forbidden transitions (which do exhibit a parity change), in addition to the expected spin changes $\Delta J = \pm 1, \pm 2$, may also have no spin change, $\Delta J = 0$.

The difference in the energy spectrum for an n^{th} -order forbidden transition with respect to an allowed decay is given by a shape factor $S_n(Z', E)$. In general this factor depends upon ratios of nuclear matrix elements as well as kinematic variables. But for n^{th} -order forbidden decays with $\Delta J = \pm (n + 1)$, called "unique decays," only one matrix element is involved. It can in this case be absorbed into the overall normalization coefficient. For example, the



Figure 2.2: The idealized β particle kinetic energy spectra for two hypothetical cases in which the decay product is presumed to be an isotope of lead with an effective nuclear radius R = 9.5 fm. At left, the maximum kinetic energy of the β is supposed to be 100 keV; at right, 1 MeV. The solid lines show the energy spectra for an allowed β^- decay; the dashed lines show the spectra for a unique first-forbidden β^- decay, Equation (2.10). These curves have been normalized such that the area under each is one.

shape factor for the unique first-forbidden transition $(\Delta J = \pm 2; \pi_i \pi_f = -1)$ is [59]

$$S_{1T}(Z',E) = \frac{(s+1)}{24m_e^2 c^4} \left[(E_{\max} - E)^2 + A(Z',E) \left(E^2 + 2Em_e c^2 \right) \right]$$
(2.10)

where

$$A(Z',E) \equiv \frac{s_1+2}{2s+2} \left(\frac{12\Gamma(2s+1)}{\Gamma(2s_1+1)}\right)^2 \left[4\rho^2(E^2+2Em_ec^2)\right]^{s_1-s-1} \left|\frac{\Gamma(s_1+i\eta)}{\Gamma(s+i\eta)}\right|^2 s_1 \equiv \sqrt{4-(\alpha Z')^2}$$

and the symbols s, η, ρ have the same meanings as in Equation (2.9). The spectra for hypothetical 1-MeV and 100-keV unique first-forbidden decays are compared to spectra for allowed transitions in Figure 2.2. It is also worth noting that the spectra of *non*unique first-forbidden β decays are usually closely approximated by the statistical shape of Equation (2.8).
Electron capture

It may be energetically favorable for an isotope to convert one of its protons to a neutron, but with a Q-value of less than 1.022 MeV. In this case β^+ decay is not possible. The nucleus may instead capture an electron from an inner orbital of its electron cloud:

$${}^{A}_{Z}X + e^{-} \rightarrow {}^{A}_{Z-1}Y + \nu_{e} \tag{2.11}$$

A decay of this type is called electron capture (symbolized ε). It may occur even when the Q-value is greater than 1.022 MeV, in competition with positron emission.

Unless the daughter isotope is produced in an excited state, electron capture is invisible to scintillation detectors because nearly all of the energy released is carried away by the neutrino. The slight kinetic energy imparted to the daughter nucleus, the nuclear recoil, is on the order of $1-10 \,\text{eV}$, far below the thresholds of Borexino and similar detectors.

γ emission

Any of the decay modes $\alpha, \beta^{\pm}, \varepsilon$ may produce a daughter isotope in an excited state. These excited states are called isomers. They usually decays to the ground state by emitting one or more γ rays. Therefore all energy released by a decay, except that bound up in neutrinos, may in principle be observed by the detector. Typically the half-life of the isomer is less than 1 ns, so the γ rays are inextricable from the parent isotope decay in the observed scintillation light. In some cases the parent isotope may decay to several different excited states of the daughter, as with the decay of ²¹⁴Bi, resulting in a very complex energy spectrum. Each component of the spectrum consists of a pure β decay spectrum, shifted to higher energies by the energy of the associated γ rays.

Because γ rays may travel tens of cm through the Borexino scintillator, isotopes whose decay results in the production of high-energy γ rays are very problematic for the detector. The higher in energy the γ rays are, the farther they are likely to travel. Some isotopes

Species		Decay	Q value	Half-life	Source	Method of exclusion
		mode	[MeV]	$ au_{1/2}$		
μ^{\pm}		various	large	-	Cosmic rays	Inner and Outer MVS
n		$p(n,\gamma)d$	2.223	$178\mu{ m s}$	Cosmogenic	MVS and time cut
$^{7}\mathrm{Be}$	(10%)	$\varepsilon + \gamma$	0.478	$53.3\mathrm{d}$	Cosmogenic	
$^{11}\mathrm{C}$		β^+	1.982	$20.4\mathrm{m}$	Cosmogenic	MVS, time/spatial cuts
$^{14}\mathrm{C}$		β^{-}	0.156	$5730{ m yr}$	Scintillator	Energy cut
$^{39}\mathrm{Ar}$		β^{-}	0.565	$269{ m yr}$	Air	
4012	(89%)	β^{-}	1.312	1.00 (7	Durt	
'n	(11%)	$\varepsilon + \gamma$	1.461	1.28 Gyr	Dust	
$^{85}\mathrm{Kr}$		β^{-}	0.687	$10.7\mathrm{yr}$	Air	
$^{87}\mathrm{Rb}$		β^{-}	0.273	$47.5\mathrm{Gyr}$	Dust	

Table 2.2: Summary table of potentially problematic lower-mass radioactive species in Borexino. The Q value for ¹¹C includes the electron/positron annihilation energy following β^+ decay. For decays by electron capture, only the energy of any emitted γ ray is given. MVS stands for Muon Veto System.

may produce γ rays that travel all the way into the Fiducial Volume from the scintillator containment vessels and even from the photomultiplier tubes (a distance of over 3 m). Most of these γ rays are already reduced in energy when they enter the scintillator; they are observed as a continuous energy spectrum whose details may be calculated only by Monte Carlo simulations. Furthermore, because γ rays deposit energy in several different places within the scintillator, the reconstructed position of a γ -ray-producing event will be much less accurate than that of a pure α or β decay.

In a few cases an isomer lasts significantly more than a nanosecond. It may then be designated by an "m" attached to the atomic number in the isotope symbol. For instance, the β -decaying isotope ⁸⁵Kr decays into the excited state ^{85m}Rb with 0.4% probability; ^{85m}Rb, with a half-life of 1.01 μ s, then emits a 514 keV γ ray to reach the ground state of ⁸⁵Rb.

2.2.2 The isotopes present in Borexino

The isotopes with the potential to cause problems for Borexino data analysis may be broken down into several categories. There are naturally occurring, long-lasting radioactive isotopes such as ⁴⁰K and ⁸⁷Rb. Less long-lived isotopes may still be found because they are continually recreated in matter by the impact of cosmic rays, as with ¹⁴C, or byproducts of human activities, such as ⁸⁵Kr. Cosmic rays may penetrate to the detector itself during operation, leaving behind short-lived cosmogenic isotopes such as ¹¹C. Finally, three longlived naturally occurring isotopes each form the head of a long and complex decay chain involving over a dozen species.

The special problem of radiocarbon

The β^- -decaying isotope ¹⁴C is by far the largest source of background in a detector based on organic scintillator. This isotope is cosmogenic. At the surface of Earth, the ¹⁴C concentration in air is continually renewed by the interaction of cosmic rays with carbon dioxide gas in the upper atmosphere. Since the half-life is relatively long—5730 years—the concentration that builds up is not negligible. Indeed, it forms the basis of radiocarbon dating. During life, all organisms respirate, metabolizing CO₂ into their genetic and structural makeup. At death, respiration ceases, so the amount of ¹⁴C left in the remains follows an exponential decay curve. It is therefore possible (with some difficulty, due to long-term variations in atmospheric ¹⁴C levels) to determine the ages of materials derived from living things (bone, ivory, wood, fabric, etc.) as long as they are not too recent or too old.

This panacea for archaeologists presents a grave problem to low-background physics experiments, such as Borexino, that are based on an organic scintillator. Unlike most contaminants, ¹⁴C cannot be removed by purification as it is chemically indistinguishable from stable carbon, and present in molecules of the scintillator itself. (The other such radioisotope one might conceivably worry about, ³H or tritium, decays by β^- emission with a Q-value of 19 keV, which is at the detector's low-energy threshold.) In natural air, the mass ratio of ¹⁴C to ¹²C is on the order of 10^{-12} g/g. If the Borexino scintillator contained ¹⁴C at this isotopic ratio, it would exhibit a ¹⁴C activity of 150 Bq per kg of scintillator.

Fortunately, the Borexino scintillator (like most commercially manufactured organic materials) is derived from petroleum. Petroleum products have spent millions of years deep underground, protected from cosmic rays; the ¹⁴C isotopic abundance measured in the scintillator was only $(5\pm 2) \times 10^{-18}$ g/g [44]. Nevertheless, this implies an activity of 0.75 ± 0.3 mBq/kg of scintillator, for a total expected ¹⁴C event rate in the Fiducial Volume of 75 ± 30 Bq, or 6.4 million events/day. It will be impossible to separate any neutrino events below or (due to finite energy resolution) within 50–100 keV of the end-point from this enormous background. The Q-value for the ¹⁴C β -decay is 156 keV. Taking into account the finite energy resolution of the detector as in Section 2.1, we predict that the rate of ¹⁴C events in the Fiducial Volume with an observed energy > 250 keV will be 0.35 events/day. If the energy threshold is lowered to 240 keV, this rate grows to 1.9 events/day, and at 230 keV, to 8.8 events/day. See Figure 2.3.

A second concern is the fact that if two ¹⁴C decays occur in close proximity in time, they may look like a single higher-energy event ("pile-up"). For this to happen, they must occur within an interval short enough that the event pulse shapes summed over PMTs (as functions of time) overlap. The scintillator volume has a diameter of 8.5 m and therefore it takes a scintillation photon (in a fluid with a refractive index of 1.5) about 43 ns to cross it. Hence light from an event at the edge of scintillator will reach the most distant PMT 43 ns after it reaches the nearest. Also taking into account the time distribution of the scintillator response function (refer to Section 7.5) and the additional delays introduced by scattering effects, we may estimate the maximum pulse duration to be no more than 80 ns. The expected rate of pairs of ¹⁴C events within the Fiducial Volume that occur within 80 ns is 38 per day; for two ¹⁴C events anywhere in the scintillator, it is 344 per day. (A factor of 3 increase in mass implies the same increase in individual event rate, for a factor of 3^2 increase in pile-up event rate).



Figure 2.3: The total rate of ¹⁴C events in the Borexino Fiducial Volume, in decays/day (dpd), as a function of the lower limit of the energy range under study. (That is, the function shown represents the integral of the ¹⁴C energy spectrum above each given energy.) The solid line represents individual ¹⁴C events, while the dashed line represents two ¹⁴C events that occur sufficiently close together in time to be detected as a single higher-energy event, a phenomenon called "pile-up." The dotted line is their sum. This figure incorporates the finite energy resolution of the detector; the actual ¹⁴C spectrum end-point is 156 keV, off the left edge of the graph.

The reconstructed position estimates that will be obtained for such double events are not meaningful. If the reconstructed position is completely random, it statistically should have a probability of roughly $\frac{1}{3}$ of lying within the Fiducial Volume, giving about 120 pile-up events/day that pass the volume cut. However, if it is approximately given by the average of the positions of the two individual decays, it will have a probability more like 80% of being inside the Fiducial Volume, for about 275 pile-up events/day that pass the volume cut. To determine which value is more correct would likely require a full Monte Carlo simulation of the detector. To be conservative, we use the higher figure here. Nevertheless, it is possible that a large fraction of pile-up events may be rejected by some method, for instance, with a cut on the goodness of fit of the reconstructed event position.

The fraction f of pile-up events that will appear to have a total kinetic energy greater than a given value \mathcal{E} is

$$f = \int_{E=0}^{\infty} P(E_1 \in [E, E + dE]) P(E_2 > \mathcal{E} - E_1)$$

=
$$\int_0^{\infty} dE \,\rho(E) \int_{\mathcal{E} - E}^{\infty} dE' \,\rho(E'), \qquad (2.12)$$

where $\rho(E)$ is the observed kinetic energy spectrum of the emitted electron in Borexino. To be precise, $\rho(E)$ is not the same as the ideal spectrum; it also takes into account the finite energy resolution of the detector and the reduced detector efficiency for observing very low-energy events (below about 50 keV). However a conservative estimate may be obtained by ignoring the detector efficiency at low energies, and using for $\rho(E)$ the function of Equation (2.8) convoluted with the detector energy resolution as in Equation (2.3). Doing so, and integrating numerically, gives f = 0.13 when \mathcal{E} is set to the ¹⁴C end-point of 156 keV. Setting \mathcal{E} to the lower end of the neutrino window, 250 keV, yields f = 0.004. Therefore the rate of these overlapping events within the neutrino window should not be greater than $275 \times 0.004 \approx 1.1$ per day in the Borexino Fiducial Volume, weighted heavily toward the lower energy bound. It should be emphasized, though, that this value is proportional to the square of the ¹⁴C contamination. If it is only a factor of three worse in Borexino than the value measured in the Counting Test Facility, the lower neutrino window end-point would have to be increased by as much as 40 keV.

Other light cosmogenic isotopes

The observed event rate in Borexino due to the isotope ¹⁴C will be essentially constant over time due to its long half-life. However, because of the depth of the experiment underground, ¹⁴C will not be regenerated at a rate comparable to its decay. The same cannot be said for some other, shorter-lived, cosmogenic isotopes. Those that are most potentially problematic in Borexino include ⁷Be and ¹¹C. Other isotopes produced by muon interactions with the organic scintillator fluid all have half-lives of at most 15 s, and Q-values greater than 10 MeV. Muons can be detected with 99.98% efficiency by the Borexino muon veto system (Section 3.3.2). When a muon is observed to pass through the detector, events following within a few seconds may be excluded from the data sample. Furthermore, most decays of these other isotopes will have an observed β energy far greater than the upper limit of the interesting energy range.

The isotope ⁷Be has a half-life of 53.3 d, so cannot be excluded by a muon coincidence cut. It decays by electron capture, but 10.4% of the time, it reaches an excited state of ⁷Li which emits a monoenergetic 478 keV γ ray, right in the middle of the neutrino energy window. The predicted σ_E at this energy is 46 keV, so it could potentially obscure a region of spectrum about 100 keV wide in the neutrino energy window. However, the expected event rate is only 0.4 events/day within the Fiducial Volume [42]. It should therefore not pose a big problem.

The cosmogenic ¹¹C has a half-life of 20.4 min. It decays by β^+ emission, with a Q-value of 1.982 MeV. Hence the observed energy of the decay is always between 1.022 and 1.982 MeV. Though no problem for observing ⁷Be neutrinos, this range overlaps an interesting energy region at 0.8–1.3 MeV where the *pep* solar neutrinos could potentially be observed. Furthermore, ¹¹C will be produced within the Fiducial Volume at a rate of 15 ± 2 events/day; 35% of these decays will fall into the *pep* energy range [42]. Considerable effort has gone into the study of the production of this isotope, therefore [22, 40, 41, 42].

Through detailed simulations, it was determined that in 95% of cases, the cosmogenic production of a ¹¹C atom is accompanied by emission of a neutron. (The other 5% are "invisible channels" through which the ¹¹C production cannot be readily detected.) The neutron is easily seen when it is captured by a hydrogen atom, yielding deuterium with production of a 2.2 MeV monoenergetic γ ray. The time scale for this process is $\tau =$ $257 \pm 27 \,\mu$ s [41]. A spherical volume, centered about the reconstructed position of the neutron capture, is defined, and this volume is monitored for likely ¹¹C decays for several half-lives. The ¹¹C rate may be reduced to about 20% of its original value while retaining 93% of the *pep* neutrino data sample simply by excluding these "dead" volumes surrounding each neutron capture event for several ¹¹C mean lives [40]. This gives a signal-to-background ratio for *pep* neutrinos compared to ¹¹C events of S/B = 0.9 [22]. The optimum values of the cuts in this case are r < 76 cm and t < 108 min [22]. Of course, the ¹¹C rate may be reduced further (asymptotically to 5%) at the cost of losing more neutrino data sample by increasing the cut radius or wait time. If the ¹¹C decay can be positively identified within the dead volume, or if the dead volume can be reduced by also reconstructing the path of the progenitor muon, the neutrino sample loss may be mitigated.

Radioactive noble gases

Three noble gas isotopes, present in all unpurified air, are of particular concern in Borexino. These are ³⁹Ar, ⁸⁵Kr, and ²²²Rn. The first two of these are β^- emitters; they produce decays which are essentially indistinguishable from neutrino scattering events. The last, ²²²Rn, decays by α emission, so it can be identified with α/β discrimination techniques, but it produces a long series of even more radioactive daughters. ²²²Rn and its daughters will be described further in a later section.

The isotope ³⁹Ar is produced, like ¹⁴C, by cosmic rays interacting with stable isotopes in the Earth's atmosphere. It has a half-life of 269 yr, and is present in normal air at a concentration of 13 mBq/m^3 at STP [60]. Purified argon gas (present in air at a volume fraction of 0.93%) therefore has an intrinsic activity of 1.4 Bq/m^3 . It is a β^- emitter, with a Q-value of 565 keV, and decays by a unique first-forbidden transition directly to the ³⁹K ground state (no accompanying γ rays). There is no way to distinguish decays of the isotope from neutrino scattering events in Borexino, so it is a very dangerous contaminant.

The requirement set for ³⁹Ar activity in Borexino is no more than one event per day $(0.1 \,\mu\text{Bq/m}^3)$ in the Fiducial Volume. To reach this goal, the Borexino nylon vessels are in the process of being purged repeatedly, first with high-purity nitrogen, and eventually

with special Low Argon-Krypton (LAK) nitrogen. Given equal volumes of nitrogen gas and pseudocumene liquid, argon will partition itself between the two in a ratio of 4.1:1 [43]. Therefore, the volume fraction of argon in the LAK nitrogen must be no greater than 3×10^{-7} , a reduction factor of 3×10^4 from that in normal air.

The isotope ⁸⁵Kr, on the other hand, is anthropogenic; it is released from nuclear fuel reprocessing plants. The half-life is 10.8 yr, and the Q-value is 687 keV. Like ³⁹Ar, it decays by β^- emission to the ground state of the daughter isotope (in this case, ⁸⁵Rb) via a unique first-forbidden transition. Its decays are similarly indistinguishable from neutrino events, and the spectra of the two isotopes within the neutrino window look quite similar.

The activity of air due to 85 Kr is currently ~ 1 Bq/m^3 , and slowly increasing. Hence the activity of pure krypton gas (found in air at volume fractions near 10^{-6}) is on the order of 10^6 Bq/m^3 . With a partitioning ratio for krypton between nitrogen and air of 1.3:1, the volume fraction of krypton in Borexino's LAK nitrogen must be less than about 10^{-13} . This is a factor of 10^7 less than in air.

There is one method by which the rate of ⁸⁵Kr decays may in principle be directly measured. With 0.43% probability, an atom of ⁸⁵Kr may decay into an excited state of ⁸⁵Rb. This excited state, known as ^{85m}Rb, has a half-life of 1.01 μ s before emitting a monoenergetic 514 keV γ ray to reach the nuclear ground state. This short-lived excited state makes it possible to look for "coincidences" in which two events that occur within a few μ s have the correct energies ($E_1 < 173$ keV, $E_2 \approx 514 \pm 50$ keV). The rate of such coincidences can be extrapolated to estimate the rate of all ⁸⁵Kr decays in the Fiducial Volume.

It should be noted, though, that if the ⁸⁵Kr activity meets the goal of one event per day in the Fiducial Volume, the number of these coincidences observed throughout the scintillator during the ten-year lifetime of the Borexino experiment will be only 48 ± 7 . The method is not very sensitive at all. Because the first event of the coincidence falls into the energy range dominated by ¹⁴C, it will also be susceptible to false positives. For instance, the number of accidental coincidences caused by a ¹⁴C event followed within 3 μ s by a neutrino event within 1σ of 514 keV will be about 54 ± 7 over ten years. (This assumes a neutrino event rate of 22 events/day in the range 514 ± 50 keV over the entire volume of scintillator.)

The Borexino experiment has found commercial suppliers of liquid nitrogen that improves upon the required argon contamination by factors ranging from 10–600 [43], and upon the required krypton contamination by up to a factor of three [42]. It is thought, therefore, that these two noble gas isotopes represent a solved problem. Great care is required nonetheless. A 30 cm³ bubble of outside air that somehow entered the scintillator would exceed the requirement for krypton all by itself; an air bubble ten times larger (still only a third of a liter) would seriously damage the sensitivity of the experiment.

Medium-weight long-lived isotopes

The category of long-lived isotopes includes those that are no longer being generated on Earth by any important mechanism, but which have survived since the origin of the solar system due to their exceedingly long half-lives. Here we mention only those that decay immediately to stable products. The heavy-element decay chains will be described later. As metals, these isotopes may be found in dust particles, or simply be present as dilute solutions within the scintillator.

One isotope in this family is ⁸⁷Rb, a pure β^- emitter that decays via a non-unique thirdforbidden transition. The abundance of rubidium in the Earth's crust is roughly 78 ppm by weight [61], and the isotopic abundance of ⁸⁷Rb is 28%. Combined with the isotope's half-life of 47 billion years, its expected activity in typical rock dust is about 65 mBq/kg. However, the Q-value of the decay is only 273 keV. The majority of decays will yield a β energy well below the neutrino energy window starting at 250 keV, therefore, and this isotope is not considered much of a concern.

More problematic is ⁴⁰K. This isotope decays by pure β^- emission (in a unique thirdforbidden transition) with 89.3% probability. The abundance of potassium in the upper crust of the Earth is 2.7% by weight [62], with an isotopic abundance for 40 K of 117 ppm. The half-life is 1.28 billion years, implying a total activity in rocky material of 810 mBq/kg, more than ten times that of 87 Rb. Furthermore, its β spectrum endpoint is 1.312 MeV, so it produces events throughout the neutrino energy range.

It will be possible, and fairly easy, to get a good measurement of the amount of 40 K in the Borexino scintillator. This is because the other decay branch (10.7% probability) is an electron capture yielding 40 Ar. (Though positron emission is energetically possible, it occurs very rarely, with a branching ratio of 0.001%.) Nearly all of the time (98%), the 40 Ar nucleus is produced in an excited state, and immediately emits a monoenergetic 1.46 MeV γ ray that is easily visible in the energy spectrum. (This measurement is performed for the prototype Counting Test Facility in Sections 8.5.2 and 9.4 of this work, with inconclusive but somewhat worrying results.) The presence of this γ ray also, however, means that potassium in structural components of the detector, external to the scintillator, may produce an additional γ -ray background within the Fiducial Volume.

The maximum concentration of potassium (all isotopes) considered acceptable in the Borexino scintillator is on the order of 10^{-14} grams per gram of scintillator. This contamination level yields 1.5 events/day in the Borexino Fiducial Volume at ⁷Be neutrino energies, 250– 800 keV [12].

Two β -emitting rare earths, ¹³⁸La and ¹⁷⁶Lu, also fall into the category of long-lived natural radioisotopes. Each of these has an expected activity in rock on the order of 30 μ Bq/kg, though, so their contributions to radioactive background in Borexino should be negligible. The radioactive heavy metals (uranium and thorium) present many more difficulties.

2.2.3 The heavy element decay chains

There are three naturally occurring heavy decay chains. Within each chain, every isotope has the same atomic number, modulo four (both α and β decay preserve this modulus).



Figure 2.4: Pictorial representations of the ²³⁸U and ²³²Th decay chains. Secular equilibrium is likely to hold only within the sets of isotopes grouped by dotted lines. Energies shown are Q values for β emitters, and α kinetic energy for α emitters. Times shown are half-lives. Isotopes shaded blue are β emitters with a spectrum endpoint above the 250 keV lower limit of the neutrino energy window. They are most problematic. Isotopes shaded yellow are α emitters. Due to α quenching in the scintillator (Section 3.1.2), their decays cause scintillation events that appear to have energies in the neutrino window. However they may be excluded from the neutrino data sample via pulse shape discrimination with an efficiency of about 95%. ²¹²Bi is shaded green as it decays both by α and by β emission.

Like 40 K, the parent isotope of each chain has a sufficiently long lifetime that some amount remains from the origin of the solar system. The uranium series, whose parent is 238 U, is the A = 4n + 2 chain; the thorium series, whose parent is 232 Th, is the 4n chain; the actinium series, whose parent is 235 U, is the 4n + 3 chain. The 4n + 1 chain, or neptunium series, does not include any progenitor long-lived enough to survive until the present day. Even isotopes in the actinium series are rare enough that they present no problem to Borexino (but see Section 8.2.3). The uranium and thorium series are depicted in Figure 2.4.

Many nuclides in these decay chains decay by α emission at high energies. Due to the phenomenon of α quenching, unfortunately, nearly all of these will be observed to have β equivalent energies within the neutrino energy window. But most (90–95%) will be possible to exclude from the data sample by means of the α/β discrimination techniques discussed in Section 8.1.3.

The uranium series

The uranium series isotopes are as follows. Refer to Figure 2.4 or Table 2.3 for their decay energies and half-lives.

- 238 U, an α -emitter, is the progenitor of the chain. It is present in typical rock dust at 2.5 ppm by weight, resulting in an activity of $30 \,\mathrm{mBq/kg}$ [62]. Although less than that of 40 K by a factor of nearly 30, this activity must be multiplied by the number of daughter isotopes (14) when 238 U is in secular equilibrium with them.
- 234 **Th** is not a problem due to its low Q-value (199 keV). It always decays to an isomer called 234m Pa of its daughter isotope.
- 234m **Pa** is peculiar in that it nearly always (99.87%) decays by β emission rather than relaxing to the isotopic ground state. The decay mode is pure β^- with 2.29 MeV end-point. Hence many of its decays are observed within the neutrino energy window.

Species	Decay	Q value	E_{α} or $E_{\beta+\nu}$	E_{γ}	Branching	Half-life		
	mode	[MeV]	[MeV]	[MeV]	ratio $[\%]$	$ au_{1/2}$		
$^{238}\mathrm{U}$	α	4.270	4.196	-	77	$4.47\mathrm{Gyr}$		
234 Th	β^{-}	0.199	0.199	-	72	$24.1\mathrm{d}$		
$^{234\mathrm{m}}\mathrm{Pa}$	β^{-}	2.29	2.29	-	98	$1.18\mathrm{m}$		
$^{234}\mathrm{U}$	α	4.856	4.774	-	72	$245\mathrm{kyr}$		
230 Th	α	4.771	4.687	-	76.3	$80\mathrm{kyr}$		
226 Ra	α	4.871	4.785	-	94.5	$1600{ m yr}$		
Radon 222 and daughters:								
222 Rn	α	5.591	5.490	-	99.9	$3.82\mathrm{d}$		
218 Po	α	6.115	6.002	-	~ 100	$3.05\mathrm{m}$		
$^{214}\mathrm{Pb}$	β^{-}	1.024	0.672	0.352	48	$26.8\mathrm{m}$		
^{214}Bi	β^{-}	3.270	1.51	1.76	40	$19.7\mathrm{m}$		
214 Po	α	7.833	7.687	-	~ 100	$164\mu{ m s}$		
$^{210}\mathrm{Pb}$	β^-	0.063	0.016	0.047	81	$22.3\mathrm{yr}$		
$^{210}\mathrm{Bi}$	β^{-}	1.161	1.161	-	~ 100	$5.01\mathrm{d}$		
210 Po	α	5.408	5.305	-	~ 100	$138.4\mathrm{d}$		
$^{206}\mathrm{Pb}$	stable							

Table 2.3: Summary table of the Uranium-238 decay chain. The energy carried by the α particle or β and neutrino, and the total energy released in γ rays, respectively, are shown for the decay mode with the greatest branching ratio.

- $^{234}\mathbf{U}$ is another α emitter, as are daughters $^{230}\mathbf{Th}$ and $^{226}\mathbf{Ra}.$
- ²²²**Rn**, another α emitter, is a special isotope. As a noble gas, atoms of this element are seldom found ionized; they are free to diffuse through most materials. It may travel long distances during its lifetime ($\tau_{1/2} = 3.8$ days). When it does, this breaks secular equilibrium in the decay chain. The concentration of radon in a volume is no indication that it contains an equal activity of the higher isotopes. In Borexino, in particular, radon produced by the decay of radium in the nylon vessels, steel support structures, photomultipliers, or even steel sphere may make its way into the scintillator volume. Delaying this travel of radon long enough for it to decay before reaching the scintillator is in fact the primary purpose of the nylon vessels.
- ²¹⁸**Po** is the last α emitter in the row. As it has a half-life of only 3.1 minutes, it may well be possible to tag the successive decays of radon, this isotope, and the following

three. Their half-lives are short enough to guarantee that all are in secular equilibrium together. Section 8.4 provides a detailed discussion of this possibility.

- ²¹⁴**Pb** is an undistinguished β^- emitter which frequently produces γ rays during its decay. The most common are emitted at 351 and 295 keV; γ energies up to 839 keV are (rarely) possible.
- ²¹⁴**Bi** (β^-) and ²¹⁴**Po** (α) are usually easily detected due to the very short lifetime of the latter isotope. However, atoms of ²¹⁴Bi located outside the scintillator may produce high-energy γ rays that travel into the scintillator or even Fiducial Volume. The most common have energies of 609 keV, 1.76 MeV and 1.12 MeV. Energies up to 2.15 MeV may occur.
- ²¹⁰Pb is essentially undetectable due to its low Q-value of 63 keV. However, its long halflife (22 years) gives it ample time to travel into the detector and produce progeny. Some studies have been performed (refer to Section 4.4.4, for instance) on its affinity for sticking to various surfaces and then washing off into scintillator that passes over those surfaces.
- ²¹⁰**Bi** is a pure β^- emitter with a Q-value of 1.16 MeV, putting many of its decays in the neutrino energy window.
- ²¹⁰**Po** is an α emitter and the final radioisotope in the uranium series. Like ²¹⁰Pb, it has some affinity for adhering to surfaces.

To summarize, the major concerns in this decay chain with respect to scintillator contaminants are the pure β^- emitters ^{234m}Pa and ²¹⁰Bi. Although each isotope will be in secular equilibrium with its parent, the parent is either indistinguishable from ¹⁴C background (in the former case) or nearly invisible to Borexino (in the latter case). The α emitter ²¹⁰Po is also a concern, just because it may be a large background as a result of "wash-off" of ²¹⁰Pb and ²¹⁰Po atoms from metal surfaces over which the scintillator passes. Finally, atoms of

Species	Decay	Q value	E_{α} or $E_{\beta+\nu}$	E_{γ}	Branching	Half-life
	mode	[MeV]	[MeV]	[MeV]	ratio [%]	$ au_{1/2}$
232 Th	α	4.081	4.011	-	77	$14.1\mathrm{Gyr}$
228 Ra	β^{-}	0.046	0.039	0.007	60	$5.76\mathrm{yr}$
^{228}Ac	β^{-}	2.137	1.11	1.03	53	$6.13\mathrm{h}$
228 Th	α	5.520	5.423	-	72.7	$1.91{ m yr}$
224 Ra	α	5.789	5.686	-	95.1	$3.66\mathrm{d}$
220 Rn	α	6.405	6.288	-	99.9	$55.6\mathrm{s}$
216 Po	α	6.907	6.779	-	~ 100	$145\mathrm{ms}$
$^{212}\mathrm{Pb}$	β^{-}	0.573	0.334	0.239	83	$10.6\mathrm{h}$
²¹² Bi (64%)	β^{-}	2.246	2.246	-	55	$60.6\mathrm{m}$
²¹² Po	α	8.954	8.784	-	~ 100	$299\mathrm{ns}$
²¹² Bi (36%)	α	6.207	6.051	0.040	25	$60.6\mathrm{m}$
$^{208}\mathrm{Tl}$	β^{-}	4.992	1.794	3.198	51	$3.05\mathrm{m}$
208 Pb	stable					

Table 2.4: Summary table of the Thorium-232 decay chain.

²¹⁴Pb and ²¹⁴Bi outside the inner nylon vessel may emit high-energy γ rays that travel all the way into the Fiducial Volume. All of these problems must be guarded against.

The maximum level of uranium contamination considered acceptable in the Borexino scintillator is about 10^{-16} g/g. Under the assumption of secular equilibrium, this contamination level would lead to an event rate of ~100 events/day in the Fiducial Volume and ⁷Be neutrino energy window [12], but most (~90) of these would be easily identified α decays.

The thorium series

Isotopes in the thorium series are described below. Refer to Figure 2.4 or Table 2.4 for their decay energies and half-lives.

 232 **Th**, an α -emitter, is the progenitor of the chain. It is present in typical rock dust at 10.3 ppm by weight, resulting in an activity of 41 mBq/kg [62]. This activity must be multiplied by a factor of ten when accounting for all the other isotopes in the series.

- 228 **Ra** is unimportant in Borexino, with a Q-value of only 46 keV.
- ²²⁸**Ac** emits a β^- with a Q-value of 2.14 MeV. This isotope produces a complex γ spectrum, including γ rays at energies 911, 969, and 338 keV. γ ray levels up to about 2 MeV may occur.
- ²²⁸**Th** and daughters ²²⁴**Ra**, ²²⁰**Rn**, and ²¹⁶**Po** are α emitters. The last three of these may be tagged as a "triple- α coincidence" given the short half-lives of ²²⁰Rn and ²¹⁶Po.
- ²¹²**Pb** is a β^- emitter with an end-point at 573 keV, in the neutrino window. It may produce γ rays, the most common at 239 keV.
- ²¹²**Bi** may decay either by α (36%) or β^- (64%). A β decay is easily seen because it is immediately followed by the α decay of the extremely short-lived ²¹²**Po**. In this case it may conceivably be possible to identify the ²¹²Pb event between this coincidence and the preceding triple- α coincidence.
- ²⁰⁸**Tl** is on the lower-probability decay branch from ²¹²Bi. It has a high-energy β^- decay that always releases a 2.615 MeV γ ray; 99.97% of the time, at least one additional γ ray with a minimum energy of 0.583 MeV is also produced. When it occurs within the scintillator, this high-energy decay and relatively short (3.1 minute) half-life may be sufficient to identify the ²¹²Bi/²⁰⁸Tl pair of decays. However, the penetrating γ rays make it the most dangerous isotope in either decay chain for detector components outside the inner nylon vessel. The exact rate of ²⁰⁸Tl events may be crucial in determining whether or not it will be possible for Borexino to observe *pep* neutrinos.

Within the thorium series, potentially troublesome isotopes within the scintillator volume include the β^- emitters ²²⁸Ac and ²¹²Pb. Most other isotopes in the series will be identifiable either as α decays through α/β discrimination, or through the method of coincidences (in some cases, extended over several minutes). The most difficult challenge, however, will be the penetrating high-energy γ rays produced by ²⁰⁸Tl atoms embedded in the photomultipliers and other pieces of detector hardware. The maximum level of thorium contamination considered acceptable in the Borexino scintillator is about 10^{-16} g/g. Under the assumption of secular equilibrium, this contamination level would lead to an event rate of about 23 events/day in the Fiducial Volume and ⁷Be neutrino energy window [12]. About 19 of these events would be easily detectable α decays.

2.2.4 Categories of radioactive background

The isotopes described in this section, and possibly others that have not yet been considered, contribute to three types of radioactive background in the detector, based upon the location of the decaying nuclides.

Internal background consists of the decays of radioactive atoms that happen to be impurities in the scintillator fluid or that are generated cosmogenically by interactions of muons with the scintillator. These decays will most probably occur with a homogeneous distribution throughout the scintillator volume. Some internal background rates, such as that due to ³⁹Ar, are predictable since the background comes entirely from materials (the LAK N_2) with known concentrations of radioisotopes. Some background rates, such as those of the cosmogenic isotopes, are predictable because they depend only upon measured physical quantities such as muon interaction cross-sections and the muon flux at the depth of Borexino. Other background rates, such as that of ⁴⁰K, are not really predictable. The numbers of atoms of these species found in the scintillator depend upon both the scintillator's initial level of contamination and upon the effectiveness of purification methods at removing contaminants. Neither are really known *a priori*. For this reason, estimating the total expected background rate in Borexino before the detector begins data acquisition would be a fairly meaningless exercise, though the rates resulting from certain individual species may be calculated with reasonable confidence.

Attempts to measure the rates of these isotopes with the 4-ton Borexino prototype, the Counting Test Facility (CTF), have been inconclusive for two reasons. First, the history of the CTF has been complex. Numerous purification tests have been performed on its scintillator, making it unclear whether the presence of certain contaminants at a given moment results from an inadequacy in the most recent purification method, or from their introduction into the scintillator by accident during a previous test. Some contaminants (the mass-210 isotopes, for instance) are known to adhere readily to various nylon and metal surfaces, and later "wash off" back into purified scintillator. Second, as the CTF has a mass of only $\frac{1}{25}$ that of the Borexino Fiducial Volume, the inability to detect a certain isotope in the CTF says nothing about whether it will be a problem for Borexino.

Despite these facts, the CTF is a most invaluable tool. Though it can never conclusively say that a certain isotope will not cause trouble for the Borexino detector, it *can* point out the most problematic areas. It has been used to test materials to be used in Borexino, scintillator purification techniques, and perform the most sensitive measurement ever made of the 14 C levels in petroleum derivates. The CTF is described in great detail in Chapter 6, and the following chapters discuss results for internal contamination that *can* unequivocably be concluded from CTF data.

A second type of background in Borexino may be classified as surface background. This background is caused by radioactive atoms either adsorbed on the inner surface of the nylon vessel that contains the scintillator, or embedded within it. Those atoms sufficiently close to the inner surface of the nylon film can release an α or β particle into the scintillator. Any decay in the nylon may produce a γ ray that travels inward, to be observed as an event. Radium atoms in the nylon will continually decay into radon, which may migrate into the scintillator and itself decay there. This process is called emanation. Finally, radon atoms from outside the inner nylon vessel may even travel all the way through it, eventually to end up in the Fiducial Volume. These possibilities are all discussed in relation to the Borexino nylon vessels in Section 4.4. Investigations of surface background in the CTF are performed in Chapter 9. Finally, one of the most difficult types of background to deal with is the external background caused by high-energy γ rays produced outside the inner nylon vessel, which nevertheless travel all the way into the volume of scintillator. The specific activities for various detector components are known fairly well, for instance in [63]. However, they produce a continuous energy spectrum in the scintillator whose precise shape and amplitude is difficult to predict, even with Monte Carlo methods. The presence of a γ -ray continuum has posed serious difficulties for analyses of internal contaminant concentrations based upon the observed energy spectrum of the CTF. Some discussion on the problem in the CTF is given in Section 9.4 of this work. Monte Carlo simulations of the expected Borexino external background are presented in, for instance, [12, 48, 49].

2.3 Prospects for observing solar neutrinos

The solar neutrino signal that can be measured in Borexino will depend upon the target mass and lifetime of the detector; the flux of neutrinos at Earth; the cross-section for each neutrino flavor to react in a visible way; and the fraction of surviving electron neutrinos. The latter three of these are functions of the incident neutrino energies. We will assume a detector lifetime of 10 years. The flux of neutrinos at Earth is given essentially by their rate of production in the Sun, which was discussed in Section 1.3.

As already mentioned, the means of observing solar neutrinos in Borexino is through elastic scattering on electrons, $\nu + e^- \rightarrow \nu + e^-$. (Neutrino capture by a ¹²C nucleus, ¹²C + $\nu_e \rightarrow$ $^{12}N + e^-$, cannot occur unless the neutrino energy is greater than 17 MeV. The threshold for capture by ¹³C is only 2.2 MeV, but the isotopic abundance of ¹³C is fairly low, 1.1%.) Given a certain incident energy E_{ν} for the neutrino, the kinetic energy spectrum of the electron is uniquely determined by the differential cross-section for the interaction. When the electron is scattered in the direction in which the neutrino was initially traveling, it is imparted the maximum kinetic energy possible for a given neutrino energy. This kinetic



Figure 2.5: First-order Feynman diagrams for neutrino-electron elastic scattering. All neutrinos may scatter on electrons via a neutral-current interaction involving the exchange of a virtual Z^0 (left). Only electron neutrinos may scatter on electrons by the mediation of a charged virtual W particle (right).

energy is given by classical kinematic arguments as

$$E_{\rm max} = \frac{E_{\nu}}{1 + m_e c^2 / (2E_{\nu})}.$$
(2.13)

2.3.1 The ⁷Be neutrinos

The 862 keV ⁷Be solar neutrinos will be observed in Borexino as an electron recoil spectrum that is nearly constant up to an energy of 667 keV, at which point it descends sharply to zero. This feature, the Compton edge, will in practice be smeared by the finite energy resolution of the detector. (Because the electron capture decay of ⁷Be has a 10.4% branching ratio to an excited state of ⁷Li, a fraction of ⁷Be neutrinos are produced with an energy of 384 keV, implying a second Compton edge at 230 keV. This signal will be obscured by ¹⁴C background. Below we consider only the 862 keV ⁷Be neutrinos.)

Below the maximum recoil energy, the differential cross section for a given neutrino energy E_{ν} is given by

$$\frac{d\sigma}{dE}(E;E_{\nu}) = \frac{\sigma_0}{m_e c^2} \left[g_{\ell}^2 + g_r^2 \left(1 - \frac{E}{E_{\nu}} \right)^2 - g_{\ell} g_r \frac{m_e c^2 E}{E_{\nu}^2} \right],$$
(2.14)

where $\sigma_0 \equiv 2G_F^2 m_e^2/(\pi \hbar^4) = 8.81 \times 10^{-45} \text{ cm}^2$. The value $g_r = \sin^2 \theta_w \approx 0.222$ for all neutrinos. The value g_ℓ is $\sin^2 \theta_w + \frac{1}{2} \approx 0.722$ for electron neutrinos, and $\sin^2 \theta_w - \frac{1}{2} \approx -0.278$ for other neutrino flavors. The difference comes from the ability of electron neutrinos

to scatter from electrons by a charged-current reaction, whereas other neutrinos can scatter from electrons only via a neutral current (Figure 2.5). The angle θ_w itself is the weak mixing angle.

The total cross section for a neutrino of energy E_{ν} is simply the integral of this expression from E = 0 up to E_{max} :

$$\sigma(E_{\nu}) = \sigma_0 \frac{E_{\max}}{m_e c^2} \left[(g_{\ell}^2 + g_r^2) - \left(\frac{g_r^2}{E_{\nu}} + g_{\ell} g_r \frac{m_e c^2}{2E_{\nu}^2} \right) E_{\max} + g_r^2 \frac{E_{\max}^2}{3E_{\nu}^2} \right].$$
(2.15)

Suppose that the flavor composition of the neutrino beam is known: a fraction P_e of the beam consists of electron neutrinos. Then we may define effective differential and total cross sections by replacing g_{ℓ} and g_{ℓ}^2 in Equations (2.14) and (2.15) by their average values:

$$\langle g_\ell \rangle = \sin^2 \theta_w - \frac{1}{2} + P_e \approx P_e - 0.278$$
 (2.16)

$$\langle g_{\ell}^2 \rangle = \sin^4 \theta_w - \sin^2 \theta_w + \frac{1}{4} + 2P_e \sin^2 \theta_w \approx 0.444 P_e + 0.0773.$$
 (2.17)

For the 862 keV ⁷Be neutrinos, Equation (2.15) yields $\sigma_e = 0.659 \sigma_0$ for electron neutrinos $(P_e = 1)$, and $\sigma_{\mu,\tau} = 0.147 \sigma_0$ for μ - or τ -neutrinos $(P_e = 0)$. In the limit $E_{\nu} \gg m_e c^2$, the respective results become $\sigma_e \to 0.538 \sigma_0 (E_{\nu}/m_e c^2)$ and $\sigma_{\mu,\tau} \to 0.094 \sigma_0 (E_{\nu}/m_e c^2)$.

With a monoenergetic neutrino beam of energy E_{ν} , the number of interactions observed to yield electron kinetic energies in a specified narrow range [E, E + dE] is given by the energy width dE multiplied by the incident neutrino flux Φ , the differential cross-section $d\sigma/dE$ (weighted appropriately for electron neutrinos and other flavors), number density of electrons in the target n, volume V of the target, and time T for which it is observed. That is:

$$\frac{\mathrm{d}N}{\mathrm{d}E}\,\mathrm{d}E = nVT\,\Phi\,\frac{\mathrm{d}\sigma_{\mathrm{eff}}}{\mathrm{d}E}\,\mathrm{d}E.\tag{2.18}$$

This can be rewritten as $nVT \Phi \sigma_{\text{eff}} \rho_r(E; E_{\nu})$. Here, $\rho_r(E; E_{\nu}) \equiv (1/\sigma_{\text{eff}}) (d\sigma_{\text{eff}}/dE)$ is the normalized (integral = 1) electron kinetic energy spectrum of the interaction.



Figure 2.6: The spectrum of electrons scattered by ⁷Be neutrinos that should be observed in Borexino, given the two extreme cases of (top curve) no neutrino oscillations; (bottom curve) complete oscillation into μ - and τ -neutrinos. The middle curve shows the expected result when $P_e = 0.55$ as predicted for energies where the MSW effect does not contribute much and vacuum oscillations dominate. Values on the *y*-axis are given in units of events per day per 100 keV. Solid lines are the ideal theoretical spectra, with a sharp Compton edge at 667 keV. Dashed lines represent smearing effects due to finite energy resolution. The steeply descending dotted curve that falls to zero near 300 keV indicates the expected ¹⁴C background (both individual and "pile-up" events).

The observed kinetic energy spectrum ρ_d will be the convolution of the true spectrum ρ_r with the energy resolution function, as described previously in Section 2.1. Hence, the total rate of ⁷Be neutrinos observed in the Borexino Fiducial Volume with energies greater than 250 keV will be $nV\sigma_{\text{eff}} \int_{250 \text{ keV}}^{\infty} \rho_d(E; E_{\nu}) dE$. (Below 250 keV, ¹⁴C background will dominate, hiding the neutrino signal.) To simplify the calculation, ρ_d in this integral may be replaced by ρ_r with little error, as the function $d\sigma_{\text{eff}}/dE$ is fairly flat over most of its range (Figure 2.6).

In the case of Borexino, $nV = 3.31 \times 10^{31}$ electrons in the Fiducial Volume. The flux of 862 keV ⁷Be neutrinos at Earth is $89.6\% \times 4.84 \times 10^9 \text{ cm}^{-2}\text{s}^{-1}$ [5, 10]. (The estimated error in this figure, due to uncertainties in the Standard Solar Model, is a relatively high 10.5%.) If no neutrino oscillations occurred, the rate of ⁷Be neutrinos observed in the Fiducial Volume

with E > 250 keV would be 43 per day. If, on the other hand, *all* neutrinos were converted to μ and τ flavors during their travels, the expected rate of ⁷Be neutrinos above 250 keV would become only 10 per day. The expected value of P_e is about 0.55, a result obtained from the prediction (as discussed in the previous chapter) that the MSW effect is small for ⁷Be neutrinos even at the densities found in the Sun's core. For $P_e = 0.55$, the expected ⁷Be neutrino event rate is 28 per day at energies above 250 keV. (Taking into account the *pep* and CNO neutrino signals described below in Section 2.3.3 will add roughly four additional events per day in the 250–800 keV range to this figure.) These numbers indicate that the actual observed neutrino event rate will be a sensitive indicator of the probability of neutrino oscillation.

In the best possible case (no non-removable background), assuming the value of P_e given above, a total of about 10⁵ ⁷Be neutrinos would be observed in events over the 250 keV threshold during ten years of detector operation. The statistical error in the calculated neutrino flux would be 0.3%. Of course, this situation cannot be achieved, and the real question for the detector sensitivity is what level of background noise may be attained. One may consider several types of background. The most harmless is that which may be "tagged" and excluded from the data sample *a priori*, for instance the ²¹⁴BiPo and ²¹²BiPo coincidence events caused by the very short half-lives of those two polonium isotopes. α/β discrimination, removing 90–95% of α events from the data sample, is a second example of event tagging. Of the remaining background, some of it may be assumed to have a known rate because it is in secular equilibrium with taggable events. This is the case, for instance, with ²¹²Pb in the thorium chain. Other backgrounds (⁴⁰K, ⁸⁵Kr, etc.) have imprecisely known rates but accurately known spectral shapes. The most troublesome backgrounds are those which exhibit neither, for instance external γ rays whose spectrum can be at best simulated with Monte Carlo methods.

Consider the case where the background spectrum is known perfectly, so it can be statistically subtracted from the neutrino signal. We suppose that events which can be individually tagged have already been purged from the data. Let the total number of events observed in the neutrino window be N_t , of which N_s are the neutrino signal and N_b are well-understood background. Hence $N_s = N_t - N_b$. Propagating the errors, we find

$$\delta N_{s} = \sqrt{\delta N_{t}^{2} + \delta N_{b}^{2}} = \sqrt{N_{t} + N_{b}} = \sqrt{N_{s} + 2N_{b}}$$

= $\sqrt{N_{s}} \sqrt{1 + 2N_{b}/N_{s}}.$ (2.19)

That is, if the S/N ratio is 10, the statistical error in the result will be at least 10% greater than in the ideal case. If the S/N ratio is two—the known background rate is 50% of the neutrino rate—then the statistical error will be more than 40% greater. Most likely, in either case the systematic error (due to the existence of background with a poorly understood spectral shape) will be larger than the statistical error.

2.3.2 Annual variations in the ⁷Be signal

Unfortunately there is no way to obtain an ideal measurement of the background rate in Borexino. The solar neutrino signal cannot be turned off to disentangle it from the background. It can, however, be varied—slightly. The Earth has an orbit with an eccentricity $\epsilon \approx 0.0167$ and semi-major axis $a \approx 1.496 \times 10^8$ km. The distance between the center and a focus of an ellipse is ϵa . At aphelion, in July, the Earth-Sun distance is 3.4% greater than at perihelion, in January. Flux received from the Sun varies as the inverse square of the distance: the maximum neutrino flux is 6.9% higher than the minimum flux.

This fact will permit us to measure the solar neutrino flux without potential ambiguities due to background with poorly known spectral shapes. Suppose that the background in the neutrino window 250–800 keV has rate B, and that the average neutrino signal has rate S_0 . In January, we have a signal of rate $(1+2\epsilon)S_0$, and in July, neutrino rate $(1-2\epsilon)S_0$. The total signal over n Januaries (each of length T = 31 days) becomes $N_{\text{max}} = nT[B + (1 + 2\epsilon)S_0]$, and over n months of July becomes $N_{\text{min}} = nT[B + (1 - 2\epsilon)S_0]$, with respective statistical errors given by the square roots of those values. The difference is then $N_{\text{max}} - N_{\text{min}} =$



Figure 2.7: Simulated data, showing the accumulated neutrino signals over one-month periods for three years as a function of time. This figure illustrates the expected annual variation in the neutrino flux due to the eccentricity of the Earth's orbit. Though this simulation was performed assuming no oscillations, the actual expected data should be similar, although with only about 65% of the numbers of events shown. Figure taken from reference [64].

 $4 \epsilon nTS_0 \pm \sqrt{2nT(S_0 + B)}$. The statistical error as a fraction of the difference is

$$\frac{\delta(\Delta N)}{\Delta N} \equiv \frac{\delta S_0}{S_0} = \frac{\sqrt{2}}{4\epsilon} \frac{1}{\sqrt{nTS_0}} \sqrt{1 + \frac{B}{S_0}}.$$
(2.20)

Plugging in the numbers, assuming $S_0 = 28$ events/day and n = 10 (years), we obtain $\delta S_0/S_0 \approx 0.23\sqrt{1+B/S_0}$. Though a 20% statistical error is quite large in comparison with the error quoted earlier (due to the small value of ϵ and the much reduced period of data collection), it must be noted that the value obtained in this way will include no systematic error due to background spectral shapes. It may also be reduced further by including other months in the data sample and fitting a periodic function to the results,

$$S(t) = S_0 \left[1 + 2\epsilon \cos\left(\frac{2\pi t}{1 \text{ yr}}\right) + O(\epsilon^2) \right], \qquad (2.21)$$

as shown in Figure 2.7.

Observation of the neutrino flux variation as a function of time is not only useful as a crosscheck for systematic errors caused by background. It may also point to new physics if the observed variation does not match the expected value of ~ 7%. Such unexpected annual variations might indicate that, contrary to theory, the ⁷Be neutrino oscillates in vacuum over length scales on the order of $2 \epsilon a \approx 5 \times 10^6$ km.

2.3.3 The *pep* and CNO cycle neutrinos

Borexino may have the ability to observe the *pep* solar neutrinos, and (depending upon their rate of production in the Sun) also the CNO cycle neutrinos, in the energy range 0.8–1.3 MeV. The *pep* neutrinos are produced at the single energy 1.44 MeV, implying a Compton edge at 1.22 MeV. Their expected flux at Earth is $\Phi_{pep} = 1.42 \times 10^8 \,\mathrm{cm}^{-2} \mathrm{s}^{-1}$, with an uncertainty of only 2% [5, 10].

Unlike the ⁷Be and *pep* neutrinos, those produced in the CNO cycle originate in the β^+ decays of the species ¹³N, ¹⁵O, and ¹⁷F. The CNO neutrinos therefore have continuous spectra that range from zero to the Q-values of these decays (less $2m_ec^2$, to account for positron annihilation). Their spectra are given by Equation (2.8) with the substitution $E \rightarrow Q - 2m_ec^2 - E_{\nu}$; that is, they are mirror images of the positron kinetic energy spectra.

The respective end points of the ¹³N, ¹⁵O, and ¹⁷F neutrino spectra are at 1.20, 1.73, and 1.74 MeV, giving end points for the recoil electron spectra at 0.99 and 1.51 MeV. The predicted fluxes of these species in the 2005 SSM are roughly $\Phi_{13N} = 3.05 \times 10^8$, $\Phi_{15O} =$ 2.31×10^8 , and $\Phi_{17F} = 5.83 \times 10^6$ cm⁻²s⁻¹ [5, 10]. (Hence the ¹⁷F neutrino will be essentially impossible to observe separately from the ¹⁵O neutrino; it will have the effect of increasing the observed ¹⁵O signal by about 2.5%.) It should be noted that these estimates for the ¹³N and ¹⁵O fluxes are about half those in the 2004 SSM due to a newly measured value for the cross section of the fusion reaction ¹⁴N(p, γ)¹⁵O. The estimated errors in the values are also large (30%+) due mainly to uncertainties in the solar composition for heavy element abundances [5]. A measurement of the CNO neutrino flux would therefore be invaluable in improving the accuracy of the SSM.



Figure 2.8: Expected spectrum of solar neutrino events in the 100-ton Fiducial Volume, in events per day per 0.1 MeV, for the energy range 0.7-1.4 MeV. The edge at about 750 keV is the tail of the ⁷Be neutrino distribution. Note that the vertical scale is much enlarged from Figure 2.6! The two dashed curves represent (upper) the ¹⁵O and (lower) ¹³N neutrinos from the CNO cycle; their amplitudes are uncertain by roughly 30%. The mainly horizontal solid curve, coming to an end near 1.35 MeV, is the *pep* neutrino spectrum. The dotted curve that peaks near 1.4 MeV is the ¹¹C radioactive background, suppressed by about a factor of five with delayed coincidence time and radial cuts. The uppermost curve is the expected total of all these signals.

For a neutrino with a continuous spectrum, Equation (2.18) and the definition of the electron kinetic energy spectrum $\rho_r(E; E_{\nu}) \equiv (1/\sigma_{\text{eff}}) (d\sigma_{\text{eff}}/dE)$ must be generalized to integrals over the normalized energy spectrum ρ_{ν} of the neutrino:

$$\frac{\mathrm{d}N}{\mathrm{d}E} = nVT \Phi \int_0^\infty \frac{\mathrm{d}\sigma_{\mathrm{eff}}}{\mathrm{d}E} (E; E_\nu) \rho_\nu(E_\nu) \,\mathrm{d}E_\nu$$
(2.22)

$$\rho_r(E) = \int_0^\infty \rho_r(E; E_\nu) \,\rho_\nu(E_\nu) \,\mathrm{d}E_\nu = \frac{\int_0^\infty (\mathrm{d}\sigma_{\mathrm{eff}}/\mathrm{d}E) \,\rho_\nu(E_\nu) \,\mathrm{d}E_\nu}{\int_0^\infty \sigma_{\mathrm{eff}}(E_\nu) \,\rho_\nu(E_\nu) \,\mathrm{d}E_\nu}.$$
(2.23)

In principle the differential and total cross sections within the integrals depend upon the neutrino energy, both directly through E_{ν} and E_{max} as shown in Equations (2.14) and (2.15), and indirectly through the energy dependence of P_e , the survival probability $P(\nu_e \rightarrow \nu_e | E_{\nu})$ for neutrinos traveling to meet us from the core of the Sun. However, in this energy regime

 P_e is still rather less than the vacuum-matter transition energy around 2 MeV. It is therefore predicted to be only weakly energy-dependent in this range, and may be approximated as ~ 0.55.

These considerations permit us to conclude that, within the 100-ton Fiducial Volume, the rate of *pep* neutrinos in the 0.8–1.3 MeV energy range will be 0.9 per day. That of the ¹³N and ¹⁵O neutrinos (combined) in this range will be 0.5 per day. The rate of ¹¹C radioactive background in this range, however, will be about 1 event/day, even if 80% of the ¹¹C background is suppressed by the delayed coincidence time and radial cuts described earlier. Figure 2.8 shows the expected energy spectrum in this range.

Let us assume a total data-taking time of 10 years. The ¹¹C exclusion cuts will reduce this figure by 7% to 9.3 years, giving total expected values for $N_s \approx 5000$ neutrino events in the 0.8–1.3 MeV range, and for a ¹¹C background in the same energy window of $N_b \approx$ 3500 events. From Equation (2.19), the statistical error in the measured *pep*+CNO neutrino rate will be about 2%. If we also treat the *pep* neutrinos as "background" in order to get a measurement of the CNO cycle neutrinos, we will see roughly 1700 CNO neutrinos with a statistical error of about 7%. This will still provide a significantly better estimate than the current uncertainties in the CNO rates of 30%.

It should be noted that, due to external background from the PMTs and other sources, the Fiducial Volume for *pep* neutrino analysis may have to be reduced to only 70 tons of scintillator. In this case, the statistical errors quoted in the preceding paragraph must all be multiplied by a factor of 1.7. It should also be noted that we have made no methodical consideration of potential systematic errors from background other than ¹¹C in the *pep/*CNO neutrino energy window.

2.4 Prospects for observing other neutrinos

2.4.1 Geoneutrinos and reactor antineutrinos

Borexino may observe antineutrinos from two sources. The first is the radioactive decay of elements in the Earth's crust and mantle. Each radioactive decay by β^- emission produces an electron antineutrino. The second source of antineutrinos consists of nuclear reactors in Europe, within a few hundred km of the detector. Just as with solar neutrinos, low-energy antineutrinos may in principle be detected via electron scattering:

$$\bar{\nu} + e^- \to \bar{\nu} + e^- \tag{2.24}$$

However, any antineutrino signal in the sub-MeV range is expected to be swamped by the ⁷Be solar neutrinos, particularly since the cross section for $\bar{\nu}_e e^-$ scattering is much smaller over most of this energy range than for $\nu_e e^-$ scattering.

A distinctive signal for antineutrinos with sufficiently high energies is available through the inverse β -decay reaction on protons,

$$\bar{\nu}_e + p \to n + e^+. \tag{2.25}$$

The reaction can occur only for electron antineutrinos. The minimum antineutrino energy required for the reaction is $(m_n + m_e - m_p)c^2 \approx 1.8$ MeV. Since most of the kinetic energy is carried away by the positron, the spectrum is nearly monoenergetic given a specific antineutrino energy. The immediately following annihilation of the positron means that the energy observed for the reaction is $E \approx E_{\nu} - 782$ keV. Thus, the observed energy is always greater than $2m_ec^2$.

By itself, this fact would be insufficient to detect the antineutrino signal; many radioactive backgrounds may be present in this energy range. However, the neutron travels less than a meter before being captured by another proton, forming a deuterium nucleus. In the process a 2.2 MeV γ ray is emitted. The mean time until capture is about 250 μ s [41]. The



Figure 2.9: The antineutrino signals observed by the KamLAND detector. The points are the experimentally observed values for 0.17 MeV-wide bins, shown with 1σ vertical bars. Total expected signal is given by the thin dotted black line. The solid bold line is the theoretical sum of all backgrounds. Individual backgrounds shown include reactor antineutrinos (blue dashed curve), the reaction ${}^{13}C(\alpha, n){}^{16}O$ (yellow dotted curve), and accidental coincidences (low-amplitude violet dashed curve at far lower left). Signals include the ${}^{238}U$ chain antineutrinos (dot-dashed red curve) and ${}^{232}Th$ chain antineutrinos (dotted green curve). Figure taken from reference [65].



Figure 2.10: Predicted spectrum of antineutrinos from Earth-based radioactive decays and nearby nuclear reactors in Borexino. Figure taken from reference [51].

signal of an antineutrino event therefore consists of a coincidence with $E_1 > 1.02 \text{ MeV}$ and $E_2 \approx 2.2 \text{ MeV}$. With appropriate radial and time cuts, background noise may be removed with high efficiency.

The rate of inverse β decay caused by geoneutrinos has been measured by the KamLAND experiment to be $5.1^{+3.9}_{-3.6} \times 10^{-31} \bar{\nu}_e$ events per target proton per year [65]. In the 300 tons of scintillator in Borexino (there is no need to restrict analysis to the Fiducial Volume due to the high efficiency with which most background are excluded), containing 1.35×10^{31} hydrogen atoms, this corresponds to 7 ± 5 events/yr. However, the rate will be rather different in various parts of the world, due to the varying radioactivity levels in continental and oceanic crust. Values of 10 events/yr (Figure 2.10) and 24 events/yr have been estimated for the specific location of Borexino in references [51, 52], respectively.

In KamLAND, the only significant background to interfere with detection of these antineutrinos (see Figure 2.9) is the reaction ${}^{13}C(\alpha, n){}^{16}O$, caused when α particles emitted by radon daughters collide with nuclei of ${}^{13}C$ (present in normal carbon at 1% isotopic abundance). The number of such background events that passed the coincidence selection cuts, for KamLAND, was estimated at 42 ± 11 over 749 days of live time in a spherical volume with 5-m radius [65]. The observed rate of 210 Po α -decays (the main source of α particles) in KamLAND was 33 Bq within a 5.5-m radius sphere [66]. Hence this reaction must produce an event rate of about 2.5×10^{-8} that of the α decay rate in a detector.

For the Borexino S/N for geoneutrinos to exceed ten, we therefore require no more than $7 \times 10^4 \alpha$ decays per day in the 300 tons of scintillator. Given that the target cleanliness level for Borexino is more on the order of $< 90 \alpha$ decays per day in the central 100 tons, the detector should be able to fulfill this additional requirement automatically. Given that geoneutrinos should easily be seen in Borexino, reactor antineutrinos (which, as seen in Figures 2.9–2.10, extend up to 8 MeV in energy) also will certainly be detected. The expected signal due to European reactor power plants is about 28 per year, of which only about $\frac{1}{4}$ will be below 3 MeV [51, 52].



Figure 2.11: Simulated energy spectrum from supernova neutrinos in Borexino. In a lowenergy window (top), many events will be due to the 2.2 MeV γ ray emitted during neutron capture after an inverse β decay reaction. (νp elastic scattering events were not included in this graph.) In the high-energy regime (bottom), there should be a prominent peak due to neutral-current reactions on ¹²C nuclei. The ability to tag β decays following chargedcurrent ¹²C interactions should also make those events easily identifiable. Figure taken from reference [53].

2.4.2 Supernova neutrinos

If a supernova were to occur in or near our galaxy, Borexino would see a burst of events with a duration of about ten seconds. Several detection mechanisms are possible. The first is the simple neutrino-electron scattering that may occur for all flavors of neutrinos and antineutrinos (with electron neutrinos having the highest cross-section). Another is the just-discussed inverse beta decay, sensitive only to electron antineutrinos. Three reactions on 12 C nuclei may also occur:

• The reaction ${}^{12}C(\nu_e, e^-){}^{12}N$ has a threshold of 17.3 MeV. The isotope ${}^{12}N$ quickly decays again by β^+ emission with a half-life of 11.0 ms.

- The reaction ${}^{12}C(\bar{\nu}_e, e^+){}^{12}B$ has a threshold of 14.4 MeV. The isotope ${}^{12}B$ decays by β^- emission with a half-life of 20.2 ms.
- The reaction ${}^{12}C(\nu_x,\nu_x){}^{12}C^*$ (a neutral current inelastic scattering that promotes the ${}^{12}C$ nucleus to an excited state $J^{\pi} = 1^+$) has a threshold of 15.1 MeV. The excited state of ${}^{12}C$ returns to the ground state immediately by emitting a monoenergetic 15.1 MeV γ ray. Since this is a neutral current reaction, any flavor of neutrino or antineutrino may be involved.

As described in reference [53], the total number of these events seen in the 300 tons of scintillator for a typical Type II supernova (total energy release of 3×10^{53} ergs) at a distance of 10 kpc should be approximately 110 (Figure 2.11). Of these, ~ 80 should be inverse β decay reactions, indicating a $\bar{\nu}_e$. These will be easily recognizable because of the following 2.2 MeV γ ray produced by neutron capture on a proton. Another 4–5 will be charged-current reactions indicating ν_e 's and $\bar{\nu}_e$'s; also easily identified by the following high-energy β^{\pm} decays. Of the remainder, about 22 will be neutral-current interactions with ¹²C (identifiable by the single monoenergetic 15.1 MeV γ ray), and about five will be electron scattering reactions; both channels are available to all flavors of neutrinos and antineutrinos. The great majority of these 110 events should be distinct and easily identifiable as supernova neutrino signals.

The inverse β decay and charged-current reactions with ¹²C are only relevant to ν_e and $\bar{\nu}_e$. None of the neutral-current reactions described above really give an indication of the energy spectrum of the original neutrinos, so it is impossible from these data to determine the spectrum of $\nu_{\mu,\tau}$ and $\bar{\nu}_{\mu,\tau}$. To remedy this deficiency, J. Beacom, W. Farr and P. Vogel proposed to look for neutral-current neutrino-*proton* scattering [54]. This process has a high-energy end point for the proton recoil kinetic energy at $E_p \approx E_{\nu}^2/m_p c^2$, and the differential cross section is weighted toward higher energies. However, as is the case with α particles, the energy observed for protons moving in a liquid scintillator is heavily quenched; see for instance Section 3.1.2. The expected number of events resulting from νp scattering

above an observed energy threshold of 200 keV, for a standard supernova at 10 kpc, is nevertheless 20/kton total for ν_e and $\bar{\nu}_e$, and 253/kton total for all other neutrino flavors [54]. This implies a total of ~76 ν_p proton-scattering events to be observed in Borexino. Similar considerations for KamLAND suggest that (with at least two operational low-energy neutrino detectors) it will be possible to make a reasonably good observation of the energy spectrum of ν_{α} and $\bar{\nu}_{\alpha}$ ($\alpha \neq e$) supernova neutrinos from a nearby stellar event.

The actual mix of neutrino flavors, and the energy spectra observed for each, will depend on the details of physical processes occurring in the supernova. These are not yet completely understood, and the study of neutrinos originating in a supernova would provide much important new data.

Chapter 3

Design and Hardware of Borexino

The Borexino experiment is among the next generation of neutrino detectors. It is located in the Gran Sasso National Laboratory (LNGS), a facility in the Apennine Mountains operated by Italy's National Institute of Nuclear Physics (INFN). The laboratory is underground (3800 meters water-equivalent), easily accessible from a highway tunnel.

In order to reach a sufficiently low rate of radioactive background events, the design of the Borexino neutrino detector is based on the principle of graded shielding. As one peers deeper into the heart of the detector, each region in turn is shielded and completely surrounded by the previous one, and must exhibit still lower levels of radioactive contaminants. Though the Borexino experiment comprises a large array of chemical plants, sensors, and control systems as well, the main focus of this chapter is on the detector itself. (Descriptions or overviews of various ancillary Borexino plants may be found, for instance, in references [42, 55, 67, 68].)

The general layout of the detector, a set of nested spheres, is shown in Figure 3.1. At the very center is the Fiducial Volume, a region containing 100 tons of scintillator fluid (an organic liquid whose precise composition will be described below). This region, from which neutrino data will be acquired and analyzed, is nominally 6.0 m in diameter. However, it is defined only by a volume cut in analysis software. If the rate of radiation from the outer parts of the detector (γ rays, radon gas diffusing inward, etc.) within the Fiducial Volume is lower than expected, it may be taken to cover a larger region and thus provide better


Figure 3.1: A schematic cross-section of the Borexino detector. From the interior going outward, regions shown include the Fiducial Volume; the remainder of the scintillator; the inner nylon vessel; the inner buffer region; the outer nylon vessel; the outer buffer region; the photomultiplier tubes; the Stainless Steel Sphere; the water-filled muon veto region; and the outer tank. Scintillator fluid is shown in bright yellow, buffer fluid in pale yellow, and water in blue. Figure courtesy of F. Calaprice and M. Leung.

statistics. Conversely, if external radiation is more of a problem than is hoped, the Fiducial Volume may have to be restricted to a smaller region to reduce background noise in the neutrino spectrum. The Fiducial Volume forms the central portion of the Inner Volume, the region of active scintillator fluid in Borexino. Since the remainder of the Inner Volume behaves as an active buffer, Borexino must have the capability of position reconstruction, which will be studied further in Chapter 5. These 273 tons of scintillator are contained in a leak-tight and transparent spherical nylon bag, the Inner Vessel, 8.4 m in diameter.

Enclosing this region is a spherical shell consisting of another 319 tons of buffer fluid, the Inner Buffer. The Inner Buffer is contained by a second nylon bag, the Outer Vessel, of diameter 10.9 m. Beyond the Outer Vessel is the Outer Buffer, a further 586 tons of buffer fluid. (These values are all nominal; uncertainties on the order of 1% in diameters, or 3% in masses, are possible.) The two nylon vessels, in addition to performing the necessary function of separating the scintillator and buffer fluids, serve to reduce the diffusion rate of radon gas into the Inner Volume by many orders of magnitude. Because the three volumes of fluid have nearly equal densities, the vessels need only be strong enough to support their own (buoyancy-reduced) weights and the buoyant forces resulting from a temperature gradient of no more than 5°C. They could therefore be made quite thin (125 μ m each) to reduce both the probability of brittle cracking, addressed in Chapter 4, and their contributions to radioactive background.

These three volumes of organic liquid and two nylon vessels are encompassed by the Stainless Steel Sphere (SSS), a rigid container with diameter 13.7 m and thickness 8 mm. The two vessels are attached at four end regions (two for each vessel) to the top and bottom of the sphere by cylindrical steel and nylon structures, the so-called tube assemblies. As well as the vessels, the SSS also supports 2240 inward-pointing photomultiplier tubes (PMTs) used mainly for data acquisition, and 208 outward-pointing PMTs used as part of a Muon Veto system. The SSS stands inside an outer tank filled with about 1000 tons of high-purity water used to detect the Čerenkov light produced by high-energy muons, and to reduce the number of γ rays originating from the rocks that surround the detector. Despite the location of the detector 1400 m underground, approximately 5000 muons pass within the SSS per day; they must be essentially eliminated from the data set in order to detect the neutrino signal, expected to be on the order of tens of events/day.

3.1 The scintillator and buffer fluids

3.1.1 The scintillator fluid

The scintillator to be used in Borexino is a solution of the fluor and wavelength shifter 2,5-diphenyloxazole (PPO) dissolved at a concentration of $1.5 \text{ g/}\ell$ in 1,2,4-trimethylbenzene, more commonly called pseudocumene (often abbreviated "PC"). For their physical properties, refer to Table 3.1. Pseudocumene is a hydrocarbon consisting of a benzene ring with three



Figure 3.2: The emission spectra of pseudocumene (blue) and 2,5-diphenyloxazole or PPO (red), shown in arbitrary units. These are shown together with the quantum efficiency of the Borexino phototubes as a function of wavelength (green). The wavelength shifter PPO is required so that scintillation light is maximally visible to the phototubes. Taken from reference [69].

Compound	Formula	CAS	Mass	Melting	Boiling	Density	Index
		registry $\#$	[amu]	pt. $[^{\circ}C]$	pt. $[^{\circ}C]$	$[g/cm^3]$	of refr.
water	H_2O	7732-18-5	18.0	0	100	0.997	1.333
\mathbf{PC}	C_9H_{12}	95-63-6	120.2	-44	170	0.875	1.504
PPO	$C_{15}H_{11}NO$	92-71-7	221.3	71	360	1.04	n/a
DMP	$\mathrm{C_{10}H_{10}O_4}$	131 - 11 - 3	194.2	0 - 5	282	1.19	1.515
PXE	$C_{16}H_{18}$	6196 - 95 - 8	210.3	n/a	295	0.988	1.565
pTP	$C_{18}H_{14}$	92-94-4	230.3	215	383	n/a	n/a
bis-MSB	$\mathrm{C}_{24}\mathrm{H}_{22}$	13280-61-0	310.4	180	n/a	n/a	n/a

Table 3.1: Physical properties of various compounds used in the Borexino and CTF scintillators and buffer fluids. Densities are given at room temperature. From top to bottom: water, pseudocumene (PC), 2,5-diphenyloxazole (PPO), dimethylphthalate (DMP), phenyl*o*-xylylethane (PXE), *p*-terphenyl (pTP), and 1,4-bis-(2-methylstyryl)-benzene (bis-MSB). See the text in this section and in Section 6.4.2 for structural formulas and more information. From references [44, 70, 71, 72, 73]. "n/a" represents a value not available in the literature.

methyl groups, while PPO consists of a heteroatomic five-membered ring with two benzene rings attached. These benzene rings play a vital part in the scintillation of these compounds. Their chemical structures are shown below:



A radioactive decay or neutrino interaction in the scintillator causes ionization: the passage of the high-energy charged particle through the medium strips electrons away from their original molecules. As these electrons recombine with molecular ions, scintillator molecules are excited electronically, resulting in the emission of scintillation photons. Some of these photons are detected by the PMTs. The authoritative reference in the field of organic scintillators is [74], which is unfortunately out of print and very difficult to obtain. Summaries of the relevant theories are provided in references [46, 75, 76]. Here, only a basic discussion of the scintillation mechanism is given.

Each carbon in an aromatic ring is sp^2 -hybridized (bonded to three other atoms), locking up three of its four valence electrons. The remaining electron is referred to as a π electron and is more or less free to travel around the ring. Scintillation in aromatic compounds is a consequence of the electronic level scheme of the π -electron cloud. The level structure of aromatic compounds consisting of one ring or several fused rings can be understood in an approximate way with the Perimeter Free Electron Orbital (PFEO) model. In the simplest approximation, the π electrons are presumed not to interact with each other, nor with the ring atoms, but to behave simply as standing waves on the ring. It is these delocalized electrons that are a prerequisite for organic compound scintillation; molecular rings with fully sp^3 -hybridized atoms (only single bonds) such as the cyclohexane molecule do not scintillate. The π electrons obey the one-dimensional Schrödinger equation on a loop of radius r,

$$\left(\frac{\hbar^2}{2mr^2}\partial_\theta^2 + E\right)\psi(\theta) = 0, \tag{3.1}$$

with the single-value condition, $\psi(\theta) = \psi(\theta + 2\pi n)$ for all integer *n*, applied. (In multiring systems, the "loop" in question represents the entire perimeter of the system of linked rings.) The single-value condition implies that the normalized single-electron basis states are discrete and look like

$$\begin{cases} |0\rangle \equiv \phi_0(\theta) = (2\pi)^{-1/2} \\ |\pm q\rangle \equiv \phi_{\pm q}(\theta) = (2\pi)^{-1/2} e^{\pm iq\theta}. \end{cases}$$
(3.2)

Since r, the radius of a benzene ring, is about 1.4 Å, the basic unit of energy is $\varepsilon \equiv \hbar^2/(2mr^2) \approx 1.95 \,\text{eV}$. Energies of the single-electron basis states are of course given by $\hat{H} |\pm q\rangle = q^2 \varepsilon |\pm q\rangle$; note that each energy level except the lowest is degenerate with multiplicity two. As electrons are spin- $\frac{1}{2}$, two (spin-paired) electrons may reside in the lowest energy level, and four in each of the other levels. Therefore the ground states of aromatics with six carbons in a single ring have completely-filled π -electron orbitals up through |q| = 1. We would predict that the first excited state of benzene, with one electron promoted from a |q| = 1 to a |q| = 2 energy level, is at an energy $E_1 = (2^2 - 1^2)\varepsilon = 3\varepsilon \approx 5.8 \,\text{eV}$ above the ground state. Due to electron spin-orbit interactions and the separation of $|\pm q\rangle$ states into different linear combinations because of the atomic potentials, the first excited state in benzene actually splits into six different energy levels with E in the range 3.7–6.7 eV above the ground state [74]. Still, the lowest-order PFEO approximation is not so far off.

The most important fact about the level splitting of the first excited state results from electron spins. If an electron is promoted from |q| = 1 to |q| = 2 but remains spinpaired with its original counterpart, the state is a singlet (total spin angular momentum $|\mathbf{s}| = 0$). On the other hand, if the electron is spin-flipped, the system has total spin angular momentum $|\mathbf{s}| = 1$ unit of \hbar , and therefore may be in one of three states, a triplet, with differing values for s_z . By convention the singlet states are designated S₀ (the ground state), S₁, S₂, etc. in order of increasing energy, while the triplet states are T₁, T₂, and so on. (An



Figure 3.3: Energy levels of pseudocumene. From left to right: in part a is shown the π -electron configuration of an S₀ ground-state molecule; in part b is the S₁ excited state; in part c is the T₁ excited state; and in part d is the molecular level scheme. The energies shown in part d are relative to the ground state; the values are taken from the experimental data (for the S₁, S₂ and T₁ levels) and computational models (for the rest) reported in reference [77]. It should be noted, however, that values reported for the pseudocumene energy spectrum in the literature differ from each other by up to 0.25 eV [77, 78]. Some possible level transitions are also shown: i, energy absorption from a radioactive decay; ii, internal conversion from S₃ to S₁; iii, fluorescence; iv, radiationless intersystem crossing from S₁ to T₂; v, internal conversion from T₂ to T₁; vi, phosphorescence. Figures adapted from reference [75].

alternate nomenclature has the ground state labeled ${}^{1}X$, with excited singlet states ${}^{1}X^{*}$, ${}^{1}X^{**}$ and excited triplet states ${}^{3}X^{*}$, ${}^{3}X^{**}$.) The first excited triplet state is always lower in energy than the first excited singlet state. This is an instance of Hund's Rule: since the spin component of a wave function with parallel electron spins is symmetric, the spatial component must be antisymmetric to preserve total antisymmetry of the wave function. Hence the electrons, which repel each other via the Coulomb interaction, may be farther apart more often and therefore are lower in energy. The energy levels of a pseudocumene molecule, for example, are shown in Figure 3.3.

The ions and secondary electrons produced when a fast charged particle travels through a scintillator recombine very quickly (< 0.1 ns), yielding highly excited molecules in states S_{2+} and T_{2+} . Other molecules are excited through mechanisms such as transfer of vibrational energy, which only produces excited singlet states. However they are excited, these molecules typically lose energy non-radiatively, falling very rapidly into the S_1 and T_1 states, respectively. A scintillator molecule in the S_1 state will decay to the ground state S_0 . (The intersystem crossing $S_1 \rightarrow T_1$, while energetically possible, is rare. This transition requires one electron to flip its spin and is termed "spin-forbidden.") Any photons emitted in the process $S_1 \rightarrow S_0$ are referred to as fluorescence. This immediate fluorescence of S_1 makes up a so-called prompt, or fast, component of scintillation light with a duration on the order of several ns. The decay to S_0 may also occur through a radiationless process called internal conversion, a transition to higher vibrational states of a lower-energy electronic state. The ratio of the number of fluorescent decays of S_1 to the total number of S_1 decays is termed the fluorescence quantum efficiency $q_{\rm fm}$ of the scintillator compound.

A molecule in the T_1 state could, like an S_1 molecule, eventually decay to the S_0 ground state by photon emission ("phosphorescence") or internal conversion. However, because phosphorescence is spin-forbidden, it occurs on a time-scale too long to observe in scintillation detectors, on the order of ms or even seconds. It is much more likely that two molecules in triplet states will collide, pushing one into an excited singlet state that can then emit scintillation light: $2T_1 \rightarrow S_1 + S_0 + vibrational energy$. This process contributes to a slow component of the scintillation pulse which lasts several μ s, much longer than the few-ns fluorescence decay time.

Although it would technically be possible to use pure pseudocumene as a scintillator, this would not be an ideal state of affairs. The pseudocumene molecule has three properties that make it, by itself, an unattractive scintillator.

• The mean lifetime of the excited S_0 state is $\tau_m = 22 \text{ ns}$ [75]. In a liquid with a refractive index of 1.504, this corresponds to a photon travel distance of 5.4 m, which

is on the same order as the Borexino detector size. Attempting to perform position reconstruction of events with such a poor resolution would be difficult.

- The fluorescence quantum efficiency for pseudocumene is at best $q_{\rm fm} \approx 40\%$ [79]. Since already only about 5% of the kinetic energy of a β particle in scintillator fluid is eventually converted to electronic excitation energy [75], this would imply a scintillation light yield of about 4500 photons/MeV for β events.
- The peak emission wavelength of scintillation light is about 290 nm, compared to the wavelength of about 360 nm at which the Borexino PMTs have the highest quantum efficiency (Figure 3.2).

The compound PPO solves all of these difficulties. The mean lifetime of its first excited state is 2.5 ns, and it has a fluorescence quantum efficiency of about 80% [75]. Furthermore, it is possible to transfer energy from an excited pseudocumene molecule to a molecule of PPO more than 95% of the time through a non-radiative mechanism that happens much more quickly than photon emission [75]. The PPO molecule then rapidly scintillates with a peak emission wavelength of 360 nm, near the point of maximum quantum efficiency for the PMTs (Figure 3.2). With an infusion of $1.5 \text{ g/}\ell$ PPO in the pseudocumene, the scintillation light yield is on the order of 11,500 photons/MeV for β events [80]. It should be noted that increasing the concentration of PPO would not improve the light yield any, as the fluor would begin to absorb too much light.

3.1.2 α particle quenching

As mentioned above, about 5% of the kinetic energy of a β particle (including those produced by neutrino interactions in the scintillator) or a γ ray is eventually converted to electronic excitation energy in the molecules of scintillator. For α particles, however, the sensitivity is roughly an order of magnitude worse, a phenomenon called α particle quenching. It occurs because α particles lose their kinetic energy over a much shorter path length, causing



Figure 3.4: The rate of photon emissions over time for typical α and β events in scintillator fluid. The *y*-axis, shown in arbitrary units, is logarithmic. These curves were reconstructed from a large number of events. The α events have a higher-amplitude "tail," caused by their greater proportion of scintillator molecules that decay radiatively after interactions between the longer-lived triplet states. Figure taken from reference [81].

them to produce a higher specific ionization (ionization in the scintillator per unit path length). When the density of ionized molecules is high, electronically excited molecules in high-energy states are likely to interact with an ion and return to the ground state nonradiatively instead of falling into the S₁ or T₁ states [76]. Such interactions are referred to as ionization quenching. As a result, α particles from radioactive decays, which typically have kinetic energies in the range 4–8 MeV, produce an amount of scintillation light similar to β events in the 400–800 keV range, overlapping the ⁷Be neutrino energy window.

Fortunately, there is a means by which α and β events may be discriminated. A qualitative description will suffice to illustrate the idea, though the argument is made quantitatively by the Voltz model of scintillation time dependence [82]. Consider a β event whose track has a low ionization density and a higher-energy α event (with a high specific ionization) that produce equal amounts of prompt scintillation light. That is, the initial number of molecules in the S₁ state is the same in each track. At low ionization density, a large proportion of excited states are produced simply by molecular excitation. At high ionization density, the proportion shifts toward those formed by ion-electron recombination, which favors the triplet states by the statistical 3:1 ratio. Therefore the initial number of molecules in the T_1 state is greater for the α particle. In addition, the α event track is much shorter. Due to these effects, the density of T_1 molecules along it is much higher. The rate of the reaction $2T_1 \rightarrow S_1 + S_0$, which depends upon the density of T_1 states, is therefore much faster in the case of the α . But by the time this reaction occurs, the original ions have all recombined: ionization quenching cannot compete with the $S_1 \rightarrow S_0$ fluorescence. Thus an α scintillation event will exhibit a higher-amplitude slow component of scintillation light than a β event of comparable light yield; see Figure 3.4.

The time dependence of scintillation pulses is discussed further in Section 7.5. Data acquisition of observables such as the tail-to-total ratio [83], as described in Section 8.1.3, can generally tell α and β events apart with efficiencies of better than 90%.

3.1.3 The buffer fluid

Near the PMTs, Borexino requires a volume filled with inert material, or "buffer fluid." The PMTs (Section 3.3) have a much higher level of radioactivity than other components of the inner detector, and immersing them in scintillator would flood the data acquisition system with a high trigger rate. Even pure pseudocumene, despite the absence of a wavelength shifter, would still scintillate sufficiently in the spectral range of the PMTs to overwhelm the neutrino signal in Borexino [84]. However, there are numerous advantages to using another pseudocumene solution as the buffer fluid, so a compound that could be dissolved in pseudocumene and quench its scintillation was sought. Such compounds typically act by increasing the probability of non-radiative relaxations from excited states.

Tests of several candidate quenchers found that the substance dimethylphthalate (DMP) exhibits excellent properties such as chemical compatibility with nylon, low light absorption in the visible and near-UV range, and relatively low cost. The buffer fluid will therefore be a solution of $5 \text{ g}/\ell$ of DMP in pseudocumene.

The chemical structure of DMP is shown below:



This solution, when tested with a ²⁰⁷Bi source, was quenched below the light yield of pure pseudocumene by a factor of 6.8 ± 0.4 . When the photon contribution due to Čerenkov radiation from β particles was neglected (via comparison with a sample of cyclohexane, an organic liquid that does not scintillate), the quenching factor was found to be 28 ± 2 [85].

The possibility of scintillator solution leaking slowly into the buffer fluid (or vice versa) has also been considered. The leak rate of the Inner Vessel has been measured as no greater than 5×10^{-3} cm³/mbar/s for pseudocumene [49]. During the expected ten-year lifetime of the experiment, at the planned operational over-pressure of one mbar for the Inner Vessel, this is equivalent, at worst, to an exchange of 1.5 m^3 of fluid in either direction between the Inner Volume and Inner Buffer. The corresponding resulting concentrations of PPO in the buffer fluid and DMP in the scintillator fluid would be 7 ppm and 27 ppm by weight, respectively. At less than 100 ppm of DMP in the scintillator solution, the light yield would be reduced by less than 1%. On the other hand, 7 ppm of PPO in the buffer fluid would reduce the quenching factor from about 7 to about 3 [85]. Though not desirable, this should not greatly hinder the experiment. If the buffer fluid were pure pseudocumene (no quencher), the addition of PPO at 7 ppm would increase the light yield to 2.5 times that of pure pseudocumene, or equivalently, set the quenching factor to 0.4 [85].

The advantages to using a second pseudocumene solution as the buffer fluid come about mainly because many physical properties of the two solutions are almost identical. The densities differ by less than one part per thousand [44], which would correspond to a temperature difference between two samples of pure pseudocumene of about 1°C. (The nylon vessels were designed to survive a temperature differential of 5° C.) The indices of refraction are very similar, so photons can travel freely from the scintillator, through the buffer regions, to the PMTs, all the while maintaining constant directions of travel. As will be explained in detail in Chapter 7, a detector comprising two regions with significantly different indices of refraction can greatly complicate data analysis. A previously proposed option for Borexino, a combination of phenylxylylethane for the main scintillator component and water for the buffer fluid [73], was less satisfactory with respect to the density difference (1%) and would have had a serious index of refraction mismatch (1.565 versus 1.333).

As an aside, it should be noted that the oxygen molecule is also a fairly effective quenching compound. The presence of oxygen dissolved in pseudocumene and PPO scintillator has altered the energy spectrum in tests with radon sources performed in the second and third versions of the Counting Test Facility (Sections 7.7 and 7.8). A radioactive source consisting of radon dissolved in a pseudocumene and PPO scintillator can experience energy quenching up to a factor of three [85].

3.2 Scintillator containment vessels

The two nested spherical nylon vessels in Borexino serve the dual functions of separating the active scintillator fluid from the buffer fluid, and preventing radon gas from diffusing into the volume of scintillator. They are described in detail in references [12, 44, 49]. Figure 3.5a shows one of the vessels under construction, and Figure 3.5b shows the inflated vessels as seen by the Borexino camera system.

3.2.1 Requirements for the vessels

The vessels must satisfy a number of mechanical, chemical, optical, and radiopurity requirements in order to be suitable for use in the Borexino detector.



Figure 3.5: The Borexino nylon vessels. At left (part a), one of them is shown under construction in the Class 100 radon-filtered clean room in Jadwin Hall, Princeton University. Here, two panels of nylon film are being glued together at their edges. At right (part b) is a photograph of the vessels taken with the Borexino camera system shortly after installation inside the Stainless Steel Sphere. The Outer Vessel and Inner Vessel are clearly visible, as are the system of ropes and belts on the Outer Vessel and the south end region of the Inner Vessel. This photograph was taken while the vessels were filled with air. When surrounded by pseudocumene, which is a better match for the index of refraction of nylon, they will appear much more transparent.

The operational stress on the nylon vessels in Borexino will be about 1.8 MPa [12]. However, the mechanical strength of the nylon vessels must be sufficient to withstand a short-term temperature differential of up to 5°C between any two of the volumes they separate. The resulting difference in density would cause buoyant forces to produce a stress of up to 20 MPa on the nylon film [12]. In addition to strength, they cannot be very brittle. This is not only a theoretical worry; the nylon vessel installed in the second version of the Counting Test Facility (a Borexino prototype) cracked open in two places shortly after its installation. The material properties of tensile strength and brittleness are discussed in detail in Sections 4.2 and 4.3.

The vessels must also not leak. This is particularly important for the Inner Vessel, which separates scintillator and buffer fluids. The leak tightness requirement for the Inner Vessel was set at a maximum of 10^{-2} cm³ of pseudocumene per second per mbar of differential pressure, and as stated above, it was measured to meet the requirement by a factor of two [49]. The Outer Vessel, which separates two regions of the same composition, does not have nearly as strict a requirement. Its leak rate specification was arbitrarily set to be $1 \text{ cm}^3/\text{s/mbar}$, and was measured at better than $0.1 \text{ cm}^3/\text{s/mbar}$ [49].

Of course, the vessels must exhibit chemical compatibility with both the buffer and scintillator solutions. This was one of the requirements that ruled out the use of an acrylic shell such as that used in the Sudbury Neutrino Observatory (whose detector is filled with heavy water). They must also be able to withstand a period of several months immersed in water. At the time of this writing, the possibility of filling Borexino first with water, followed later by replacement of the water with buffer and scintillator fluids, was still under consideration as a means of cleansing the vessels. As a complicating factor, it is known that the presence of water has the effect of making nylon films less brittle, yet weaker, a reversible phenomenon described in Sections 4.2 and 4.3.

The nylon vessels need to be transparent at the wavelength of scintillation light. Their transmittance in pseudocumene at 366 nm has been estimated at roughly 98% based on measurements in air [12]. A marginally less important criterion is that their index of refraction should be similar to that of pseudocumene (1.504). The refractive index of nylon film is about 1.53, less than a 2% difference [12]. These parameters are both well within acceptable ranges.

As the solid components of the detector infrastructure in the most intimate contact with the scintillator fluid, the nylon vessels must also be the cleanest components. Any dirt adhering to the vessel surfaces, or radioactive atoms adsorbed on them, may wash off into the scintillator. Specific activities of γ -ray emitting isotopes such as ⁴⁰K and various uranium and thorium daughters must be very low. The vessels must act as a barrier to radon diffusion. Furthermore, their intrinsic contamination of ²³⁸U and ²²⁶Ra must be sufficiently low that the amount of radon gas produced within the Inner Vessel, which then diffuses into the Inner Volume of scintillator, is small. The ability of radon to permeate through nylon is much greater when the nylon is wet. Still, it will be shown in Section 4.4 that if the water content of the scintillator and buffer fluids is kept below about 30% of the solubility of water in pseudocumene, these constraints should not prove difficult to meet.

3.2.2 Fabrication and design of the vessels

To prevent radon daughters and dust from collecting on the surfaces of the nylon vessels, they were constructed inside a radon-filtered Class 100 clean room located in Jadwin Hall, Princeton University. The clean room and radon filter, as well as details of the vessel construction process, are described in reference [44]. The construction took approximately one year, summer 2001 through summer 2002, and involved on the order of 5–10 skilled personnel at any point.

The nylon film

Each of the two vessels was constructed from a large number of individual flat panels (36 for the Inner Vessel and 40 for the Outer Vessel). First, unrolled nylon sheets were covered with clean plastic film (to protect from deposition of radon daughters) and allowed to reach equilibrium over several days with the ambient humidity of the clean room, about 50%. (Nylon absorbs water from its environment and expands, hence the need to attain equilibrium before performing any precision operations.) They were then cut, using templates, into the shape of the cross-section of a convex lens, with rough dimensions $15 \text{ m} \times 80 \text{ cm}$. Mechanized carts designed by A. Nelson, running along magnetic tracks laid down over the template edges, were used to semi-automate the cutting. The exact compositions and properties of the nylons used are described in Section 4.1.4 and following.

Once all the panels were cut, they were glued together at the edges, one at a time, using a Resorcinol-based solvent [44]. As little film surface was exposed to the air as possible during this operation. Again, mechanized carts were used to semi-automate the process (Figure 3.5a). Hundreds of C-clamps were used to hold each seam together while the glue cured. As a vessel was put together, each newly attached panel was laid over the stack of already assembled panels, forming an accordion-like pile. For ease of operations, the pile was kept as tidy as possible; indeed, the vessels were never inflated to their final spherical shapes until they had been installed within the Stainless Steel Sphere at LNGS. The stacks of panels were kept soaking wet as much as possible, so they would remain supple and easier to handle. When not being actively worked upon, the stacks of panels were kept covered by a metallized film. This minimized exposure of the Resorcinol to light, which eventually causes it to develop a green tinge. (The amount of this effect caused by scintillation light during the operation of the Borexino experiment will be negligible.)

Some small defects discovered in the film panels (embedded pieces of foreign material, dents that could become a puncture, creases in the seams, etc.) were fixed by cutting around each of them and replacing the excised film with a patch. The patch was glued to the surrounding material, and then clamped for curing, in a way very similar to how the film panels were joined together.

The ends of the vessel panels—the parts that would ultimately become the top and bottom of each vessel—were truncated prior to the gluing operations. The structures at the top and bottom of a vessel, generally referred to as the north and south end regions, are special because they are the points of attachment for almost all the hardware, sensors, and mechanical supports of the vessel. They will be described further in Section 3.2.3. Following attachment of the Inner Vessel end regions to the film panels, the final seam between two panels was closed, forming a complete (albeit still-folded) sphere. Then, a system of holddown ropes, described immediately below, was strung along the vessel and attached to the north and south end regions. Light sources for vessel position monitoring (also described below) were attached at various points on the vessel, with the optical fibers that lead to them being routed away through the north end region. Once the Inner Vessel was thus finished, it was nested inside the Outer Vessel. The Outer Vessel end regions were attached (both to the Inner Vessel end regions and to the Outer Vessel film panels), and the Outer Vessel was closed at the final seam. It was then made complete by the installation of its system of ropes and optical fibers.

The rope and belt support system

Each vessel is supported by three sets of ropes. All the ropes are made from multithreaded Tensylon, an ultra-high molecular weight polyethylene made by Synthetic Textiles. The material was chosen due to its tensile strength (1.45 GPa) and its low level of contamination with potassium (0.2 ppm total) and heavy metal isotopes [49]. In addition, the ropes were specially woven by Cortland Cable without a coating; coatings of nylon ropes typically include potassium salts, which must be avoided due to the presence of the naturally occuring radioactive isotope 40 K. The ropes were cleaned with alcohol, and pre-stretched in the clean room before being cut to their proper lengths.

Two sets of ropes on each vessel make up the hold-down system, which prevents the two vessels from moving upward or downward in response to positive or negative buoyant forces resulting from density changes (most probably due to temperature gradients). On each vessel, $\frac{1}{2}N$ ropes pass over the top and are attached at each end to a vertical tube below the bottom of the vessel (N is the number of panels: 36 for the Inner Vessel and 40 for the Outer Vessel). At the top of the vessel, each rope passes twice through a nylon ring, so it cannot slip away. These ropes keep the vessel from moving upward. Another $\frac{1}{2}N$ ropes are arranged in the opposite way, passing over the bottom of the vessel and fixed above the top. The setup is illustrated in Figure 3.6.

At the points of attachment, the ropes are grouped in bundles of six for the Inner Vessel (five for the Outer Vessel) and each group is connected to a single load cell. Hence, the total buoyant force on a vessel may be found by summing over the tensions T_i of the ropes (or



Figure 3.6: Diagram of the Borexino hold-down system. The concentric circles represent the Inner Vessel (IV), Outer Vessel (OV), Stainless Steel Sphere (SSS), and top of the Water Tank (WT). The north and south end regions of the vessels, shown as gray vertical bars, both carry instrumentation and tubing between the external world and the inside of the detector, and hold the vessels in place within the sphere. Each vessel has two sets of Tensylon ropes wrapped around it meridionally: a set of ropes affixed at both ends to the north end regions (shown in bold), constraining the vessels against downward movement; and a set affixed at both ends to the south end regions (shown as dashed lines), constraining the vessels against upward buoyant forces. At top is shown a blown-up diagram of the point of attachment of one bundle of six rope ends to the north end region via a load cell. Figure taken from reference [44].

the load cell readings) and multiplying by the cosine of the attachment angle ($\theta_{IV} = 58^{\circ}$; $\theta_{OV} = 53^{\circ}$):

$$F_{\text{buoy}} = \sum_{i} T_i \cos \theta_x. \tag{3.3}$$

In the specification limit of a 5°C temperature difference between two volumes, the difference in density would be about 4 kg/m^3 [44], or a maximum possible upward or downward force of 13 kN on the Inner Vessel and 27 kN on the Outer Vessel. Additionally, the scintillator is less dense than the buffer fluid at the same temperature by 0.9 parts per thousand [44] (due to the slightly different compositions of the fluids), for an upward force on both vessels of 2.5 kN. Thus the maximum tension on any Inner Vessel rope, at the specification limit, should be 790 N, and for the Outer Vessel, 1250 N. The diameter of the ropes was chosen such that the maximum tension on the Inner Vessel ropes would be on the order of 10% of the rope yield strength, 5.8 kN [49]. It is actually closer to 15%; however, the development of a 5° temperature gradient is considered extremely unlikely, so the specification limit is rather a pathological case.

Originally it was planned for each rope to be routed along most of its length through holes in nylon tabs glued directly onto the outer surface of each vessel. A rope acting against upward buoyant force would run from its point of attachment at the south end region, diagonally upwards, to come into contact with the vessel surface; curve up and around the vessel along the center line of a film panel, constrained in place by tabs; pass through two holes in a solid nylon ring at the north pole; curve back downward along the center line of the film panel opposite the first one; separate from the film; and run in a straight shot back to another attachment point at the south end region. Likewise, a rope acting to prevent downward motion would follow a mirror-symmetric course from north end region, around the south pole of the vessel, back to north end region.

However, attaching the vessel film to the end regions such that the panel center lines met the end region at a perpendicular proved impossible. The panels had to be tilted slightly, and hence, they could not be relied upon to follow lines of constant longitude once the vessels were inflated. The meridional ropes were therefore freed from most of the tabs. Instead, a third set of ropes, called "belts," that follow lines of constant latitude, was added to each vessel. The two sets of meridional ropes were tied to the belts (using strings of Tensylon, the same material as the ropes and belts themselves) at fixed distances apart in order to prevent them from all bunching up on one side of a vessel. Together, the three sets of ropes form a complete mesh of latitude and longitude lines over each vessel.

The optical fiber calibration system

While the Stainless Steel Sphere is empty, monitoring the shapes and positions of the vessels is easily done using the Borexino camera system, as seen in Figure 3.5b. The system, designed at the Virginia Polytechnic Institute, consists of seven computer-controlled inward-pointing digital cameras and halogen lamps mounted on the inside of the SSS [86]. When the Borexino detector is filled with buffer and scintillator solutions, however, the vessels will (by design) be nearly invisible. Therefore, a number of point-like light sources, described more extensively in [44, 87], have been placed on the outer surfaces of both vessels. These will permit the positions of numerous sample points on the vessels to be determined to an accuracy of better than 2 cm [86].

These light sources consist of a set of lasers located outside the detector; a number of fiber-optic cables that pass into the Stainless Steel Sphere through the top in a leak-tight manner; and 78 $^{3}/_{32}$ "-diameter Teflon spheres, one attached at the end of each cable, that act to diffuse the laser light in an almost isotropic manner. 38 light diffusers are arrayed at various known points on the outer surface of the Inner Vessel, and the remainder are on the Outer Vessel. They are affixed to the vessel film by means of nylon tabs similar to those that were used to route the vessel ropes.

In addition, eight light sources for calibration of the Borexino event position reconstruction software are provided by 355-nm lasers shining into fiber-optic cables that end at the Inner Vessel outer surface, and are constrained to point directly into the scintillator fluid. These sources have no light diffusers. Instead, the optical fiber ends are embedded within nylon "buttons" glued to the vessel film in order to have the fibers point inward at a 45° angle to the vessel surface. While the lasers are turned on, PPO molecules in a region of scintillator a few cm in diameter will be excited directly. Performing event position reconstruction on these regions will show the behavior of the event reconstruction code near the edge of the scintillator volume, where it is expected to have the poorest resolution, without the potential danger of bringing a radioactive source near the nylon film. These calibration sources will also be used to monitor the quality of the scintillator over time and to study effects of light trapping near the Inner Vessel. As shown in Section 9.2.1, these effects are significant in the Counting Test Facility. It is hoped that in Borexino, where the scintillator and buffer fluids have almost identical indices of refraction, the light trapping will be negligible.

3.2.3 The vessel end regions

The end regions of the Borexino vessels serve numerous functions, so have quite a complicated design. They hold the vessels at a defined position within the Stainless Steel Sphere, and keep the hold-down system of ropes fixed in place against positive or negative buoyant forces. They form the points at which all the nylon panels on each vessel meet. They allow the passage of cabling for instrumentation—the fiber-optic cables for the calibration system, electric wires for reading out values on the load cells and temperature sensors, and sampling ports for retrieving scintillator and buffer fluid samples—as well as the insertion and removal of massive quantities of fluid during detector filling. The north end regions must permit the passage of a long rod carrying a radioactive source at the tip in order to allow source calibration measurements, a process described in [86]. Despite all these requirements, they must be as light as possible, to reduce γ rays from radioactive events; and as small in solid angle as possible, so as not to block light from reaching the PMTs.



Figure 3.7: The south end region of the Inner Vessel, shown during vessel construction shortly before the Inner Vessel was placed inside the Outer Vessel. (The purple strap wrapped around the central part of the end region is attached to a small crane in center background.) The other three end regions are similar in design. After installation, this structure was located near the bottom of the Stainless Steel Sphere, with the central steel and nylon tubes oriented vertically, and the left side of the end region pointed downward.

The support structures

The four support structures (two for each vessel, top and bottom) are all fairly similar in design. However, the Outer Vessel support structures are composed mainly of stainless steel, for strength, while the Inner Vessel structures nearest the scintillator are made mainly of copper and nylon, for reduced levels of radioactivity. The requirement for the γ ray background produced by the Inner Vessel end regions was that it be no more than 0.1 events/day in the 250–800 keV neutrino window within the 100-ton Fiducial Volume. If the vessels had no end regions, each of the lens-shaped film panels would come to a point at the north and south poles, just like lines of longitude on the Earth. However, the panel ends have been truncated so that each vessel has two circular holes at its polar regions. On the Inner Vessel, these are each about 60 cm in diameter. Two solid nylon rings make a rim for each Inner Vessel aperture. One is a "collar ring" that faces into the volume of scintillator. The other has a protruding rim that faces into the Inner Buffer volume, and is therefore L-shaped in cross section. The ends of the film panels are sealed between these rings with a combination of Resorcinol glue and a formic acid paste. (As noted above, the film panels do not meet the rings exactly in the radial direction, leading to slight twists in the vessels.) This transition from nylon film to bulk nylon required the greatest amount of research and development in the end regions' construction.

Each of these apertures of the Inner Vessel is kept sealed by an annular sheet of transparent nylon, the "bridge film," stretched across it. The bridge film is sealed between the collar ring and L-ring, like the film panel ends, on its outer edge. It is transparent to permit the passage of scintillation light to photomultiplier tubes near the top and bottom of the Stainless Steel Sphere. On its inner edge, it is sealed between two "clamp rings" that are fastened directly to a nylon pipe. (The bridge film was required because extending the 36 film panels all the way to the nylon pipe, where they would have been extremely narrow, would have been impractical.) Passing perpendicularly along the z-axis through the Inner Vessel aperture, the 68-cm long pipe is a 10-cm inner diameter bulk nylon tube that permits scintillator fluid to be added to or drained from the Inner Volume. To prevent compression of the L-ring and collar ring in the xy-plane, three solid nylon spokes connect the rings to the nylon pipe at a perpendicular. Furthermore, eight copper struts are bolted diagonally to the L-ring and to the pipe exterior, farther away from the vessel, to support the rings against vertical and shear forces. At the same point, the pipe makes a transition from nylon to stainless steel.

Both sets of ropes making up the hold-down system are fastened to each end region. At the south end region, the ends of each rope that acts against an upward buoyancy force are attached, in bundles of six, to load cells. Each load cell is fixed by steel shackles to the stainless steel portion of the central vertical pipe. On the other hand, the ropes that act against downward forces on the vessel remain close to the vessel film; they are each threaded through two holes, situated opposite each other in the protruding rim of the L-ring, in order not to slip away from the vessel south pole. The situation is exactly mirrored at the Inner Vessel north end region.

The design of the south end region of the Inner Vessel is shown as it is laid out in the Princeton clean room in Figure 3.7, just before the Inner Vessel was inserted inside the Outer Vessel. Many of the features described above are clearly visible.

The Outer Vessel end regions are similar in design but slightly less complex; the distance farther from the scintillator permitted the use of a stronger if less radiopure design. For instance, instead of nylon bridge film, an opaque bulk nylon disk fills the gap within each Outer Vessel L-ring and collar ring. The disk is attached at the center to a stainless steel central pipe. In this case the pipe consists of two nested cylinders; scintillator fluid will pass through the inner one into the Inner Vessel nylon pipe and then the Inner Volume, while buffer fluid will pass through the outer cylinder directly into the Inner Buffer region. The design of the rope attachment points, load cells, and support struts (though steel instead of copper) is basically the same as the Inner Vessel case. The Outer Vessel south end region is shown in Figure 3.13.

Each of the four end regions is required to withstand a vertical force of about 3 kN so that in the event of a catastrophic structural failure, the nylon vessel film will tear rather than the end regions breaking [44]. Even at full load, the leak rates of mock-up versions of the end regions were found to be better than the nominal specification of a $10^{-3} \text{ cm}^3/\text{s/mbar}$ liquid leak rate by five orders of magnitude [44, 49].

The instrumentation

In addition to the optical fiber calibration system, described already, the instrumentation within the Borexino vessels includes load cells to measure buoyant forces on the vessels and temperature sensors to monitor the scintillator and buffer fluid at different points.

There are 28 load cells in all attached to the hold-down system of ropes. Eight, each holding five rope ends, are attached to each Outer Vessel end region; and six, each holding six rope ends, to each Inner Vessel end region. The load cells are the Sensotec model 34, a custom-built model with a mass of 70 g and an expected systematic error (due to non-linearities and hysteresis effects) of at most 0.2% of its full-scale range. Each of the Outer Vessel load cells may experience a maximum load (in the highly unlikely limit of a 5°C temperature difference between fluid volumes) of 6.3 kN, or 4.7 kN for the Inner Vessel load cells. They are therefore tuned to have respective full-scale ranges of 10 and 5 kN.

Eight more load cells (two for each of the four end regions) are intended to measure elongation under load of the main stainless steel pipes. These cells are each attached at their ends to two fixed points along the outside of the steel pipe. They have nominal full-scale capacities of 110 N, although the measured values will be converted to measures of elongation rather than stress.

Eight temperature sensors are mounted at the north and south end regions as well. These are custom-made Pt100 RTDs produced by Fisher Rosemount (Italy), rated to have precisions of better than $\pm 0.1^{\circ}$ C [49]. They are positioned in order to measure vertical temperature gradients within the buffer fluid: two measurement points at each end within the Outer Buffer volume, and two at each end within the Inner Buffer. (No sensors are located within the Inner Volume due to radiation concerns.) The temperature sensors are not visible in Figure 3.7, since they were not installed until after the vessels were shipped to LNGS.

Though the vessels have not yet had any fluids inserted as of this writing, the sensors are already producing useful results. They show that air temperatures remain up to $5^{\circ}C$

warmer at the top of the Stainless Steel Sphere compared to the bottom. The stratified temperatures thereby implied indicate that convection within Borexino fluids should not occur, so it may be possible to tag related radioactive coincidences at relatively long time scales, on the order of a few hours. Refer also to Section 8.4.

Finally, each end region has two $\frac{1}{2}$ "-diameter tubes attached that put the Inner Volume and Inner Buffer in potential communication with the external world. One is for fluid sampling purposes; the other is connected to a pressure gauge so that relative pressures between the three volumes may be measured. The Inner Vessel tubes are flexible nylon, while the Outer Vessel tubes are stainless steel.

To prevent leaks of possibly contaminated external air into the scintillator or buffer fluids, the end regions are designed so that all fluids are doubly contained. Therefore, the optical fibers, sampling ports, and read-out wiring of the temperature sensors and load cells are routed out of the end regions to the external world through a complicated double set of feed-throughs. Most of these feed-throughs were installed while the vessels themselves were being installed inside the Stainless Steel Sphere; see the following section for a description of the installation. Details on the feed-throughs may be found in reference [49].

3.2.4 Shipping and installation of the vessels

The completed nylon vessels were sprayed with many liters of deionized water so they would retain moisture during shipping. They were next nested inside two layers of precisioncleaned polyethylene film and one layer of metallized film. Finally, the covered vessels were folded back on themselves several times. The middle, folded, portion of the vessel bundle was hung inside a plastic cage, or so-called "film cassette." The north and south end regions were suspended from a scaffolding in order to be packaged for shipping (Figure 3.8). They were shipped from Princeton University on August 6, 2002, arriving at the Italian underground laboratory about one week later.



Figure 3.8: The Borexino nylon vessels, sealed, folded back upon themselves, and attached to a scaffolding for ease in shipping. The north end region is at right. In this photograph, dated August 2, 2002, they are leaving the clean room at Princeton University, and are about to be packaged in a wooden crate. The author is at bottom center of the photograph.

Unfortunately, on August 16, a small quantity of pseudocumene (~ 50ℓ) was accidentally drained into a small stream just outside the underground laboratory. The repercussions of this incident included legal difficulties that embroiled the laboratory for many months, and perception of a need to completely reconstruct the underground water drainage system [88]. As a result, the nylon vessels remained packaged in their original shipping container for almost two years, far longer than intended. Some discoloration of the nylon film occurred at the end regions during the long period of inactivity. This was thought to result from slight oxidation of copper support struts included in the Inner Vessel end regions, and is not believed to be a problem. The vessels were not completely installed within the Stainless Steel Sphere until March 26, 2004.

Installation of the vessels in the sphere was a complicated and challenging process. Some problems to overcome included the weight of the end regions (total mass on the order of 50 kg at each pole), the 13-m length of the film panels ($\pi/2$ times longer than the eventual diameter of the inflated vessels), and the fragile PMTs that were already installed everywhere on the sphere's inner surface except near the bottom. The description here has been synthesized from reference [89] and the personal experience of the author. At the time of the procedure, access to the Stainless Steel Sphere was controlled through a Class 100 clean room "CR 1" at the base of the detector, from which a temporary hallway passed through the Water Tank. A ramp led from the end of this hallway into a 3-m diameter aperture near the bottom of the sphere. Inside the hallway and sphere, scaffolding had already been constructed. An elevated horizontal I-beam was attached at the top of the scaffolding along the length of the hallway; it acted as a track along which several sets of wheels were constrained to roll. Much of the first part of vessel installation involved suspension of the vessels from these wheels.

To begin with, the vessels, still bundled up inside the frame in which they had been shipped, were moved into CR 1. They were then transferred from the original shipping frame onto the CR 1 scaffolding: first the north end region, then the film cassette, and finally the south end region were suspended from wheels on the I-beam track. The two end regions were suspended horizontally, parallel to the track. During this stage, additional hardware, the tube assemblies, was attached at each end region (Figure 3.9); the purpose of these is to carry tubing and instrumentation wiring from the Outer Vessel end regions through the Water Tank to the external world. The north end region was then pulled up the ramp and hoisted, with the aid of various winches, up through the scaffolding inside the sphere. It was transferred from winch to winch, in the process being rotated through 90° until it hung vertically. The transfers reached a conclusion as the end region finally was suspended from a single rope leading up through the top of the sphere (13.7 m height) to a winch near the top of the Water Tank. As the vessels followed after the north end region, they were unpeeled from the three layers of film protecting them, and were meanwhile continually kept moist by spraying with deionized water.

At this point, with the north end region suspended about $\frac{1}{3}$ of the way to its destination at the top of the Stainless Steel Sphere (Figure 3.10), a complicating factor was discovered.



Figure 3.9: The Borexino vessels, unpacked from their shipping frame and suspended from the scaffolding in CR 1 at LNGS. The south end regions are in the foreground near the center of the photograph. The south tube assembly (on the floor at bottom right) is about to be attached. At the back of the room on the left is the entrance to the temporary hallway through the Water Tank, leading to the Stainless Steel Sphere 3-m aperture.



Figure 3.10: The vessels, with the north end regions having been raised about $\frac{1}{3}$ of the way to the top of the Stainless Steel Sphere. The scaffolding used to rotate the north end region to its vertical position is clearly visible. This photograph was taken with the Borexino camera system.



Figure 3.11: Delicate surgery on the load cell attachments at the Inner Vessel north end region. The bolts of the attachments needed to be secured so they could not come loose. This necessitated opening up the Outer Vessel.

One of the bolts holding a load cell in place had come unfastened, most likely as a result of vibrations during the shipping of the vessels. After returning the errant bolt to its position, the installation team, led by C. Galbiati and A. Pocar, decided to solve the problem in a more permanent way. Each bolt had a small hole in its head. A Tensylon string was threaded through the hole and looped around the load cell in such a way that the bolt could make no more than one full rotation. This process had to be performed for load cells of both the Inner and Outer Vessels. To reach the Inner Vessel load cells, 30-cm slits were cut in the Outer Vessel (one at the south end, two at the north end). The task was

particularly delicate for the Inner Vessel load cells at the north end region, which was now hanging inside the sphere (Figure 3.11); an accidentally dropped piece of hardware here might fall onto the Inner Vessel bridge film, puncturing it. Happily, the work was finished without mishap. The slits in the Outer Vessel were patched by the standard procedure of gluing and clamping. The procedure was more difficult than when performed on the flat tables in the Princeton clean room; special clamping fixtures were built and hung from the scaffolding.

The next step was to raise the north end regions to the top of the sphere. This was done with relatively little difficulty. Meanwhile, the rest of the vessel, still suspended from the I-beam track, was advanced forward through the 3-m opening. The film cassette was disassembled at the same time. The remainder of the vessel film could not all be allowed to hang freely from the north end regions, however; in its still-folded state, it would drag on the sphere floor, being $\pi/2$ times longer than the inflated Outer Vessel diameter. The excess length of film was stored on a platform atop the scaffolding in the sphere. When enough had been gathered up, the south end regions were hoisted upward through the scaffolding in the sphere bottom.

Finally, over a period of several days, the vessels were inflated from the top with synthetic air (a mixture of high-purity oxygen and nitrogen). As inflation progressed, the bundled-up section of film naturally rose, and the scaffolding in the Stainless Steel Sphere was disassembled in stages (Figure 3.12). A raised plastic railing and steel structure were installed around the Outer Vessel south end region (Figure 3.13) in order to prevent the hold-down ropes attached there from dragging onto the bottom of the sphere, possibly becoming entangled with the PMTs.

Following the inflation, the vessels were purged with high-purity nitrogen, the regular supply for the underground laboratory, for several cycles, until the concentrations of argon and krypton inside them reached the limiting values in the nitrogen. The respective concentrations of those gases in the high-purity nitrogen are 11 ppm and 40 ppt by volume [43],



Figure 3.12: Disassembly of the Stainless Steel Sphere scaffolding begins as the inflation of the vessels reaches a critical point. Several of the belts on both Inner and Outer Vessels can be seen. Photograph taken with the Borexino camera system.



Figure 3.13: The south end region of the Outer Vessel, seen as the process of vessel inflation nears its end. The two sets of Outer Vessel ropes are visible: the set acting against upward buoyant force, affixed near the bottom of the stainless steel pipe, and the set acting against downward forces, threaded through the L-ring and stretched taut over the Outer Vessel near the top of the photo.

corresponding to respective specific activities for the isotopes ³⁹Ar and ⁸⁵Kr at STP of 15 and $40 \,\mu\text{Bq/m}^3$. This is still far too high, as the requirement for the Borexino scintillator is about $0.1 \,\mu\text{Bq/m}^3$ for each isotope. Further purging cycles will therefore occur once special nitrogen, low in argon and krypton (LAKN₂) is available; refer to Section 2.2.2 for more information.

3.3 Photomultiplier tubes

3.3.1 The inner detector

The interior of the Stainless Steel Sphere supports 2240 inward-facing photomultiplier tubes (PMTs), arrayed in 42 rings (Figure 3.10). The tubes are the 8" model ETL 9351 produced by Electron Tubes Limited (formerly EMI). The model has a high quantum efficiency (the probability that an incoming photon is converted to a photoelectron) of 26.5% at a wavelength of 420 nm [68], with a broad spectral dependence shown in Figure 3.2. However, the ETL 9351 has a collection efficiency (the probability that a photon emitted from the photocathode is collected by the first dynode) of only about 60% [90], for an overall efficiency of 16%. The duration of a PMT signal caused by a single photoelectron is approximately 15 ns [91]. Detailed information about the PMTs is provided in references [68, 92, 93].

The tubes consist of low-radioactivity Schott 8246 glass. The base of each is encased by a cylindrical stainless steel housing, attached with a pseudocumene-resistant epoxy. A steel feed-through welded to the base carries the PMT voltage supply and signal cable through the Stainless Steel Sphere into the Water Tank. Each of these feed-throughs was leak-checked individually. The photocathode of each PMT is shielded from the Earth's magnetic field by a foil of μ -metal, a nickel-iron alloy, formed into a conical shape. Without it, PMT characteristics such as the peak-to-valley ratio would be degraded on the order of ten percent [94]. The μ -metal, which is chemically incompatible with pseudocumene, is covered with a thin layer of clear phenolic paint.



Figure 3.14: One of the Borexino inner detector photomultiplier tubes. From left to right, the components visible in this photograph are the aluminum light concentrator; the μ -metal magnetic shield; the stainless steel housing; and the feed-through for the high-voltage electrical connection.

Of the PMTs, about 80% make up the inner or neutrino detector. They are equipped with so-called "light concentrators" (Figure 3.14), curved aluminum surfaces that are roughly parabolic in cross-section. Light concentrators are designed according to principles detailed in [95, 96] such that light coming from within the Inner Vessel that hits the inner surface of the light concentrator is nearly always reflected into the photocathode. On the other hand, light originating outside the solid angle subtended by the Inner Vessel (as seen by the photocathode) that hits the light concentrator is reflected back out with high probability. Therefore the light concentrators act both to increase the effective photocathode coverage of the Stainless Steel Sphere (at a cost much less than that of installing more PMTs), and partially to filter out unwanted light in the buffer regions produced, for instance, by muons crossing through the detector.

The total optical coverage of the PMTs with light concentrators for a scintillation event occurring at the center of the Inner Vessel is 32%, taking into account the 90% reflectivity of the light concentrators [68]. Considering the scintillator light yield, detector optical coverage, and PMT efficiency, one would expect the conversion from β particle kinetic energy to photoelectrons to be about 590 photoelectrons/MeV. However, additional factors such as light absorption by pseudocumene reduce this value to 400 photoelectrons/MeV [55].

The PMTs and light concentrators are unfortunately, and unavoidably, the most radioactive components within the Stainless Steel Sphere. The Schott glass used is a factor of 10 better in radiopurity than ordinary glass, but it still has concentrations of 30 ppb U, 10 ppb Th, and 20 ppm K by mass [69]. The main danger is the 2.615 MeV γ -ray emitting isotope ²⁰⁸Tl in the thorium decay chain. Estimates of the resulting event rate within the Borexino Fiducial Volume, determined by Monte Carlo methods, are 0.6 events/day from the PMTs in the 250–800 keV energy range, and another 0.6 events/day in the same energy range caused by the light concentrators [49]. These event rates increase rapidly if the volume of scintillator taken as the Fiducial Volume is expanded.

3.3.2 The Muon Veto System

Recall that approximately 5000 muons are expected to pass through the Borexino detector per day. In order not to confuse these with neutrino events, at most one muon event per day may contaminate the neutrino data sample, implying that it is necessary to veto them with an efficiency of roughly 99.98%. This is done with the aid of two Muon Veto Systems, described further in references [37, 38, 97, 98].

Roughly 400 PMTs of the 2240 inside the Stainless Steel Sphere do not have light concentrators. Collectively they make up an Inner Muon Veto System. Since they have an unrestricted field of view, the number of photoelectrons collected by these PMTs divided by the number collected by the PMTs with light concentrators will be larger for an event in the buffer fluid than for an event within the Inner Volume. This light collection ratio can be used after the fact in reconstruction or analysis software to exclude Čerenkov light produced by muons from the data sample. (A high-energy muon that travels through the detector will of necessity spend at least 38% of its time in the buffer regions.)

On the outside surface of the Stainless Steel Sphere, pointing outwards, and on the floor of the Water Tank, pointing upwards, are located an additional 208 PMTs. These make up the Outer Muon Veto System. These PMTs are the same ETL 9351 model as the inner detector tubes. Unlike the PMTs inside the sphere, each outer PMT is completely encapsulated from the surrounding water in a stainless steel housing, which has a glass window in front of the photocathode. A special electronics subsystem is set up for data acquisition on the Outer Muon Veto System. This subsystem has the capability to abort data acquisition in the inner detector immediately following detection of a muon track in the Water Tank.

Since the ultra-pure water found in this volume does not scintillate, the only near-ultraviolet light that can be observed by these PMTs is Čerenkov radiation. They are therefore in an ideal position to detect muons. In order to increase the probability of observing Čerenkov photons from a muon track, large parts of the outside of the sphere and the inside of the Water Tank are lined with highly reflective Tyvek sheeting.

A muon which has a short path length within the Stainless Steel Sphere deposits a relatively low amount of energy, and therefore is most in danger of falling within the neutrino energy window. The nature of such an event, however, makes it most detectable by the Inner and Outer Muon Veto Systems. Since it travels only near the edge of the inner detector, it is picked up with high efficiency by the inner muon system. Furthermore, a short path length in the inner detector implies a long path length within the Water Tank, so such a muon track is more easily observed by the outer muon system. The combination of these effects, together with other indicators such as the time delay and time width of an event, is expected to yield the required muon veto efficiency [37, 97].

3.4 Data acquisition

The Borexino detector data acquisition system is designed to perform four functions when an event of a programmable minimum energy occurs within the detector:

• Measure the number of photoelectrons detected by each PMT;
- Measure the time at which each PMT saw a photoelectron;
- Measure the sum over all PMT signals as a function of time (the "pulse shape");
- Determine whether the event was caused by a muon passing through the detector.

The first two of these data are used to determine the position and energy of the event. The position is required to decide whether or not it is in the Fiducial Volume. The energy is necessary for more obvious reasons. The pulse shape is used to decide whether an event is a β -like signal that could be a neutrino, or an α -like event that must be a radioactive decay. Finally, the presence of muons is important both because they may be mistaken for neutrino events if their path length within the Stainless Steel Sphere is short enough; and because a muon may produce radioactive nuclides (cosmogenics) that do not decay until some time afterward. Naturally, all of these results have finite accuracy. For reference, the expected energy resolution of a 1 MeV β -like event is about $\sigma_E/E = 7\%$. (A perfect detector would have 5% energy resolution at that energy; we presume some inefficiency, as detailed in Section 2.1.) The expected position resolution (1 σ) is about $\sigma_x = 8.0 \text{ cm}$ [55].

The data acquisition system can be subdivided logically into the electronics boards (hardware), the computers that run software controlling them (the "producers") that sends data over a network, and the computers receiving these data (the "consumers") that store the data to disk. These subsystems of data acquisition on the set of 2240 inner PMTs are covered in detail in references [68, 92, 99, 100]. For data acquisition on the 208 outer PMTs of the Outer Muon Veto System, information may be found in [22]. The following summary, shown in schematic form by Figure 3.15, is only a skeleton outline of the flow of data from PMT to hard disk, and it oversimplifies many things.

The high-voltage cables that carry signal read-outs from the PMTs feed into a set of analog Front End boards. These boards, described in reference [91], have the task of decoupling the power supply and signal, both of which travel over a single PMT cable. Additionally, each board performs analog integration of the PMT signal over time. For each PMT, a



Figure 3.15: A schematic of the flow of data within the Borexino data acquisition system. The Front End boards are represented in green, the digital boards in blue, and the Borexino Trigger Board in turquoise. The red HV Mainframe is the power supply to the PMTs. The yellow boxes are the VME crate controllers. Taken from reference [92].

board provides an amplified copy of the PMT signal and a copy of the time-integrated signal as output. In case of very high-energy events (such as might be caused by supernova neutrinos), the board also outputs an analog sum of all channels.

The analog signals from the Front End boards are fed into a system of digital electronics. The digital boards, described in reference [101], are manufactured by Laben S.p.A. and designed for the VME bus. Each VME crate can hold 20 boards, with each board capable of handling eight independent electronic channels (pairs of input signals). Eight channels per crate are unused, giving a total of 160 PMTs handled per crate, for 14 crates in all. For each electronic channel, the analog PMT signal first passes to a discriminator that performs two tests: first, whether the signal crosses a programmable threshold (currently set to $\frac{1}{4}$ of the amplitude of a single photoelectron signal), in order to veto dark noise of the PMTs; second, whether the signal is an increasing function at that moment. If both conditions

are met, the discriminator fires. The time of the firing is digitized by a Time to Digital Converter (TDC) with a resolution of 0.5 ns, and stored to a circular memory buffer on the board. The incoming time-integrated PMT signal (which is delayed slightly relative to the other signal) is sampled at the time of the discriminator firing, and again after an 80-ns delay. The difference between the sampled values, which is simply the accumulated charge on the PMT (proportional to the number of photoelectrons), is digitized by an Analog to Digital Converter (ADC) with 8-bit resolution and stored to the memory buffer. Very highenergy events could saturate these 8-bit values, hence the need for the analog sum output of the Front End boards. The memory buffer stores data for $7 \,\mu$ s, after which it is overwritten by newer data.

Once a discriminator has fired, it is purposely disabled for 140 ns in order to avoid a race condition between sampling and data storage. Therefore, timing data for a second photoelectron detected by a PMT less than 80 ns after the first one are lost, but its existence is disclosed by the fact that the value recorded for the charge is about twice that expected for a single photoelectron. A second photoelectron that arrives between 80 and 140 ns after the first one will be completely lost. Even for events at the high end of the neutrino energy window, on average only 7% of hit PMTs will see more than one photoelectron, and nearly all additional photoelectrons will be detected within 30 ns of the first one. Furthermore, even if two events occur within a 140-ns time window, the number of channels temporarily disabled by the first event (approximately 320 for events at the high end of the neutrino window) will be small compared to the total number of PMTs. The second event can still be detected, albeit with a slightly reduced number of photoelectrons. This dead time is therefore expected to have negligible effect on the data sample acquired by the detector.

A number of trigger boards sum together the number of times that a discriminator on a timing input of the digital electronics has fired. This sum eventually is fed to a master board, the Borexino Trigger Board (BTB). Every $16^{2}/_{3}$ ns, the BTB determines whether the number of photoelectron hits within the previous 50 ns is greater than a programmable threshold value. If the threshold is set sufficiently high, the data acquisition system will

ignore low-energy events, the vast majority of which are uninteresting ¹⁴C decays, and they will be automatically discarded before they can ever consume disk space.

If the sum boards do exceed the threshold of the BTB, it sends a trigger to all digital boards to flag the event. Each board, upon receiving the trigger, sends a signal on its VME crate bus. A program running under the operating system of the VME crate controller saves the 7- μ s contents of the boards' memory buffers, beginning with the data causing the initial trigger, to RAM. A second process running concurrently, which may be termed the producer, builds the data in RAM into a data structure, also called an event. Unfortunately, the process of event building itself takes several μ s, so after an event has been built, the BTB trigger generator is disabled for a period of 10 μ s. Therefore any events occurring between 7 and 17 μ s after the initial event are invisible. This dead time also should not impact the ability of the detector to see most coincidences due to radioactive decays. The mean coincidence time of the low-probability ⁸⁵Kr decay to ^{85m}Rb is 1.5 μ s (\ll 7 μ s), while that of the ²¹⁴BiPo coincidence is 237 μ s (\gg 17 μ s). Any *third* event that occurs within 100 μ s of the first is also ignored, due to the inability of the digital boards to handle more than two pending data reads, but the impact of this dead time should be even more negligible.

Both the VME crate controllers (the MVME 2302 model, designed by Motorola) and the data receiving workstation (a standard computer with Intel-compatible CPU) run the Debian distribution of the GNU/Linux operating system. The operating system of the VME crate controllers has been customized specifically to fit their environment; the workstation has a more-or-less standard setup. Once an event is built in RAM, each VME crate controller sends its part of the event across the network to the workstation. There, finally, the complete event is saved to disk. The data acquisition may be controlled via a Web-based interface (Figure 3.16) accessible only to local operators.

The Outer Muon Veto System is largely independent of this data flow. It, too, has a number of analog boards serving as Front End boards and digital boards that convert the signals of the PMTs in the Water Tank to digital values. When a sufficiently large number of



Figure 3.16: The control interface to the Borexino data acquisition system. This screen of the interface shows a map of the status of the inner detector PMTs at the moment the screenshot was taken. Taken from reference [92].

Water Tank PMTs are hit within a short time-span, the muon system electronics send a flag directly to the Borexino Trigger Board. The presence of this flag is saved in the data output.

When the total charge and the hit time is known for every PMT that detected a photoelectron during an event, the spatial position and the time of a neutrino interaction or radioactive decay within the scintillator can be determined. This determination is referred to as event reconstruction. The official reconstruction software of Borexino, known as Echidna, is a modular code based in C++. It is intended to convert raw charge and timing information into a list of events with definite position and total energy. These will be stored in a format readable by the popular data analysis software ROOT [102], developed at CERN, and distributed to all members of the Borexino Collaboration.

Chapter 4

Nylon Film of the Borexino Vessels

As discussed in the previous chapter, the two Borexino vessels are among the most critical parts of the experimental design. The Outer Vessel acts as a barrier to dissolved radon gas that might be produced by radioisotopes in the photomultiplier system and Stainless Steel Sphere. The Inner Vessel, in addition to presenting a second barrier to radon entering the Fiducial Volume, also prevents the mixing of scintillator with buffer fluid. These vessels must neither leak nor allow radon to diffuse through; either possibility would raise the background signal rate in Borexino sufficiently to prevent meaningful neutrino observations. As such, there are a number of stringent requirements that the material composing them must meet.

First, the vessels must be chemically compatible with the fluids that will be used in Borexino: pseudocumene scintillator and buffer solutions, and possibly also a preliminary filling with water. Needless to say, they must also be transparent at the wavelength of the scintillation light produced within the Inner Vessel. A thin nylon film meets both of these requirements. The good transmittance and low light scattering of the nylon types selected for Borexino, as well as the close match between their index of refraction (n = 1.53) with that of pseudocumene (n = 1.505), has been thoroughly detailed elsewhere [12]. Details of the chemical properties of nylons will be discussed below.

To prevent mechanical failure, the nylon film making up the vessels must by design be able to withstand a temperature gradient of up to 5°C between any two of the three volumes inside

the Stainless Steel Sphere. The important issue is not the temperature difference itself, but the difference in buoyancy it would cause between the different volumes of scintillator and buffer fluid. Such a temperature gradient has been calculated to result in stress on the film of up to 20 MPa [12]. Furthermore, stresses at this level may cause a phenomenon called creep, an irreversible stretching of the film that would lead to deformation of the vessels.

The film must also not be brittle enough to crack during complicated operations such as vessel construction and installation. It is a major concern that a Borexino prototype nylon vessel installed in the Counting Test Facility experienced mechanical failure in December 1999. Due to the inherent geometrical difficulties entailed by building spherical vessels from a flat sheet of plastic, the nylon vessels are constructed from a large number of separate panels. During their construction and shipping, a great deal of folding and creasing occurred. Additional stresses were put on them when the folded vessels were installed into Borexino and inflated into spheres. In both cases, two intersecting folds in the film, which yield a point, might potentially have caused a puncture to form in the nylon. Cracking along a crease, which appears to have been the cause of failure in the CTF, was another possibility.

For these two reasons, it was vitally important for the vessel building group at Princeton to understand the material properties of plastics such as nylon. Parts of this chapter will describe how varying the relative humidity of the environment at a constant temperature affects the material properties of two different types of nylon, and whether the type of environment (air versus pseudocumene) is relevant. From these results, it was possible to predict how the type of nylon used, the thickness of the nylon film, and the humidity of the film's environment during construction, installation, and data acquisition would affect the material properties of the Borexino vessels.

Finally, the vessels must act as effective barriers against radon diffusion. To prevent radon emanation from the film itself, they must also have very low levels of intrinsic radioactive impurities. Several studies have examined radon diffusion through nylon, and from that model have been able to determine the level of the radon progenitor 226 Ra in the films as

well. As a complicating factor, nylon becomes much more permeable to radon when very humid.

4.1 Chemical properties of nylon

The word "nylon" refers to a large group of polymers that are built from carboxylic acids¹ and amines.² They are therefore a type of polyamide.³ Many members of the nylon family have the properties of chemical compatibility with pseudocumene; high light transmittance as a thin film; and sufficient material strength to make them good candidates for use as the Borexino vessels.

4.1.1 Nomenclature and structure

Nylons may be grouped into two major families depending upon their monomers. The AB family is produced from amino acid⁴ monomers, while nylons in the AABB family consist of monomers comprising a dicarboxylic acid joined to a diamine. In brief, AB nylons result

¹carboxylic acid: an organic molecule on which at least one carbon is singly bonded to a hydroxy

⁽⁻OH) group and doubly bonded to another oxygen (=O). That is, $\begin{array}{c} O \\ C \\ R \end{array}$, or as usually written

inline to save space, R–COOH. The molecule is an acid because the hydrogen of the hydroxy group may dissociate: R–COOH $\stackrel{\text{H}_2\text{O}}{\rightleftharpoons}$ R–COO⁻ + H⁺.

A note on notation: Most organic molecules can be thought of as an undistinguished mass of carbon atoms with chemically reactive "functional groups" sticking out. To a first approximation, the chemical properties of the molecule depend upon the functional groups. When the exact arrangement of the carbon skeleton is unimportant, it may be abbreviated as an "R group" with the letter R, as above.

²**amine**: an organic molecule which has at least one $-NH_2$ functional group. The molecule is a base since it may pick up a hydrogen: $R-NH_2 + H^+ \stackrel{H_2O}{\leftarrow} R-NH_3^+$.

³amide: an organic molecule which has separate sets of carbon atoms linked by an amide group, consisting of a nitrogen atom and carbonyl (C=O). That is, R-NHCO-R'.

⁴**amino acid**: an organic molecule with a carboxylic acid group and an amine group. That is, $H_2N-R-COOH$. Amino acids are well-known as the components of the proteins found in all living organisms.

from the overall reaction

$$n \operatorname{H}_2 \operatorname{N} - \operatorname{R} - \operatorname{COOH} \longrightarrow \operatorname{H}_{--} \operatorname{HN} - \operatorname{R} - \operatorname{CO}_{-} \operatorname{In} \operatorname{OH} + (n-1) \operatorname{H}_2 \operatorname{O}$$
(4.1)

while AABB nylons are produced via

$$n \operatorname{H}_{2}\mathrm{N} - \mathrm{R} - \mathrm{N}\mathrm{H}_{2} + n \operatorname{HOOC} - \mathrm{R}' - \operatorname{COOH} \longrightarrow$$
$$\mathrm{H}_{--}[\mathrm{HN} - \mathrm{R} - \mathrm{N}\mathrm{HCO} - \mathrm{R}' - \mathrm{CO}_{-}]_{n} \operatorname{OH} + (2n-1) \operatorname{H}_{2}\mathrm{O}. \tag{4.2}$$

As an additional complication, in some cases the product of these reactions may, with the release of water, form a closed ring; one end of the polymer attaches itself back to the other.

Several types of notation may be used to describe nylons. As described in reference [103], the one most frequently used is the PA system.⁵ In this system, AB nylons are labeled PA-x, with x representing the amino acid monomer. AABB nylons are labeled PA-xy, with x representing the diamine portion (the R reactant of Equation (4.2)) and y representing the dicarboxylic acid portion (the R' reactant) of the monomer. Most nylon monomers are derivatives of straight-chain alkanes⁶ with the functional groups at each end of the chain. These are designated very simply by letting x or y be the number of carbon atoms in each component. For instance,

- PA-6: an AB nylon with formula $H \longrightarrow HN(CH_2)_5 CO \longrightarrow OH$;
- PA-66: an AABB nylon with formula $H HN(CH_2)_6 NHCO(CH_2)_4 CO OH;$
- PA-612: an AABB nylon with formula $H = HN(CH_2)_6 NHCO(CH_2)_{10}CO = OH$.

In place of "PA", often the word "nylon" is used: nylon-6 or nylon-66. When spoken, the pronunciation is in accord with the monomers: "nylon six-six" (not "nylon sixty-six"), "nylon six-twelve," etc.

⁵PA stands for "polyamide."

⁶alkane: a hydrocarbon with only single bonds. A straight-chain alkane has all carbon atoms laid out in a row, *e. g.*, $CH_3CH_2CH_2CH_3$ (*n*-butane).

Some monomers or monomer components, of course, are not straight-chain alkane derivatives; in this case the symbols used in the PA nomenclature are somewhat arbitrary. For instance, isophthalic acid and terephthalic acid, benzene derivatives with the following respective chemical structures, are symbolized as "I" and "T".



Hence the AABB nylon with monomers produced from terephthalic acid and the diamine $H_2N(CH_2)_6NH_2$ would be termed PA-6T, or nylon-6T (shown below).



So far all the examples have been *homopolymers*, made up of only one type of monomer. One may imagine a polymer built up from more than one kind of monomer; this is termed a *copolymer*. This is represented with a forward slash in the PA nomenclature; thus, a nylon made from monomers of both PA-6 and PA-6T would be the copolymer PA-6/6T. Note that the nomenclature does not give any specification of the relative amounts of each type of monomer, nor their order within the chain. Furthermore, copolymers should not be confused with *blends* — mixtures of different polymers. The blend PA-6/PA-66, consisting of chains of both PA-6 and PA-66 homopolymers, is a different creature than the copolymer PA-6/66, in which each chain is made of both PA-6 and PA-66 monomers.

A final contributor to the bulk properties of nylons is the number of monomers N making up an average polymer chain in the material. This is usually on the order of hundreds or thousands. It is impossible to produce nylon in bulk such that every chain in the material is the same length, so N is not a specific number, but instead has some distribution. Values commonly used to characterize this distribution include the *number-average degree* of polymerization $\bar{P}_n \equiv \sum_{N=1}^{\infty} n_N N$, and the mass-average degree of polymerization $\bar{P}_m \equiv$ $\sum_{N=1}^{\infty} m_N N$. Here, n_N and m_N are the mole fraction and mass fraction, respectively, of the species made of N monomers in the material.

 P_n is the first moment $\langle N \rangle$ of the number distribution of N. In the limit of a large average N (so that the mass of the H– and –OH groups at either end of a chain may be neglected), $\bar{P}_m \approx \sum_{N=1}^{\infty} n_N N^2 / \bar{P}_n = \langle N^2 \rangle / \langle N \rangle$, the ratio of the second moment to the first moment. Higher-order characteristic values for the number distribution may be defined similarly. Note that \bar{P}_n will tend toward the lower end of the distribution curve of n, while \bar{P}_m will tend toward the lower end of the broadness of the number distribution is therefore given by the *polydispersity index* (PDI), defined as $\bar{P}_m / \bar{P}_n \approx \langle N^2 \rangle / \langle N \rangle^2$.

As a trivial example, consider a monomer, dimer, and trimer of nylon-6. The molar mass of each species equals the number of monomers in the species, times the molar mass of the repeating unit (113 amu), plus 18 amu for the H– and –OH groups at either end. The number-average and mass-average degrees of polymerization for these three molecules are, respectively:

$$\begin{split} \bar{P}_n &= \frac{1}{3} \times 1 + \frac{1}{3} \times 2 + \frac{1}{3} \times 3 = 2 \\ \bar{P}_m &= \frac{131}{732} \times 1 + \frac{244}{732} \times 2 + \frac{357}{732} \times 3 \approx 2.31. \end{split}$$

The polydispersity index for this system is about 1.15. Even for these small values of N, the error involved in neglecting the masses of the end groups of the chains is negligible: $\langle N^2 \rangle / \langle N \rangle = 2.33$ and $\langle N^2 \rangle / \langle N \rangle^2 = 1.17$.

4.1.2 Crystallinity

The polymer chains of nylon may be categorized as varying from *crystalline* in one extreme, to *amorphous* in the other. A sample of nylon may be characterized by the percent crystallinity w^c , usually expressed by weight; a typical value for w^c is 50%. Amorphous and crystalline regions are distributed at random throughout the sample. A polymer chain may fold back on itself numerous times in the same crystal. Alternatively, it may begin in one crystalline region, traverse an amorphous region, and end in a different crystalline region. The crystalline and amorphous regions are therefore defined by their internal structures rather than the identity of the polymer chains within them.

In crystalline regions, polymer chains are arranged in a rigid structure, with hydrogen bonds connecting amide hydrogens on one chain to oxygen atoms in the amide groups on the adjacent chain. Two main crystalline structures, the α and γ forms (described in detail in [103]) have been well-characterized through X-ray diffraction techniques. The structure of amorphous nylon, on the other hand, is not well-understood. It is known that in amorphous regions, many amide groups are not cross-hydrogen-bonded. This allows water molecules dissolved in the nylon to form hydrogen bonds with these amide groups instead. This phenomenon results in a transition temperature, dependent upon the moisture content of the nylon, at which amorphous nylon changes its physical characteristics; this will be discussed further in Section 4.2.1.

Copolymers tend to be more amorphous than homopolymers, simply because irregularities in the polymer chains do not mesh together well. Common additives used to disturb the formation of crystalline regions in PA-6 and PA-66 are isophthalic acid and terephthalic acid, each containing a bulky and disruptive benzene ring. (Their molecular structures are depicted in Section 4.1.1.) The resulting nylon copolymer is entirely amorphous.

For the purposes of Borexino, completely amorphous nylons ($w^c = 0$) are more desirable because they have much better transparency (less "haze") than partly crystalline nylons. Nevertheless, crystals do not present a problem in thin films if the maximum size of crystalline regions is much smaller than the wavelength of the near-ultraviolet scintillation light (\sim 350 nm). Homopolymers may be created with such small crystals by very quick cooling ("quenching") from the melt state, for instance by extruding nylon film at relatively low temperatures, so that the crystals do not have time to grow [104].

4.1.3 Nylon hydrolysis

At high temperatures, the polymerization reactions Equations (4.1) and (4.2) are actually equilibrium reactions. Since they are condensation reactions, Le Chatelier's Principle predicts that the addition of water to nylon will push the reactions back toward the left. This would break the polymer chains down into monomers again with the release of water, a process called hydrolysis. Hydrolysis reduces the strength of nylon and causes it to become brittle as well. Fortunately, polymerization is also an exothermic process, so at room temperature the forward direction of the polymerization reaction is heavily favored.

Though nylon-6 and nylon-66 are among the species most susceptible to hydrolysis, it is not expected to be a problem in Borexino. Studies of nylon-66 soaked in water for some time indicates that it becomes brittle after about 2 months at 66°C [12]. Using the rule-ofthumb that reaction rates are slowed by a factor of two for a ten-degree temperature drop, hydrolysis would be predicted to become a problem at the ambient 10°C temperature of Borexino after about 8 years. This is far longer than the expected 6–12 month duration of a water-filling phase of detector operation, if it is even implemented.

4.1.4 Candidate materials for Borexino

Several different types of nylon film have been manufactured during the history of the Counting Test Facility and Borexino experiment. In 1992, C38F film was produced from pellets of Durethan C38F, a nylon-6 copolymer with proprietary formula manufactured by Bayer Corporation [105]. Films of 125, 100, and 500 μ m thicknesses were extruded at the Miles-Mobay plant (also owned by Bayer) in Pittsburgh, Pennsylvania, USA. The 500 μ m-thick film was used for the CTF vessel in CTF 1, due to the need of supporting the large buoyant load resulting from suspension of a pseudocumene-filled vessel in water. The RMS deviation from this thickness was about 30 μ m. To construct the CTF 2 vessel, the same C38F pellets were extruded in a 500 μ m thickness at a different location, the New Jersey Institute of Technology's Polymer Processing Institute. No plants willing to produce 500 μ m-thick film could be found at the time of production of the CTF 3 vessel, so it was constructed from leftover Miles-Mobay C38F film.

In 1998, three different films were extruded for testing purposes. In each case, pellets of nylon were purchased from the producers, and extrusion was done at American Leistriz, a company in New Jersey, USA [12]. They were as follows:

- Capron (also called C100) was extruded from Capron B73ZP⁷ pellets made by Allied-Signal/Honeywell [107]. These pellets are a homopolymer of nylon-6, kept transparent by quenching at extrusion time.
- **C90** film is a blend of Capron B73ZP pellets (90%) with Selar PA3426 (10%). Selar PA3426 (the "PA3426" designation is a product number, not a PA-system nomenclature) is an amorphous nylon sold by DuPont [108] to be blended with nylon-6 for preventing crystallization. It was noted that a concentration of Selar higher than 10% made the film unstable [12].
- **Sniamid** extruded at Leistriz was made from Sniamid ADS40T pellets manufactured by Nyltech. Like C38F, ADS40T is a proprietary nylon-6 copolymer.

⁷The 'B' in B73ZP indicates nylon-6. The number "73" refers to the viscosity. 'Z' means that the pellets had been screened and dusted, and 'P' that they had been washed to make the material "food-grade" [106].

In the rest of this chapter, these materials will be referred to as Capron-Leistriz, C90-Leistriz, and Sniamid-Leistriz in order to distinguish them from the materials used in the Borexino vessels.

Two types of nylon film were manufactured for use in Borexino:

- **Capron** was used in the Borexino Outer Vessel. It was extruded [106, 109] from Capron B73ZP pellets (described above) in Pottsville, Pennsylvania, USA at the Allied Signal plant in 2001. The percentage of haze measured in the just-extruded film was 0.5-1.5% (the latter after being soaked in water for 12 hours) [106]. The film has an average thickness of $125 \pm 8 \,\mu$ m.
- Sniamid was used for the Borexino Inner Vessel. It is a mixture [110] of 83% Sniamid ADS40T pellets (described above) and 17% Ultramid B4⁸ pellets by BASF [112], extruded at the mf-folien⁹ plant [113] in Germany in 2001. B4 is a nylon-6 homopolymer with an average molecular mass of 4×10^4 amu [111], implying an average degree of polymerization of 350. It is not clear whether this average is the value for \bar{P}_n or \bar{P}_m .

After extrusion, both of these films were sent to CleanFilm, Inc. in Islandia, New York, USA for surface cleaning. Their method of cleaning used ultrasonic methods to loosen particles, which were then removed from the surface by suction. The films were then certified by CleanFilm as meeting Level 25 of Military Standard 1246C [114]. This level of cleanliness corresponds to no more than 3 mg of particulate matter distributed over the entire inner surface of the Inner Vessel [44]. These materials will be referred to simply as Capron and Sniamid throughout the rest of this chapter.

For future reference, it should also be noted that following a series of corporate mergers and spinoffs, Nyltech, Allied Signal, and the nylon business of BASF are all currently part of Honeywell Nylon Incorporated.

⁸The 'B' in "B4" indicates nylon-6, while the '4' is a reference to the viscosity (greater, implying a higher molecular mass, than that of Ultramid B3) [111].

⁹The company uses only lower-case letters in its name.

4.2 Polymer mechanical properties

Amorphous nylons have a temperature T_g below the melting point at which a second-order phase transition occurs [115]. This is known variously as the glass transition or α relaxation. For nylon-6, the melting point is at ~220°C (although it may be supercooled to as low as 170°C before freezing) [103]. The glass transition temperature is in the range 0°-50°C for amorphous nylon-6, depending upon the moisture content of the sample.

Above T_g , a plastic is relatively weak, but rubbery and pliable. At lower temperatures it is stronger, but glassy and rather brittle. The glass transition temperature is dependent upon the exact composition of the plastic. Since plastics can absorb moisture from their environment, a change in the relative humidity of the plastic's environment will also affect the value of T_g . In fact, higher relative humidities generally imply lower transition temperatures. Furthermore, it has been shown that nylon which is water-saturated (and therefore at a temperature above T_g) has a much higher radon diffusivity than dry nylon [12]. Therefore the moisture content in the Borexino scintillator could have a big impact on the operational behavior of the Borexino nylon vessels.

4.2.1 Glass transition in nylons

An idealized polymer, whose molecular chains consist of a series of monomers linked by single bonds, may exist in four different states [115]. At sufficiently high temperature, the substance is molten, and the polymer chains rotate and slip past one another freely. When cooled below the melting point, the polymer becomes cross-linked by hydrogen bonds formed by water molecules between adjacent chains. However, the water molecules have some freedom of movement, so the chains may still move and flex into different conformations to some extent. Although now solid, the polymer is soft and pliable; it can be stretched and deformed easily. This phase of the polymer is referred to as the rubbery state. Further cooling prevents this breaking and reforming of hydrogen bonds, so the functional groups on a chain are now semi-permanently fixed to other substituents on an adjacent chain. The polymer chains become trapped in their configurations at the time of cooling. In general these configurations are random and chaotic. This is known as the glassy state of the polymer. The transition between rubbery and glassy states is a second-order phase transition, which occurs at a specific temperature T_g and may be measured with standard calorimetric methods. (As the transition is second-order, however, the change in physical properties on either side of T_g is gradual.) An additional possible phase is a crystalline phase in which all polymer chains are stacked in a regular lattice. It is very difficult to crystallize long polymer chains, so as noted in Section 4.1.2, a real plastic may contain crystalline domains along with regions of polymer in the glassy or rubbery states. Or, as is the case with the types of nylon considered here, it may be completely amorphous.

The glass transition temperature T_g is affected by a number of factors, including the average number of monomers in each polymer chain and the concentration of other substances dissolved in the polymer. For instance, a polymer composed of short chains will have a lower energetic barrier to conformation changes than a polymer made up of long chains of the same monomer. The short chain polymer thus has a lower glass transition temperature. This may be expressed by the semi-empirical formula [115]

$$T_g(n) = T_g(\infty) - \frac{C_g}{nM},\tag{4.3}$$

where $T_g(n)$ is the glass transition temperature of a polymer chain of n monomers; M is the molecular mass of a monomer; $T_g(\infty)$ is the glass transition temperature for an infinitely long chain of the same monomer; and C_g is an empirical parameter depending upon the monomer.

Nylons often contain plasticizers, substances which when dissolved in the polymer reduce its glass transition temperature. This makes the polymer less brittle and easier to handle during the manufacturing process. So-called copolymers, for instance, have different kinds of monomers as plasticizers. Empirically, the glass transition temperature of the mixture is roughly equal to a weighted average of the glass transition temperature of each constituent [115]. However, the weighting is not necessarily equivalent to the weight fraction of each component; it must be determined experimentally for the specific polymers used in the mix.

Water is also a plasticizer, and some empirical laws have been suggested for the dependence of T_g upon the moisture content of a polymer. These laws have two drawbacks: they are valid only in limited ranges of water concentration, and they depend upon empirical parameters which, again, must be measured for each polymer. It is therefore more useful simply to make measurements relating the glass transition temperature to the moisture content of the desired plastic. The results of a small experiment performed for the Princeton Borexino group are of interest in this context [116]. A differential scanning calorimeter was used to make measurements of T_g for a 1 mg sample of 125 μ m thick nylon-6 film dried in a silica gel bottle for one day, and for a similar sample soaked in water for several hours. The value of T_g was found to be 52°C for the dry nylon, and -6°C for the wet nylon. This demonstrates the existence of a value for moisture content at which T_g is equal to room temperature.

4.2.2 Tensile strength, Young modulus, and creep

Two major classes of mechanical properties, strength and brittleness, are relevant to the integrity of the Borexino vessels. Strength is the more straightforward to test; it can be quantified either by the tensile strength σ_t or by the Young modulus E. Measuring these properties for a sample of nylon film requires the construction of a graph of applied stress as a function of the film's fractional elongation ("strain").

To produce this graph, a long, narrow strip of film of length L is pulled from either end in a process of controlled elongation. This can be done using a device called a Tinius-Olsen machine, to be described in Section 4.3.3. At each moment, both the elongation of the film and the applied force are measured. Plotting the stress σ (force per unit cross-sectional



Figure 4.1: Diagrams of stress-strain relationships for nylon films below the glass transition temperature (left) and above it (right).

area) on the y-axis and the strain $\epsilon \equiv \Delta L/L$ on the x-axis yields a diagram something like those shown in Figure 4.1.

The tensile strength of the film σ_t is simply calculated as the maximum force applied divided by the film's initial cross-sectional area. (This is to be distinguished from the "tensile strength at break," which is defined using the force applied at the time of mechanical failure [117]; or the "yield strength," which is defined using the force applied at the first point on the stress-strain plot where $d\sigma/d\epsilon = 0$ [118]; these three values are not all necessarily identical.)

The Young modulus E is given by the slope of the stress versus strain relation, $d\sigma/d\epsilon$. For ideal elastic materials, which obey the classical Hooke's Law

$$\sigma = E\epsilon,\tag{4.4}$$

this derivative is a constant. In reality, $E(\epsilon) = d\sigma/d\epsilon$ is a function of the strain. As required by the ASTM standard for testing the tensile properties of nylon film [117], in this study the value used for E was always the maximum value attained by the function $E(\epsilon)$. It is important to note that both E and σ_t are smaller for films in the rubbery state than for films in the glassy state. As Figure 4.1 shows, nylon films also act qualitatively different on either side of T_g . Glassy films behave elastically, with a linear stress-strain relation up to a point of sudden, rapid failure. Above T_g , nylon films are not elastic all the way to the point of failure. Before failing, they will behave plastically, stretching for some time with no change required in the applied force. This is an irreversible deformation. It should also be noted that for films in the rubbery state, the transition between elastic and plastic behaviors is gradual. This causes a phenomenon called creep, a slow but irreversible deformation of the film under constant stress, even at a stress well below the tensile strength. Creep is an important phenomenon to consider in load-bearing nylon vessels that are required to survive for several (~10) years.

4.2.3 Griffith model of brittle fracture

Some quantitative notion of brittleness can be obtained from measurements of a thin film's fracture toughness, surface energy, and critical thickness. Fracture toughness K_{Ic} (also referred to in the literature as the critical stress intensity factor) is a measure of a film's ability to resist cracking. This quantity is derived from A. A. Griffith's model of brittle fracture [119]. The Griffith model is based upon the presumption that a cracked elastic film, stretched taut, will have its free energy U decreased if the crack spreads. Now, U is a function of both the stress σ on the film, and the size and geometry of the crack. Assuming a linear crack of length c, therefore, the fracture stress σ_f above which the crack spreads is given by setting $\partial U/\partial c = 0$ and solving for σ . (The quantity σ_f is, naturally, always less than the tensile strength σ_t of an intact film.)

In the simplest geometry, that of an infinitely wide sheet being stretched perpendicular to the crack by a uniform stress σ , the free energy per unit film thickness is [119]

$$U = -\frac{\pi\sigma^2 c^2}{4E} + 2c\gamma \tag{4.5}$$

where γ is the surface energy per unit area. The term on the left is a mechanical energy term, and the term on the right comes from the surface energy of the film. From this



Figure 4.2: Convenient geometry for measuring the fracture toughness of nylon films. A strip of thin film, width w and length L, has a centered crack of length c perpendicular to the length. The strips used in this study had effective dimensions L = 23 cm, w = 2.5 cm, and c ranging between 0.4 and 1.2 cm.

equation, the stress at failure is

$$\sigma_f = 2\sqrt{\gamma E/\pi c}.\tag{4.6}$$

The fracture toughness for in-plane stress is defined by

$$K_{Ic}^2 \equiv 2\gamma E,\tag{4.7}$$

so Equation (4.6) may be rewritten as

$$K_{Ic}^2 = \frac{\pi}{2} \,\sigma_f^2 c. \tag{4.8}$$

It is not practical to perform mechanical tests in this configuration (very wide sheets of nylon film would be required), so a generalization of the formula is needed. Fortunately, this is relatively simple. For a strip of finite width w, with a centered crack of length c (as shown in Figure 4.2), the equation becomes [119]

$$K_{Ic}^2 = Y^2 \sigma_f^2 c \tag{4.9}$$

where Y is a geometrical factor of order one; specifically,

$$Y^2 = \frac{w}{c} \tan\left(\frac{\pi c}{2w}\right). \tag{4.10}$$

In the limit $c/w \to 0$, Y^2 naturally reduces to $\pi/2$. It should be noted that in the literature, the strip width and crack length are frequently designated 2w and 2c, respectively, giving an extra factor of two in Equation (4.9).

This equation is only applicable to (in principle, atomically) thin cracks. However, values of Y have been tabulated for many other geometries as well. For instance, for a circular hole in an infinitely wide sheet of film, $Y^2 = 2/\pi$ [119]. This is smaller than the value of $Y^2 = \pi/2$ in Equation (4.8). From Equation (4.9), we see that for a material with a given fracture toughness K_{Ic} , Y is inversely proportional to σ_f : a smaller Y implies a larger σ_f . This corresponds to the real-life experience that a circular hole in a stressed film is less likely to spread than is a crack perpendicular to the direction of stress.

4.2.4 Surface energy and critical thickness

Not only should cracks be prevented from spreading, they must not form in the first place. The most important mechanism by which they may form is through creasing in the film. In everyday experience, thin sheets of plastic (such as plastic grocery bags) may be creased back upon themselves with no damage, while thicker chunks of plastic (such as circuit boards) will snap. In addition, dimensional analysis of two relevant quantities, the tensile strength σ_t (force per unit area) and the surface energy γ (energy per unit area) shows that they differ by a factor of length. These facts suggest the introduction of a "critical thickness," an approximate thickness above which a sheet of plastic folded back upon itself



Figure 4.3: Coordinate system used for a folded sheet of film. A dashed line indicates the surface that separates stretched and compressed regions.

will crack. The following argument is due to C. Galbiati, a member of the Borexino group at Princeton [116].

Consider a piece of plastic film of thickness t folded back upon itself, as shown in Figure 4.3. Suppose for simplicity that the middle surface of the film (dashed line in the figure) is not stretched and that the radius of curvature of the fold on the outer surface is t. Then the strain varies from 100% elongation on the outer surface to 100% compression on the inner surface. One may therefore make the approximation

$$\epsilon(r) = \frac{2r}{t} - 1. \tag{4.11}$$

Now, the energy U required to perform the work of folding is given by the volume integral $\frac{1}{2} \int \sigma \epsilon \, dV$. Within the distorted portion of the film, the strain is greater than that at the yield stress almost everywhere. Assuming an ideally plastic material, therefore, the stress may be replaced with σ_t in the stretched region and with $-\sigma_t$ in the compressed region.

Using the coordinate system shown in Figure 4.3, the integral becomes

$$U_{folding} = \frac{\sigma_t}{2} \int dz \int_0^{\pi} d\theta \left[\int_{t/2}^t r \, dr \left(\frac{2r}{t} - 1 \right) - \int_0^{t/2} r \, dr \left(\frac{2r}{t} - 1 \right) \right]$$

= $\frac{\pi}{8} \sigma_t t^2 \Delta z.$ (4.12)

The crux of the argument is to ask when the mechanical energy due to this strain is equal to the energy that would be required to crack the film. This latter value is simply the energy of the two new surfaces that would be formed:

$$U_{cracking} = 2\gamma t \,\Delta z. \tag{4.13}$$

Since the folding energy is proportional to t^2 , and the cracking energy to t, it is not energetically favorable for very thin films to crack when creased. Setting the two energies equal yields the critical thickness,

$$t_{crit} = (8/\pi)(2\gamma/\sigma_t).$$
 (4.14)

Films of nylon that are significantly thicker will crack when folded back upon themselves.

4.3 Testing material properties of nylon film

In outline, these tests were intended to measure the material properties of several different types of nylon at varying levels of relative humidity. It was also important to make sure that there were no significant differences between the properties of nylon immersed in air and pseudocumene at the same humidity.

Three types of nylon film were tested in this study: two candidate films (Capron and Sniamid) that were under consideration for use in the Borexino vessels, and a sample of film (C38F) used in construction of the Counting Test Facility (CTF) 1 and 3 vessels. (The film used in the CTF 2 vessel was chemically the same C38F nylon, but it was extruded at



Figure 4.4: The Controlled Relative Humidity Facility. From left to right, glove boxes C, A, and B.

a different plant; no samples of it remained available for testing.) These films are described in more detail in Section 4.1.4. The two films used in Borexino have an average thickness of $125 \pm 8 \,\mu\text{m}$, while the C38F film has an average thickness of $500 \pm 30 \,\mu\text{m}$.

4.3.1 The Controlled Relative Humidity Facility and nylon samples

The Controlled Relative Humidity Facility (CRHF), in which most of the experimental work took place, was a setup of three connected airtight glove boxes (Figure 4.4) filled with nitrogen and maintained at a temperature of $22 \pm 1^{\circ}$ C. The boxes were kept at a slight positive pressure in order to prevent outside air from diffusing into the humidity-controlled environment. This setup was constructed through the efforts of the Borexino group at Princeton. Two polycarbonate glove boxes (labeled A and B) were manufactured by Terra Universal, and the third (labeled C) was a larger custom-made box. Boxes A and B were used to hold nylon film samples while they reached equilibrium at a particular relative humidity. Box C contained a Tinius-Olsen machine (to be described in Section 4.3.3), used for stress testing of the nylon samples. A small fan placed in the connection between boxes A and B ensured good air circulation and mixing within the CRHF.

A Series 100 hygrometer from Nyad was used to measure relative humidity within the CRHF. This is a capacitive sensor with a resolution of 1 ppm of water vapor by volume in atmosphere. The sensor's accuracy is 50 ppm. (For comparison, a relative humidity of 20% at 22°C corresponds to 2480 ppm of water vapor.) The Series 100 hygrometer features a programmable analog output in the range 0-10 V, which was fed into a process controller. This controller automatically opened a dry vent line of nitrogen gas from a compressed nitrogen bottle when the humidity went above a maximum value, or a wet vent line of nitrogen bubbled through a glass jar of water when the humidity dropped below a minimum value. The minimum and maximum values were usually set to be 15 ppm below and above the desired relative humidity level. Thus, the moisture content of the CRHF could be stabilized to within about 30 ppm of the set point (0.008% of the maximum possible humidity at 22° C).

At the beginning of the experiment, the CRHF was set to reach zero humidity as closely as possible. A large number of nylon samples (to be described below) were placed in the glove boxes, and left to reach equilibrium with the dry atmosphere for 15 days. Measurements of material properties were performed as described in Sections 4.3.2 and 4.3.3, taking roughly two working days. Then the humidity control was ramped up to the next level, and nylon samples were given at least another 6 days to reach equilibrium again before more measurements were performed. This procedure was repeated at the relative humidity levels of 0%, 10%, 20%, 40% and 60%.

At each relative humidity level, tests were made on four different types of nylon film sample: Capron, $125 \,\mu\text{m}$ thick; Sniamid, $125 \,\mu\text{m}$ thick; C38F, $500 \,\mu\text{m}$ thick; and $125 \,\mu\text{m}$ thick Capron samples that were immersed in pseudocumene (these will be referred to as "Capron in PC"). The Capron and Sniamid types of nylon were the two contenders for use in the Borexino vessels. The thicker C38F came from the same batch of film used in the first and third versions of the Counting Test Facility; it was tested to ensure that the test conclusions were reasonably independent of the thickness of the nylon samples. For each of the four types of sample, 80 small $(3 \text{ cm} \times 3 \text{ cm})$ strips of film and 320 large $(33 \text{ cm} \times 2.5 \text{ cm})$ strips of film were prepared. The small strips were to be used for the moisture content tests, and the large strips for tests of material properties. For the first three types of sample, the small strips of film were placed on perforated racks in the glove boxes, while the large strips were hung from the ceiling of glove box A. This was done to ensure that samples felt sufficient airflow to reach equilibrium with the atmospheric humidity on both sides. Samples for the "Capron in PC" condition, both large and small, were placed in a tank of 16ℓ of pseudocumene that was open to the atmosphere of glove box B and in equilibrium with it. To prevent mix-ups, every sample was uniquely labeled with a permanent marker.

4.3.2 Tests of moisture content

For two reasons, it was desirable to test the water content of the nylon samples. First, there was some question as to whether water content of the films was a linear function of the relative humidity of the environment; second, the water content would give some indication as to whether the nylon samples had actually reached equilibrium. For consistency, the water content was always expressed as a fraction of the dry nylon mass.

The tests were performed with a Computrac 3000 moisture analyzer (Figure 4.5) produced by Arizona Instruments [120]. Samples for the Computrac 3000 must be placed inside a vial with a special lid containing a rubber septum. To perform a measurement, a sample is placed in the vial and the lid is sealed; the vial is then inserted into the Computrac 3000. The instrument punctures the septum with a hollow needle and heats the vial to a preset temperature. Gases escaping through the needle are analyzed for water vapor content until the rate of escaping water vapor falls below a programmable cutoff value. The Computrac 3000 then reports the total mass of water in the sample by time integration over the water vapor content.



Figure 4.5: The Computate 3000 moisture analyzer made by Arizona Instruments.

Humidity	H_2O content [% by dry mass]							
[%]	Ca	pron	Capro	on in PC	Sni	amid	С	38F
0	0.18	(0.01)	0.16	(0.01)	0.20	(0.02)	0.58	(0.05)
10	1.39	(0.05)	1.47	(0.01)	1.45	(0.03)	1.58	(0.07)
20	2.21	(0.10)	2.16	(0.09)	2.23	(0.07)	2.46	(0.07)
40	4.23	(0.13)	4.05	(0.16)	3.90	(0.15)	4.36	(0.14)
60	6.00	(0.29)	6.26	(0.14)	6.17	(0.20)	6.31	(0.42)

Table 4.1: Moisture content of the four types of nylon sample as a function of relative humidity in the surrounding atmosphere. These values were determined with the Computrac 3000. The percentages are relative to the dry mass. Statistical errors are shown in parentheses.

Humidity	H_2O content [% by dry mass]							
[%]	Capron		Capron in PC		Sniamid		C38F	
0	0.00	(0.00)	0.00	(0.00)	0.00	(0.00)	0.00	(0.00)
10	0.55	(0.08)	0.49	(0.15)	0.71	(0.05)	0.84	(0.14)
20	0.93	(0.13)	1.48	(0.26)	1.08	(0.04)	1.37	(0.32)
40	2.24	(0.05)	2.41	(0.08)	2.15	(0.04)	2.76	(0.13)
60	3.72	(0.13)	3.83	(0.18)	3.55	(0.02)	4.18	(0.36)

Table 4.2: Moisture content of the four types of nylon sample as a function of relative humidity in the surrounding atmosphere. These values were determined by measuring the mass gain in individually identified samples of nylon, so the values at a relative humidity of 0% are all zero by definition. The percentages are relative to the dry mass. Statistical errors are shown in parentheses.



Figure 4.6: Moisture content (% by dry mass) of the four types of nylon samples as a function of humidity, determined with the Computrac 3000.



Figure 4.7: Moisture content (% by dry mass) of the four types of nylon samples as a function of humidity, determined by measuring the increasing masses of individual samples.

The nylon samples could not be exposed to air outside the controlled environment of the CRHF, as this might skew the results. Therefore, to measure the water content of a nylon sample correctly, the procedure had to be complicated a bit. First, several clean vials and lids, as well as new septa, were placed inside the CRHF via an airlock. The vials were left to sit for at least an hour. Then at least three samples of each type of nylon were placed into vials. Each sample was a small piece (masses ranged between 20-200 mg) of the 3 cm \times 3 cm nylon strips. The vials were sealed, as well as additional "blank" vials containing only air from inside the glove boxes. Next the vials were removed from the CRHF, and the mass of water contained in each sample and blank was measured with the Computrac 3000. (The temperature of the Computrac 3000 oven was set to be 220°C, which is the temperature suggested by Arizona Instruments for measuring moisture content in most nylon brands. The cutoff rate was set to be $0.30 \,\mu g/s$ of water vapor.) From the mass of water measured for each nylon sample, the average water mass contained by a blank sample was subtracted. Each nylon piece was then assumed to be dry and was weighed on a scale with 10 μ g resolution. Dividing the mass of water in each nylon sample by the dry mass yielded the fractional water content. Results are given in Table 4.1 and Figure 4.6.

As an independent check on these measurements, water contents of nylon samples at the various relative humidities were measured based solely on mass differences. Five samples of each type of nylon were numbered and kept in the CRHF. At each humidity level, the new mass of each sample was recorded. The fractional water content in this case was calculated as the difference between each sample's current mass and its original dry mass (at 0% relative humidity), divided by the original dry mass. These results are summarized in Table 4.2 and Figure 4.7.

For the analysis of moisture content and the other properties discussed in later sections, the results of every individual test were recorded in a spreadsheet. A separate worksheet was devoted to material properties at each relative humidity, as well as to the moisture content results. From the results calculated on each worksheet, a summary worksheet was compiled containing the average values for each type of sample and relative humidity. Finally, the Origin data analysis package [121] was used to produce graphs from this summary worksheet. Error bars given in tables and shown on the graphs are one-sigma statistical errors derived from the spread of the data; no attempt was made to determine systematic errors. However, given the variability inherent in measuring bulk material properties, any systematic errors are likely to be relatively unimportant.

The fractional moisture contents of the types of nylon, as a percentage of dry weight, are shown as functions of the relative humidity in Figures 4.6 and 4.7, and tabulated in Tables 4.1 and 4.2. The functions are nearly linear, and very similar for the four types of nylon. As measured by the Computrac 3000, the C38F has significantly higher moisture content at zero relative humidity, which may be explained by a failure of the thicker film to reach equilibrium completely.

There are significant discrepancies between the values determined with the Computrac 3000 (Table 4.1), and the values obtained by measuring the increasing nylon masses (Table 4.2). The latter set of data is in good agreement with similar data measured for 18 μ m-thick C38F film by Wójcik and Zuzel in 2004, using the same method [45]. It seems unlikely that the Computrac 3000 was miscalibrated; calibration measurements were performed with it frequently and the calibration error would have to be on the order of 50%. Other explanations seem to require that "dry" nylon has a level of ~ 2% moisture content that is chemically bound to the polymer and cannot be removed by heating. As the environment becomes more humid, this moisture content would somehow be released as free water molecules, dissolved in the nylon.

4.3.3 Tests of material strength

The Tinius-Olsen machine is a device with two jaws that can be made to separate at a constant rate (Figure 4.8a). Its jaws are attached to a load cell measuring the force between them. The Tinius-Olsen machine produces analog outputs that can be plugged



Figure 4.8: (a) the Tinius-Olsen machine; (b) a graph of a stress-strain relationship produced with it.

into the axes of a plotting device, yielding a graph of the force required to separate the jaws as a function of their distance apart (Figure 4.8b). With the appropriate transformations, this becomes a graph of the stress versus strain relationship. Such tests conform to the ASTM standard D882 for measuring the strength of thin plastic sheeting [117]. They were informally referred to as "pull tests."

For tests of strength, a $33 \text{ cm} \times 2.5 \text{ cm}$ sample of nylon film was transferred from glove box A into box C, where its ends were inserted into the jaws of the Tinius-Olsen machine. The jaws are set to pull apart until the force between them reaches a maximum. This is the point of failure, at which a nylon film will break if in the glassy state, or stretch plastically if in the plastic state. The maximum force, in Newtons, is shown on a digital display on the machine. From this value, the dimensions of the nylon film, and the graph produced on the plotting device, the tensile strength and Young modulus of the film may be calculated. At each humidity level, for each of the four types of nylon, three to five nylon samples were tested in this way.

A number of parameters may be set on the Tinius-Olsen machine, including the rate at which the jaws separate, and the vertical and horizontal scales of the resulting graphs. For these tests of tensile strength, the jaw separation rate was set to 250 mm/minute. The

Humidity	Tensile strength [MPa]							
[%]	Capron		Capron in PC		Sniamid		C38F	
0	68.9	(8.8)	70.2	(1.9)	75.0	(0.9)	74.3	(2.7)
10	53.4	(4.3)	58.6	(3.7)	66.0	(1.2)	75.4	(2.7)
20	41.4	(1.7)	38.3	(0.3)	53.4	(1.8)	66.0	(1.4)
40	22.9	(0.9)	23.7	(0.5)	25.9	(0.5)	25.5	(2.4)
60	18.9	(0.8)	18.5	(0.2)	20.5	(0.4)	17.6	(0.5)

Table 4.3: Tensile strengths, in MPa, of the four types of nylon sample as a function of relative humidity in the surrounding atmosphere. Statistical errors are shown in parentheses.

Humidity	Young modulus [GPa]								
[%]	Ca	pron	Capro	Capron in PC		Sniamid		C38F	
0	1.66	(0.11)	1.61	(0.08)	1.64	(0.07)	1.27	(0.05)	
10	1.29	(0.11)	1.38	(0.14)	1.59	(0.04)	1.32	(0.09)	
20	1.02	(0.10)	0.90	(0.01)	1.22	(0.14)	1.21	(0.03)	
40	0.51	(0.05)	0.57	(0.02)	0.61	(0.04)	0.54	(0.10)	
60	0.42	(0.04)	0.42	(0.04)	0.42	(0.05)	0.34	(0.01)	

Table 4.4: Young moduli, in GPa, of the four types of nylon sample as a function of relative humidity in the surrounding atmosphere. Statistical errors are shown in parentheses.

vertical scale of the graph was set to 10 or 5 for the $\sim 100 \,\mu\text{m}$ thick nylon strips, and 50 or 20 for the 500 μm thick strips. The horizontal scale was set to 50 (corresponding to 1 mm of jaw separation being equivalent to 4 mm of horizontal scale on the graph).

Tensile strength for each sample was easily obtained by dividing the maximum force obtained during a pull test by the cross-sectional area of a nylon strip. (Note that the thickness varied slightly from strip to strip; this quantity was measured for each sample before the pull test with a micrometer.)

The Young modulus was determined as follows: a tangent line to the plot of force versus distance was constructed at the steepest part of the graph produced by the Tinius-Olsen machine. It was determined how far the nylon strip would have stretched if it obeyed Hooke's Law perfectly; that is, if the slope had been this steep all the way to the point of maximum force. The strain of this hypothetical extension was calculated as a fractional



Figure 4.9: Tensile strength, in MPa, of the four types of nylon samples as a function of humidity.



Figure 4.10: Young moduli, in MPa, of the four types of nylon samples as a function of humidity.

change in length, $\epsilon = \Delta L/L$. (It is important to note that L, the effective length of each strip, was 23 cm rather than the actual length of 33 cm, since 5 cm on each end were held firmly by the jaws of the Tinius-Olsen machine.) The Young modulus was calculated as the tensile strength of this strip divided by the hypothetical strain ϵ .

Tables 4.3 and 4.4 and the corresponding graphs (Figures 4.9 and 4.10) show that both the tensile strength and the Young modulus decrease rapidly as functions of the relative humidity. At low values of relative humidity, the different types of nylon have significantly different values for these quantities. At 60% relative humidity, however, all types of nylon behave similarly, with tensile strength of about 19 MPa and a Young modulus around 400 MPa. One can also estimate the humidity at which the glass transition temperature T_g crosses 22°C. This humidity appears to be lowest for the Capron, around 10%, and highest for the C38F, at perhaps 30%.

It was observed during these strength tests that the qualitative behaviors of the nylon films are also quite different at low and high values of relative humidity. Below the glass transition point, nylon strips obey Hooke's Law rigorously until the point of failure, at which they abruptly shatter. Above the transition point, strips act plastically and may stretch for quite some time without the measured force changing.

Finally, it should be noted that the material properties of Capron in air are very similar to those of Capron soaked in pseudocumene. In other words, the properties of Capron film are determined solely by the relative humidity of its surroundings; pseudocumene does not weaken the plastic in any way. This is of course very important to the Borexino project, in which the nylon vessels will experience both air and pseudocumene environments.

4.3.4 Tests of creep

Although the phenomenon of creep was not studied in parallel with the other measurements of mechanical strength described above, several measurements of creep were performed

Material	Stre	ess	Elongation		
	[MPa]	[psi]	(6 hours)	(1 day)	
C38F	3.4	500	1.7%	3.0%	
in	7	1000	5.8%	9.8%	
nir	10	1500	11%	23%	
all	12	1800	20%	44%	
C28F	3.4	500	2.4%	2.2%	
U30F	7	1000	3.4%	6.2%	
III Trator	10	1500	38%	68%	
water	12	1800	120%	125%	

Table 4.5: Results of creep from short-term measurements of C38F film under various constant loads in air and water environments. Elongation values are precise to about 1%. Taken from reference [122].

Material	Stress		Elongation			
	[MPa]	[psi]	(10 days)	(20 days)	(50 days)	
Snjamid	1.7	250	0.9%	1.1%	0.8%	
Loistriz in	3.4	500	1.4%	1.8%	1.4%	
Leistriz III	5	750	3.2%	3.6%	3.6%	
pseudocumene	7	1000	6.2%	6.6%	7.0%	
C00	1.7	250	$<\!0.5\%$	$<\!0.5\%$	$<\!0.5\%$	
C90- Loistriz in	3.4	500	1.7%	1.6%	1.8%	
	5	750	3.6%	3.9%	4.1%	
pseudocumene	7	1000	5.7%	7.0%	7.3%	

Table 4.6: Results of creep from long-term measurements of Sniamid-Leistriz and C90-Leistriz film under various constant loads, immersed in pseudocumene in equilibrium with outside air in the summer. Elongation values are precise to about 0.5%. Taken from reference [12].
earlier on C38F, Sniamid-Leistriz, and C90-Leistriz film (*cf.* Section 4.1.4). Some shortterm tests reported in [122] measured the amount of creep in 250 μ m-thick C38F film under various loads in both air and water. The results of these measurements are tabulated in Table 4.5. Several important conclusions can be drawn. First, the presence of water in the environment seems to make a strong difference at loads greater than ~7 MPa. Second, within a time scale of one day, the amount of elongation in nylon film under constant stress continues to increase rapidly. Finally, the amount of creep as a function of applied load increases much faster than linearly, so it is important to keep the stress level as low as possible.

Longer-term tests were also done. First, in the NSF proposal appendix, two samples of the same C38F film were hung in water at temperatures of 23°C and 8°C, respectively, with constant loads of 7 MPa for four months. Both samples were stretched by 10-15% in the first few days, but then remained at a constant length [122]. Tests in summer 1998 on Sniamid-Leistriz and C90-Leistriz nylon film samples immersed in pseudocumene exposed to the humidity in normal air also show that the elongation of load-bearing film reaches a static value after an initial ~10 day period of stretching. After 50 days, the maximum fractional elongation achieved for the Sniamid-Leistriz film was < 1% at a stress level of 1.7 MPa, and ~7% at a stress of 7 MPa [12]. Very similar results were obtained for the C90-Leistriz film (see Table 4.6). As with the short-term tests, the long-term creep values are a strongly increasing function of the applied constant load.

For comparison, the operational stress in Borexino will be about 1.8 MPa [12]. Temperature gradients between fluids in the three volumes of pseudocumene may however cause stresses up to an order of magnitude higher at the design limit of $\Delta T = 5^{\circ}$ C. It is therefore essential that these temperature differences be noticed and minimized somehow within a few hours. The values tabulated in Table 4.5 point out the importance of keeping the pseudocumene relatively dry during Borexino operation, in case of such a circumstance.

4.3.5 Tests of brittleness

Spreading of cracks in film

Tests for fracture toughness and related quantities were slightly more involved than tests of tensile strength. Each $33 \text{ cm} \times 2.5 \text{ cm}$ nylon strip used in these tests was initially prepared with a slit in the center, perpendicular to the length of the strip, as shown in Figure 4.2. Slits were of length c = 4, 6, 8, or 10 mm, except for the thicker C38F samples, which were given slits of length c = 6, 8, 10, or 12 mm. At each humidity level, for each of the four types of nylon, at least one sample with each size of slit was pulled to failure in the Tinius-Olsen machine. The maximum force applied during each test was recorded.

These measurements allowed several quantities related to brittleness to be determined. First, fracture toughness can be calculated from the formula

$$K_{Ic} = Y\sigma_f \sqrt{c},\tag{4.15}$$

obtained from the theory of the Griffith model described in Section 4.2.3. Surface energy may be derived from the fracture toughness with the formula

$$2\gamma = \frac{K_{Ic}^2}{E},\tag{4.16}$$

where E is the Young modulus. (In these results, "surface energy" will be reported as the quantity 2γ , the amount of energy required per unit area to form the *two* new surfaces of a crack.) Finally, the critical thickness t_{crit} is calculated as

$$t_{crit} \approx \frac{8}{\pi} \frac{2\gamma}{\sigma_t}.$$
(4.17)

For each strip, the value of K_{Ic} was calculated using the value of σ_f measured in a pull test and the value of Y calculated from the size of the center slit. The quantity 2γ was determined for each individual strip using that strip's value of K_{Ic} and the average Young modulus E, obtained as described in the previous section, for the appropriate sample type and relative humidity. Then K_{Ic} and 2γ were averaged over all strips of the same type of nylon at the same humidity (regardless of slit size). The critical thickness t_{crit} was produced from the average value of σ_t for this humidity, determined as described in the previous section, and the averaged values for K_{Ic} and 2γ for this humidity.

Some difficulties were encountered with this method of analysis at higher values of relative humidity. The Griffith theory of fracture toughness is well defined only for elastic materials. Above the glass transition point, nylon behaves plastically instead. Calculated results for nylon films at higher values of relative humidity were dependent upon the slit size of the sample. However, for the samples with larger slits, the stress at failure σ_f was significantly below the tensile strength σ_t and still within the range of roughly elastic behavior. It was therefore assumed that calculating fracture toughness using these samples would still give approximately correct results. Thus, at 40% and 60% humidity, data from samples with the smallest slit size were excluded from the preceding calculations.

The results show that the fracture toughness decreases for all types of nylon as relative humidity increases (Table 4.7, Figure 4.11). Surface energy also appears to decrease (Table 4.8, Figure 4.12), although the wider errors make the extrapolation of a trend more difficult. Finally, it is clear that the critical thickness of nylon is low (about 500-700 μ m) at lower ambient humidity, and substantially larger (1.0-1.4 mm) above the glass transition point (Table 4.9, Figure 4.13). In fact, the calculated critical thickness at low humidity should be taken with a large grain of salt. Recall from Section 4.2.4 that the critical thickness model is based upon the assumption of a film that behaves plastically. Films in the glassy state behave elastically, which would lower the true critical thickness at low humidity well below the calculated values. This helps to explain the cracking failure of the CTF 2 vessel in 1999. This vessel was stored in a very dry nitrogen atmosphere, and its thickness of 500 μ m would have been dangerously close to the true critical thickness, if not above it, at that humidity.

Humidity	Fracture toughness $[MPamm^{1/2}]$							
[%]	Cap	oron	Capror	n in PC	Snia	mid	Ca	B8F
0	151.2	(22.1)	131.5	(4.2)	135.9	(6.8)	163.4	(13.0)
10	125.5	(7.8)	120.0	(7.7)	126.3	(8.0)	159.8	(6.7)
20	108.5	(5.0)	113.7	(7.0)	117.0	(7.2)	149.9	(5.0)
40	70.4	(3.2)	77.3	(3.9)	81.1	(5.1)	86.0	(7.9)
60	56.9	(4.7)	55.7	(2.0)	57.5	(3.7)	57.5	(5.3)

Table 4.7: Fracture toughness, in $MPa \,mm^{1/2}$, of the four types of nylon sample as a function of relative humidity in the surrounding atmosphere. Statistical errors are shown in parentheses.

Humidity	Surface energy $[kJ m^{-2}]$							
[%]	Cap	oron	Capror	n in PC	Snia	amid	C3	88F
0	14.0	(4.3)	10.7	(0.7)	11.3	(1.1)	21.2	(3.2)
10	12.2	(1.5)	10.4	(1.3)	10.1	(1.2)	19.4	(1.6)
20	11.6	(1.1)	14.4	(1.8)	11.3	(1.3)	18.6	(1.2)
40	9.7	(0.9)	10.5	(1.1)	10.9	(1.4)	13.7	(2.5)
60	7.8	(1.3)	7.4	(0.5)	8.0	(1.1)	9.9	(1.8)

Table 4.8: Surface energy, in $kJ m^{-2}$, of the four types of nylon sample as a function of relative humidity in the surrounding atmosphere. Statistical errors are shown in parentheses.

Humidity	Critical thickness [mm]							
[%]	Ca	pron	Capro	n in PC	Sni	amid	\mathbf{C}	38F
0	0.26	(0.09)	0.19	(0.01)	0.19	(0.02)	0.36	(0.06)
10	0.29	(0.04)	0.23	(0.03)	0.19	(0.02)	0.33	(0.03)
20	0.36	(0.04)	0.48	(0.06)	0.27	(0.03)	0.36	(0.02)
40	0.54	(0.05)	0.56	(0.06)	0.54	(0.07)	0.69	(0.14)
60	0.52	(0.09)	0.51	(0.04)	0.50	(0.07)	0.72	(0.13)

Table 4.9: Critical thickness, in mm, of the four types of nylon sample as a function of relative humidity in the surrounding atmosphere. Statistical errors are shown in parentheses.



Figure 4.11: Fracture toughness, in $MPa \,mm^{1/2}$, of the four types of nylon samples as a function of humidity.



Figure 4.12: Surface energy, in $kJ m^{-2}$, of the four types of nylon samples as a function of humidity.



Figure 4.13: Critical thickness, in mm, of the four types of nylon samples as a function of humidity.

Failure due to repeated mishandling

Two additional sets of tests were performed that did not involve the Tinius-Olsen machine. The first was very qualitative, consisting of taking thin pieces of film and bending them back and forth, by hand, along the same crease. It was observed that dry film would quickly snap in two, while film that had been kept in a humid environment could be bent many times. In the latter case, the worst that would happen to the film was the development of opaque white areas of crazing along the crease. Though unaesthetic, crazing does not affect the film strength greatly.

The so-called packet tests were more quantitative tests of the ability of nylon film to survive repeated mishandling. Several nylon packets were produced, consisting of two circles of film, 28 cm in diameter, glued together at the edges. To the center of one circle could be attached a plastic tube. These packets, like the other nylon samples, were stored in the CRHF in order to reach equilibrium. At each humidity level, several packets of each type were tested in glove box B. (Packets were produced only with Capron and Sniamid films.)



Figure 4.14: Results of the packet tests: number of vacuum cycles survived by each Capron and Sniamid packet (referred to as "pillows" in the figure) at varying humidities.

For a packet test, a nylon packet was attached to a gas manifold comprising several valves and vacuum pumps. The test consisted of repeated inflation and deflation of the packet in the following cycle:

- 1. Inflate the packet at a positive pressure of 9 torr for 90 s;
- 2. Deflate the packet for 90s using a vacuum on the order of 50 torr absolute pressure;
- 3. Close the values for $50 \, \text{s}$;
- 4. Evaluate the pressure in the packet.

The test is defined to be over when the absolute pressure in step 4 of the cycle has exceeded 400 torr. At this point the number of cycles for which the packet survived was recorded.

Figure 4.14 shows the results of the packet tests for the Capron and Sniamid packet. Unfortunately, this plot suffers from low statistics, so data points for the individual packets have been plotted as well as the averages. Although the data are sketchy, it appears that the Sniamid packets fail sooner, *i. e.*, in significantly fewer cycles, than the Capron packets. This can be correlated with the observation that the critical thickness of Capron is marginally greater than that of Sniamid at all humidity levels.

4.4 Radioactive contamination of nylon surfaces

Because the nylon vessels are in intimate contact with the scintillator in Borexino, it is crucial that they be clean and low in radioactive contaminants. One can imagine four ways in which the vessels could cause (or fail to prevent) an unacceptably high event rate in the scintillator:

- The Inner Vessel may contain a level of radioactive contaminants sufficient to overwhelm the neutrino signal, even in the Fiducial Volume ("intrinsic radioactivity").
- Radon produced by external sources may diffuse through the nylon vessels into the scintillator ("diffusion").
- Radium trapped in the nylon vessel material may produce radon that migrates into the scintillator ("emanation").
- Radioactive particles (dust or adsorbed ions) adhering to the surface of the nylon film may come off into the scintillator ("washoff").

4.4.1 Intrinsic activity in nylon film

Of these four possibilities, the first is of the least concern. Recall that α and β particles travel no more than a few cm in pseudocumene; even γ rays will generally travel 10-20 cm. The spatial resolution of the detector for 250 keV events (at the lower end of the neutrino window) will be on the order of 20 cm, as can be estimated from the techniques of Chapter 5. If the level of radioisotopes embedded in the Inner Vessel film is high enough to

Sample	Bulk ²²⁶ Ra	Surface ²²⁶ Ra	Total ²²⁶ Ra
1	[mBq/kg]	$[\mu Bq/m^2]$	[mBq/kg]
Capron B73ZP (1st batch)	0.21 ± 0.03	1.1 ± 0.3	0.22 ± 0.04
Capron B73ZP (2nd batch)	0.46 ± 0.07	6.4 ± 1.6	0.55 ± 0.08
Sniamid blend	< 0.021	< 0.8	0.016 ± 0.004

Table 4.10: Radium activity in Borexino vessels film deduced via a mathematical model of radon diffusion from measurements made at Heidelberg [123]. It is believed that the second batch of Capron film was contaminated during extrusion, explaining the higher activities.

Sample	Type	226 Ra	$^{238}\mathrm{U}$	232 Th	Κ
		[ppt U equiv.]	[ppt]	[ppt]	[ppb]
Sniamid	pellets	-	1.1 ± 0.0	1.6 ± 0.1	25 ± 9
Sniamid	film	1.3 ± 0.3	2.8 ± 0.1	3.8 ± 0.2	-
Capron	pellets	-	0.46 ± 0.04	1.1 ± 0.1	25
Capron	film	18 ± 3	-	-	-

Table 4.11: Summary of radioactive contamination in the two Borexino nylon vessel films. ²²⁶Ra contamination is derived from the Heidelberg measurements, ²³⁸U and ²³²Th from ICP-MS at Tama, and K concentrations from NAA at the Technical University of Munich and the Missouri University Research Reactor. The values for K contamination from NAA are in disagreement between the two sources; the larger value has always been printed in this table. The directly measured values for U and Th contamination of Sniamid film were made on the pure ADS40T "Sniamid-Leistriz" film extruded at American Leistriz in 1998 (*cf.* Section 4.1.4), not on the 83%/17% ADS40T/Ultramid B4 blend extruded at mf-folien in 2001 and used in the Borexino Inner Vessel.

overwhelm the neutrino signal in the Fiducial Volume, 125 cm away, then radon emanation is likely to be a far worse problem. A Monte Carlo estimate by L. Cadonati, assuming film impurities (by mass) of 2 ppt ²³⁸U, 4 ppt ²³²Th, and 10 ppb potassium, yields an estimate of < 0.01 counts/day in the Fiducial Volume due to emitted γ rays [12]. Even with the revised contamination numbers given below, this value is still less than 0.05 counts/day in the Fiducial Volume.

Ultra-sensitive measurements done at Heidelberg based on a mathematical model of radon emanation [123] have yielded the measurements of 226 Ra activity shown in Table 4.10. (It should be noted that if secular equilibrium in the 238 U decay chain is assumed, $12.35 \,\mu$ Bq of ²²⁶Ra activity per kilogram of film is equivalent to 1 ppt ²³⁸U contamination.) A summary of these measurements is given in Table 4.11. Also shown in the table are direct measurements of the three troublesome elements by inductively coupled mass spectroscopy (ICP-MS) done by Tama Chemicals [124] in 1999, and neutron activation analysis (NAA) performed by R. von Hentig and T. Riedel in 1998-1999 [12, 63].

It is not known why the Capron film is a factor of ten worse than the Sniamid film in Ra contamination when the Capron pellets initially seemed better. But by the time that the radon emanation measurement on the Capron film had been made, the Borexino Outer Vessel had already been constructed from Capron, which at the time was the first-choice material due to its better pellet purity and its slight mechanical advantages discussed in Section 4.3. Since the Outer Vessel is 1.25 m from the scintillator and 2.5 m from the Fiducial Volume, this situation was judged to be tolerable. Only the Inner Vessel was constructed from Sniamid film.

4.4.2 Radon diffusion through nylon

We now examine the effectiveness of nylon films as a barrier to radon gas. First consider atoms of a stable noble gas diffusing into a medium. They obey Fick's Law, which states that the rate of diffusion is proportional to the local impurity concentration gradient:

$$\mathbf{J} + D\nabla\rho = 0. \tag{4.18}$$

Here,

- J [atoms / (area \times time)] is the molecular flux of the impurity,
- ρ [atoms / volume] is the local concentration of the impurity,
- D [area / time] is the diffusion coefficient.

When combined with the continuity equation

$$\nabla \cdot \mathbf{J} + \partial_t \rho = 0, \tag{4.19}$$

this yields the diffusion equation

$$\partial_t \rho = D \,\nabla^2 \rho. \tag{4.20}$$

Because the nylon film of the vessels is very thin compared to their radii, they may be treated mathematically as thin sheets, infinite in extent. We assume a coordinate system in which a sheet of film extends infinitely in the yz-plane, with thickness d extending from x = 0 to x = d. Let the outside of a vessel be in the negative x direction. Then Equations (4.18) and (4.20) may be restricted to one dimension:

$$J(x,t) = -D \rho_x(x,t)$$
 (4.21)

$$\rho_t(x,t) = D \rho_{xx}(x,t).$$
(4.22)

These equations govern the diffusion of stable substances through the vessels.

However, ²²²Rn has a finite mean life $\tau = 5.516$ days¹⁰: without diffusion, the radon concentration at any point would decay exponentially, following the equation

$$\rho_t = -\rho/\tau. \tag{4.23}$$

The equation that must actually be used to find the rate of radon diffusion through nylon takes both diffusion and decay into account:

$$\rho_t = D\rho_{xx} - \rho/\tau. \tag{4.24}$$

An important length scale for this equation is the characteristic diffusion length, $\ell \equiv \sqrt{D\tau}$ (also written as $1/\alpha$ in some references). This is the mean distance that a radon atom will

¹⁰As a side note, throughout this thesis the symbol τ (optionally with an identifying superscript) is used to represent the mean life of an isotope, while the symbol $\tau_{1/2}$ (also with an optional superscript) is used for the half-life. The two quantities are of course connected by the relation $\tau_{1/2} = \tau \log 2 \approx 0.69315 \tau$.

Material	Reference	Diffusion constant	Diffusion length
		$D \ [\mathrm{cm}^2/\mathrm{s}]$	$\ell [{ m cm}]$
Nylon (dry)	[125]	2.5×10^{-12}	$1.1 imes 10^{-3}$
Nylon (wet)	[125]	1.3×10^{-9}	$2.5 imes 10^{-2}$
Water	[126]	10^{-5}	2
Pseudocumene	[123]	2×10^{-5}	3
Air	[127]	0.11	230

Table 4.12: Diffusion constants for ²²²Rn in various substances. The diffusion length $\ell = \sqrt{D\tau}$ is the mean distance a radon atom will travel before it decays. The diffusion lengths for ²²⁰Rn ("thoron"), part of the ²³²Th decay chain, are not tabulated here, but are smaller by two orders of magnitude.

travel in a medium with diffusion constant D. Values of D and ℓ for radon in some relevant materials are given in Table 4.12; the first estimate of these values was in fact made with the Counting Test Facility [128]. Note that D for radon in nylon varies by almost three orders of magnitude (and ℓ by a factor of 25) depending upon the moisture content in the nylon. The specific dependence has been studied and found to be relatively flat for nylon in equilibrium with a relative humidity between 0-30% in the surrounding environment, but steeply increasing above that [45].

In Borexino, the concentration of radon in solution outside the Outer Vessel will be essentially constant, replenished by emanation from the photomultiplier tubes and Stainless Steel Sphere. Let this concentration be ρ_0 . This gives the boundary condition

$$\rho(0,t) = S\rho_0. \tag{4.25}$$

S is the ratio of the solubility of radon in nylon (number of atoms dissolved per unit volume) to the solubility of radon in the fluid outside the vessel.

We wish to find the stationary solution, in which $\rho_t = 0$ at any point within the Outer Vessel film, and in which the concentration of radon in solution within the Inner Buffer is also constant. This is the solution that will be reached asymptotically at times $t \gg \tau$. In order to do so, we must first determine the boundary condition at x = d. There are two extreme possibilities: on one hand, radon may mix thoroughly in the Inner Buffer; on the other hand, there may be no mixing and the internal radon concentration radial profile is governed by radon diffusion in pseudocumene. The former is the worst case, so it is the only one considered here.

If we assume that radon in the Inner Buffer mixes completely in a time scale much less than the radon mean life, then the radon concentration in the Inner Buffer $\rho^{\text{IB}}(t)$ is spatially homogeneous. The total number of radon atoms in the Inner Buffer, $N^{\text{IB}} = V^{\text{IB}}\rho^{\text{IB}}$, is given by the differential equation

$$\partial_t N^{\rm IB} = -N^{\rm IB}/\tau + 4\pi R^2 J(d,t).$$
 (4.26)

Substituting in Equation (4.21), the stationary version of this equation becomes

$$V^{\rm IB}\rho^{\rm IB} = -4\pi R^2 D\tau \rho_x(d) \tag{4.27}$$

(*R*, the radius of the Outer Vessel, is 5.5 m). We equate $\rho(d) = S \rho^{\text{IB}}$, obtaining the (ugly) boundary condition

$$\rho_x(d) = -\frac{V^{\rm IB}}{4\pi R^2 S D\tau} \,\rho(d). \tag{4.28}$$

The stationary solution to Equation (4.24) given the boundary conditions in Equations (4.25)and (4.28) is

$$\rho(x) = S\rho_0 \frac{\sinh\left[(x_0 - x)/\ell\right]}{\sinh\left(x_0/\ell\right)}$$
(4.29)

where $x_0 = d + \ell \tanh^{-1}(4\pi R^2 S \ell / V^{\text{IB}}).$

For our purposes, it can be shown that $4\pi R^2 S\ell/V^{\text{IB}} \ll 1$, so we may make the approximations $\sinh(x_0/\ell) \approx \sinh(d/\ell)$ and $\sinh[(x_0-d)/\ell] \approx 4\pi R^2 S\ell/V^{\text{IB}}$. (A plot of this approximation is shown in Figure 4.15a.) Therefore,

$$\frac{\rho^{\mathrm{IB}}}{\rho_0} = \frac{\rho(d)}{\rho(0)} \approx \frac{4\pi R^2 S\ell}{V^{\mathrm{IB}} \sinh\left(d/\ell\right)}.$$
(4.30)

Alternatively, as in [12] and [125], we may define an "effective permeability" for the film

$$P_{\rm eff} = DS \, \frac{d/\ell}{\sinh\left(d/\ell\right)} \tag{4.31}$$

Environment	Rel. solubility	Diffusion const.	Eff. permeability	$ ho^{ m IV}/ ho_0$
	S	$D [\mathrm{cm}^2/\mathrm{s}]$	$P_{\rm eff} \ [{\rm cm}^2/{\rm s}]$	-
Nylon in dry PC	1.0	2.5×10^{-12}	6.1×10^{-16}	3.8×10^{-20}
Nylon in humid PC	0.15	1.3×10^{-9}	1.8×10^{-10}	3.7×10^{-9}
Nylon in H_2O	0.7	1.3×10^{-9}	8.7×10^{-10}	7.9×10^{-8}

Table 4.13: Predicted ratios of radon concentration in the Inner Vessel to that in the Outer Buffer under the pessimistic assumption of complete mixing. Recall that the Outer and Inner Vessel radii are 5.5 m and 4.25 m respectively, and each vessel has a thickness of $d = 125 \,\mu\text{m}$. The values of S and D come from those tabulated for nylon 6 in reference [125], except for the "nylon in humid PC" environment, in which S is calculated from the value at 100% relative humidity in Table 1 of reference [45], scaled for an environment of pseudocumene rather than N₂.

and express this result as

$$\frac{\rho^{\rm IB}}{\rho_0} = \frac{4\pi R^2}{V^{\rm IB}} \frac{\tau}{d} P_{\rm eff}.$$
(4.32)

Combining this result with the analogous result at the Inner Vessel (radius r = 4.25 m), using $V^{\text{IB}} = 4\pi (R^3 - r^3)/3$, and letting the two vessel thicknesses be the same value $d = 125 \,\mu\text{m}$, we have

$$\frac{\rho^{\rm IV}}{\rho_0} = \frac{9R^2}{r(R^3 - r^3)} \left(\frac{\tau}{d} P_{\rm eff}\right)^2 \approx (1.04 \times 10^{11} \, {\rm s}^2 \, {\rm cm}^{-4}) P_{\rm eff}^2.$$
(4.33)

Even in the worst-case scenario of complete mixing inside moist nylon vessels, Table 4.13 shows that we can expect a reduction of > 10^7 in the radon concentration from the Outer Buffer volume to the scintillator in the Inner Vessel. To meet the requirement of no more than one count per day due to radon in the Fiducial Volume, we then require a radon concentration of no more than 1.5 Bq/m^3 in the Outer Buffer (a total radon count rate of about 900 Bq over its 580 m^3). Typical uranium concentration in dust is about 3 ppm = 35 Bq/kg [129], so with the assumption of secular equilibrium in the dust, this would be about 25 kg dust. If it were distributed as a thin layer settled onto the bottom surface of the Stainless Steel Sphere (horizontal projection in the xy-plane is 150 m^2), and using an estimated density for rock dust of 5 g/cm^3 , the layer of dust would be $\sim 30 \,\mu\text{m}$ thick.

Although that number sounds small, it is more typical of the amount of dust in a seldomcleaned basement than in a nominally class-10000 clean room. This result shows that radon diffusion from the Outer Buffer into the scintillator is not expected to be a problem in Borexino.

4.4.3 Radon emanation from nylon

Emanation of radon from the nylon vessels is a bit more of a concern. Formally, emanation may be treated in a similar way to diffusion. In this case an additional term is added to Equation (4.24) to describe the rate at which new radon atoms are produced by the decay of ²²⁶Ra embedded in the nylon. That is,

$$\rho_t = D\rho_{xx} - \rho/\tau + \mathcal{A} \tag{4.34}$$

where \mathcal{A} is the radium decay event rate per unit volume [atoms / (volume × time)].

Due to the long mean life of ²²⁶Ra ($\tau^{\text{Ra}} = 2310 \text{ yr}$), \mathcal{A} is essentially constant. However, in the experience of the Borexino project, radium in nylon is not likely to be in secular equilibrium with its progenitor uranium, due to their different chemical properties. Its abundance in the vessels must therefore be measured directly with sensitive counting techniques [123].

Let us once again assume the pessimistic case of complete mixing, and find the resulting stationary solution for the Inner Vessel. (The general time-dependent solution is treated in reference [123].) Suppose that the respective equilibrium radon concentrations in the Inner Buffer and scintillator, ρ^{IB} and ρ^{IV} , are negligible compared to $\mathcal{A}\tau$. Then the solution must be symmetric about x = d/2, go to zero at x = 0 and x = d, and in the limit of a thick slab of nylon, go to a constant value at the center: $\lim_{d\to\infty} \rho(d/2) = \mathcal{A}\tau$. The required result is found to be

$$\rho(x) = \mathcal{A}\tau \left(1 - \frac{\cosh\left[(x - d/2)/\ell\right]}{\cosh\left(d/2\ell\right)}\right).$$
(4.35)

A plot of this function is shown in Figure 4.15b.



Figure 4.15: Concentrations of radon in the nylon film for (a) diffusion (left), and (b) emanation (right). In both cases, the x-coordinate is plotted in units of the film thickness, $d = 125 \,\mu\text{m}$. The surfaces of the film are shown as vertical lines at x/d = 0 and x/d = 1. In the case of diffusion from outside (x < 0), the value of S is taken to be 3 for purposes of illustration. The upper curve is that for wet nylon ($\ell/d = 2.0$), the lower curve is that for dry nylon ($\ell/d = 0.088$), and the middle curve is an intermediate case ($\ell/d = 0.25$) corresponding to a relative humidity of about 55%. In the case of radon emanation, the upper curve is that for wet nylon. The horizontal bar at $\rho/A\tau = 1$ shows the radon concentration that would be obtained with no diffusion (all radon remaining trapped).

This solution yields a flux into the scintillator of

$$J(d) = -D\rho_x(d) = \mathcal{A}\ell \tanh(d/2\ell). \tag{4.36}$$

Substituting into the stationary form of Equation (4.26) and replacing R with r and V^{IB} with $V^{\text{IV}} = 4\pi r^3/3$ gives the expected equilibrium concentration of radon in the scintillator,

$$\rho^{\rm IV} = \frac{3\mathcal{A}\,\tau\ell}{r}\,\tanh\left(d/2\ell\right).\tag{4.37}$$

To find the equilibrium activity per unit volume in the scintillator \mathcal{A}^{IV} , we must divide again by the mean life τ . In the limit $\ell \ll d$, we then have the activity in the scintillator being $\mathcal{A}^{IV} \approx 3\mathcal{A}\ell/r$: it is proportional to the diffusion length. On the other hand, in the limit $\ell \gg d$, the activity in the scintillator becomes $\mathcal{A}^{IV} \approx 3\mathcal{A}d/2r$ (that is, half the radon generated in the film ends up in the scintillator), and only the thickness of the film is relevant.

From Table 4.10, the measured equilibrium 226 Ra activity in Sniamid (given its density of 1.14 g/cm^3 that may be extracted from [123]) is $\mathcal{A} = 18.2 \text{ mBq/m}^3$. The predicted



Figure 4.16: Expected radon decay event rate in the Fiducial Volume due to radon emanation from the Sniamid film making up the Inner Vessel. The rate, in decays per day (dpd), is shown as a function of relative humidity on the horizontal axis. Note the steep event rate increase beyond about 30-40% relative humidity. Graph reproduced from [45].

radon concentration in the scintillator from emanation is then $\rho^{IV} = 0.06 \text{ atoms/m}^3$ for dry Sniamid, and 0.38 atoms/m³ for wet Sniamid. These atomic concentrations translate to total count rates in the Fiducial Volume of 1.0 Rn decays/day and 7.0 Rn decays/day, respectively. Recall also that each radon decay is followed within a few hours by four more events further down the ²³⁸U decay chain. Though these can be statistically subtracted by tagging the decays of the radon daughters ²¹⁴Bi and ²¹⁴Po that occur in rapid succession, it will still be most important to keep the Inner Vessel dry once it has been filled with scintillator.

One may ask how the expected rate of radon decays in the Fiducial Volume behaves as the relative humidity changes. This can be predicted using the known behavior of the diffusion constant of radon in nylon. A graph of the expected radon decay rate as a function of relative humidity, assuming complete mixing, is reproduced from [45] in Figure 4.16. It can

be seen that the expected radon activity stays at a manageable level of ≤ 2 counts/day up to a relative humidity around 40%, above which it steeply increases.

One may also ask whether there is any danger from emanation of ²²⁰Rn ("thoron"), a product of the ²³²Th decay chain. Assuming secular equilibrium in that chain and using the ²³²Th concentration in Sniamid from Table 4.11 gives a ²²⁰Rn production rate of $\mathcal{A} = 17.1 \,\mathrm{mBq/m^3}$, almost the same as the production rate of ²²²Rn. The mean life is $\tau = 80.2 \,\mathrm{s}$, so the worst-case diffusion length in wet nylon is $\ell = 2.7 \,\mu\mathrm{m}$. These figures yield a predicted worst-case ²²⁰Rn activity of 0.3 ²²⁰Rn decays/day in the Fiducial Volume with the assumption of complete mixing. Though that assumption is unrealistic due to the short half-life of ²²⁰Rn, it may make some sense for the longest-lived decay product ²¹²Pb ($\tau_{1/2} = 10.6 \,\mathrm{h}$) and its daughters, some of whose energies fall into the neutrino energy window.

4.4.4 Desorption of radioactive contaminants from nylon

A potential danger that should not be underestimated is the problem of ²²²Rn daughter isotopes adhering to the surface of the Borexino vessels, and later "washing off" into the scintillator (desorption). These atoms are not intrinsic to the nylon film, but instead have attached themselves to it during vessel construction. During the few-year period between the beginning of vessel construction and the beginning of Borexino fluid operations, these atoms will have decayed into ²¹⁰Pb. It is the longest-lived ²²²Rn daughter, with a half-life of $\tau_{1/2} = 22.3$ yr.

²¹⁰Pb itself is not a problem for Borexino, since the endpoint of its β spectrum is only 64 keV. However, its immediate decay product, ²¹⁰Bi, has a β spectrum endpoint of 1.162 MeV, completely overlapping the neutrino window. This product will be formed at an almost constant rate, dependent upon the original amount of ²¹⁰Pb adhering to the nylon film, throughout the lifetime of the experiment. The same is true of the next isotope in the decay chain, ²¹⁰Po (although as an α particle emitter, it is less dangerous, for its decays may be mostly excluded from the data set by α/β separation.)

Two questions about ²¹⁰Pb must therefore be answered to determine whether it will be a problem for Borexino. First, what is the initial surface density of ²¹⁰Pb atoms on the Inner Vessel film surface? Second, given the planned filling scheme for the experiment (filling with water from the bottom, followed by filling with pseudocumene from above), what fraction of these atoms will remain attached to the film during the water filling (*i. e.*, not be drained with the water), and then later detach into the scintillator?

It should be noted that some concentration of ²¹⁰Pb and ²¹⁰Po atoms will also be added to the scintillator due to wash-off from the surfaces of the Borexino storage tanks, filling stations and filters. This section will discuss only the ²¹⁰Pb atoms originally adhering to the nylon film; other surfaces are discussed in detail in reference [46].

Expected surface concentration of ²¹⁰Pb/Po

The concentration of radon daughters is roughly constant near the center of a room. Toward the edges of the room, the walls, floor and ceiling act as a sink. It is impossible to analytically describe the situation, as the room may contain convection currents that move atoms about independently of diffusion. However, we suppose that near the surfaces in the room, the air is static so that the concentration of each species can be described by Fick's Law, Equation (4.18):

$$\mathbf{J} + D\nabla\rho = 0$$

If the concentration away from the room walls ρ_0 is constant in time and space, then the solution to the one-dimensional diffusion equation, Equation (4.22), near a wall is given by $\rho(x) = \rho_0 \left(1 - e^{-x/\ell}\right)$, where ℓ is roughly the thickness of the static air layer. Then, $J(0) = -D\rho_0/\ell$ (the negative sign simply describes the direction of flow). Defining a

"deposition velocity" as $v_d \equiv D/\ell$, we may rewrite the rate of deposition onto a surface as $J = v_d \rho_0$. Therefore, the accumulated surface density of ²¹⁰Pb atoms will be

$$\sigma(^{210}\text{Pb}) = \sum_{i=1}^{5} \rho_i v_d t, \qquad (4.38)$$

where

- $\sigma(^{210}\text{Pb})$ is the surface density of ^{210}Pb atoms,
- ρ_i is the concentration in air of the i^{th} species in {²¹⁸Po, ²¹⁴Pb, ²¹⁴Bi, ²¹⁴Po, ²¹⁰Pb},
- t is the elapsed time during which the surface is exposed to air.

In practice, it is difficult to measure the individual concentration of each isotope. As in [47], we may instead rewrite Equation (4.38) as follows:

$$\sigma(^{210}\text{Pb}) = k\rho_0 \bar{v_d} t \tag{4.39}$$

where ρ_0 is the room concentration of radon itself, \bar{v}_d is a weighted average of the deposition velocities for each species, and k is the necessary proportionality constant. An effective deposition velocity may be defined as $v_0 \equiv k\bar{v}_d$.

From an experiment done with 5 cm × 5 cm pieces of nylon film in a 2 m × 2 m × 3 m test clean room with a high radon concentration [130], an effective deposition velocity $v_0 = 3 \times 10^{-8}$ m/s was derived [47], three orders of magnitude smaller than the average deposition velocity for radon daughters in a normal room. The small value of k causing this is most likely due to the atoms never reaching the nylon surface, instead being removed by the air filters. Therefore we expect v_0 to scale inversely with λ_f , the number of clean room volumes filtered per unit time. Unfortunately this number was not available for the test clean room. Measurement of the flow rate of a HEPA filter identical to that in the test clean room yielded 90 ± 10 cfm [131], giving a value for λ_f of $(3.5 \pm 0.4) \times 10^{-3}$ s⁻¹. In the Borexino vessel construction clean room, $\lambda_f = 0.036$ s⁻¹ [44], so we expect v_0 there to be smaller by a factor of 10: $v_0 \approx 3 \times 10^{-9}$ m/s. The typical radon activity in the vessel construction clean room was maintained at 1.5– 2 Bq/m^3 through the use of an activated charcoal vacuum swing adsorption filter [44, 132], for a concentration of about 10⁶ Rn atoms/m³. The nylon vessels, during construction, were kept sandwiched between two other sheets of film as much as possible. Even in the final stages of assembly, they remained covered by plastic sheeting when not being actively worked upon. We may estimate that the average exposure time for any individual surface area on the Inner Vessel was only about 1 hour [44]. Assuming the value of v_0 obtained above, the resulting estimated ²¹⁰Pb surface density will be on the order of 10 atoms/m², for a total of ~2500 ²¹⁰Pb atoms on the inner surface of the Inner Vessel. In a worst-case scenario (complete desorption into pseudocumene and thorough mixing), these would contribute 0.07 decays/day in the Fiducial Volume from each of ²¹⁰Pb and ²¹⁰Po.

²¹⁰Pb/Po desorption into water and pseudocumene from nylon

Various substances, including nylon, teflon, and stainless steel, have been exposed to a radon-laced atmosphere, and then soaked in water for different amounts of time to determine what fraction of radon daughters are removed. The behavior of both ²¹⁰Pb and ²¹⁰Po have been studied; the former with a γ -ray germanium detector, and the latter with a silicon surface barrier α -particle detector. For the Ge detector, the data acquisition system and software were purchased with the system; for the Si detector, data acquisition was performed with a CAMAC crate controller and in-house custom software.

These tests with nylon samples exposed to radon for several months have shown that a ten-minute soak in deionized water is sufficient to remove 75% of ²¹⁰Pb atoms from a nylon surface [132]. Although it seems that 5–10% of the ²¹⁰Pb may be left even after ten days of soaking in water, this measurement is statistically dubious due to a low initial activity. In any case, assuming the value for v_0 given above, at most ~0.01 decays/day of lead should be present in the Fiducial Volume if it is decided to fill the detector with water for a few months (an optional operation) before scintillator and buffer fluids are introduced. This is about

a factor of four lower than the estimate of 0.04 decays/day arrived at through a different method of calculation in [44]. If the water-filling is not performed, though, the expected rate of 0.07 decays/day in the Fiducial Volume is nevertheless more than acceptable.

²¹⁰Po is slower to come off in water, with the amount remaining exhibiting only a logarithmic dependence on elapsed time. However, as mentioned already, ²¹⁰Po events can be efficiently discriminated by α/β separation so do not pose nearly as much of a problem. In any event, desorption of lead and polonium atoms from steel surfaces in the Borexino filling stations is expected to be more potentially dangerous than that from the nylon film itself.

4.5 Solubility of water in pseudocumene

It was demonstrated in Section 4.3.2 that the moisture content of nylon film exposed to air at a certain relative humidity is equal to that in nylon film exposed to pseudocumene at the same relative humidity. That is, if a sample of air has 30% relative humidity, and a sample of pseudocumene is exposed to air at 30% relative humidity, the two samples are equivalent environments as far as a strip of nylon is concerned. This suggests a few other questions important in Borexino:

- What is the maximum amount of water (the solubility) that can be dissolved in pseudocumene?
- Can this saturation level of water in pseudocumene be achieved by holding the pseudocumene in air at 100% relative humidity?
- Is the water content of pseudocumene a linear function of the relative humidity of the surrounding air?
- What is the temperature dependence of the solubility of water in pseudocumene?

4.5.1 Water solubility in aromatic hydrocarbons

No previous tests of the solubility of water in pseudocumene were found in the literature. However, the solubilities of water in similar aromatic hydrocarbons such as benzene (C₆H₆, molecular mass = 78.1 amu) and ethylbenzene (C₈H₁₁, molecular mass = 107.1 amu) have been measured, and even fit to a function of temperature. The solubility of water in a hydrocarbon is usually expressed as a mole fraction x_w . If x_w is much less than one, this value can be converted to a mass fraction simply by multiplying by the ratio of the molecular mass of water (18.0 amu) to that of the hydrocarbon.

One semi-empirical formula that has been derived to express the temperature dependence of x_w is given by

$$\log x_w = A + B/T + CT \tag{4.40}$$

("log" is the natural logarithm, and T is the temperature in Kelvins). For benzene, the coefficients have been determined to be A = -1.64055, B = -2029.41 K, and $C = 9.00544 \times 10^{-3}$ K⁻¹ [133]. At 22°C, this formula yields $x_w = 2.84 \times 10^{-3}$ (655 ppm), and at 10°C (the approximate temperature in the Gran Sasso underground lab), $x_w = 1.91 \times 10^{-3}$ (439 ppm).

A different semi-empirical formula has been put forth to describe a divergent high-temperature behavior of the water solubility at the temperature T_{3c} of the three-phase critical end point. This formula is given by¹¹

$$\log x_w = \alpha + \beta \left(\frac{T_{3c}}{T} - 1\right) + \gamma \left(1 - \frac{T}{T_{3c}}\right)^{1/3} + \delta \left(1 - \frac{T}{T_{3c}}\right).$$

$$(4.41)$$

The coefficients of this equation for water in ethylbenzene (with $T_{3c} = 568.1$ K) are given as $\alpha = -0.37215$, $\beta = -4.4626$, $\gamma = -0.38596$, and $\delta = -2.59850$ [134]. At 22°C, Equation (4.41) for the solubility of water in ethylbenzene yields $x_w = 2.34 \times 10^{-3}$ (394 ppm), and at 10°C, $x_w = 1.54 \times 10^{-3}$ (258 ppm).

¹¹Neglecting terms of order $(T/T_{3c})^2$ and higher, at low temperatures this formula may be rewritten in the form of Equation (4.40) by setting $A = \alpha - \beta + \gamma + \delta$, $B = \beta T_{3c}$, and $C = -(\gamma/3 + \delta)/T_{3c}$.

From these data, assuming that the value of x_w for water in pseudocumene is similar to that for water in benzene or ethylbenzene, we would predict a water content of 350–430 ppm in saturated pseudocumene at room temperature.

4.5.2 Experimental tests of water solubility in pseudocumene

The water content of the tank of pseudocumene in glove box B of the Controlled Relative Humidity Facility (described in Section 4.3.2) was measured at each of the reference relative humidity levels, 0%, 10%, 20%, 40%, and 60%. This was done only with the Computrac 3000, not by measuring the sample masses, since the mass increase was expected to be too small to determine accurately. In this case, the temperature of the Computrac 3000 oven was set at 50°C, as recommended by Arizona Instruments for measurements of moisture content in aromatic compounds. The cutoff rate for each measurement was set to be the same as for the nylon measurements, $0.30 \,\mu g/s$ of water vapor.

The fractional water content of pseudocumene is much smaller than in nylon at the same relative humidity, so it was necessary to be more careful. First a blank sample of air was measured with the Computrac 3000. Then the same vial was re-measured without opening, thereby purging virtually all water vapor from the air inside. At this point a few m ℓ of pseudocumene from the tank in glove box B were placed into the vial through the septum (a syringe was used so that the vial need not be opened). The same vial was measured twice more, and the sum of water collected in these two measurements was assumed to be the water contained in the pseudocumene. Finally, a fifth measurement was made on the vial to check for consistency. It was assumed that the minuscule mass of water measured in the second and fifth steps was due to diffusion, and therefore should be subtracted from the mass of water measured in the pseudocumene in the third and fourth steps. If the water content measured in the second or fifth step was much more than 10 μ g, the septum of the vial was presumed to be damaged and the measurement was discarded. At least three data points were collected in this way at each relative humidity.

Humidity	Moisture content			
[%]	$[ppm H_2O]$			
0	-1.1	(3.2)		
10	37.8	(4.3)		
20	75.7	(4.8)		
40	171.0	(13.7)		
60	277.6	(1.9)		

Table 4.14: Moisture content of pseudocumene as a function of relative humidity in the surrounding atmosphere. These values were determined with the Computrac 3000. The percentages are relative to the dry mass. Statistical errors are shown in parentheses.



Figure 4.17: Moisture content (ppm) of pseudocumene as a function of humidity.

Results are tabulated in Table 4.14. The fractional moisture content of the pseudocumene, in ppm, is shown as a function of the relative humidity in Figure 4.17. It is close to linear, as expected. A χ^2 fit of the pseudocumene data in Table 4.14 to a linear function gives the expected value of water content in pseudocumene at 100% relative humidity as 465 ppm. At 10°C, the saturation content of water in both benzene and ethylbenzene is a factor very close to 1.5 times smaller than that at 22°C. If the same holds true for pseudocumene, than its water content at 100% relative humidity and 10°C should be 310 ppm.

4.6 Conclusions regarding the nylon films

It should be clear from the results of this chapter that in order to handle the nylon film without danger of it cracking, the film must (1) be above the glass transition point (in its plastic phase), and (2) be significantly thinner than the critical thickness at that point. Both Capron and Sniamid films appear satisfactory with regard to the second point when limited to a thickness of $125 \,\mu$ m. The 500 μ m C38F film that was used in the CTF2 prototype vessel was near or above its critical thickness at low humidity. In fact, that vessel was stored for a long time in an atmosphere of dry nitrogen. The concept of critical thickness explains its cracking failure well.

At room temperature, in order to have a large margin of safety, the first requirement can be satisfied by staying at a high controlled relative humidity. All types of nylon appear to have completed the glass to plastic transition at humidity greater than about 40%. Thus, during construction of the Borexino vessels, the clean room in Jadwin Hall, which is at a constant temperature of 18°C, was held at a relative humidity of 45 to 50%. Humidity in the underground lab at Gran Sasso was chosen to be even higher, at 50 to 60%, since the lab is colder than room temperature, about 10°C. During construction, installation, and inflation of the nylon vessels, the room humidity was supplemented by constantly spraying the accessible film surfaces with atomizers containing deionized water. Once Borexino has been filled with pseudocumene scintillator, the possibility of dehydrating the pseudocumene is foreseen. With the vessels fully inflated, the probability of creases developing is low. It therefore makes sense to worry more about possible deformation of the vessels than about brittle failure. In the glassy state, nylon is stronger and less susceptible to irreversible deformation (creep). An additional very important benefit to this idea is that dry nylon film presents much more of a barrier to the diffusion and emanation of dissolved radon gas. A good final recommendation for the operating humidity would seem to be about 30%: at this level, the expected activity due to emanation from the Inner Vessel has not yet started to increase sharply. The film is reasonably strong (30-40 MPa tensile strength) but not yet too brittle (400-500 μ m critical thickness). A level of 30% relative humidity represents a good compromise between the increased radon emanation and lower tensile strength at higher humidities, and the increased brittleness when the film is dry.

To implement this humidity level, the pseudocumene moisture content should be fixed at 30% of its saturation value at 10°C. That is, it should be set at roughly 95 ppm. This may be monitored with the Nyad hygrometer and the Computrac 3000, both of which will be present at Gran Sasso. In fact, the moisture content may actually be increased a bit from this value, since the nylon film tends to remain in the glass phase at higher humidities when the temperature is low.

The final selection of the nylon types used in the Borexino vessels was ultimately made for chronological reasons. Originally, it appeared that Capron was the better choice due to its superior performance in the packet tests and its marginally greater critical thickness. The Outer Vessel is in fact made of Capron. However, measurements undertaken by the Borexino Heidelberg group [123] show that the Sniamid will emanate less radon than the Capron by a factor of ten. Both Capron and Sniamid are within acceptable limits for mechanical properties; however, it is important for Borexino to be as radiopure as possible, so the Inner Vessel was constructed from Sniamid instead.

Chapter 5

Position Reconstruction in Scintillation Detectors

Borexino is one of a new generation of ultra-low-background scintillator-based detectors. Such detectors are widely used for the detection of weakly interacting particles. At present the main focus of observation is on neutrinos and antineutrinos from various sources, but there are also plans to construct large optical detectors to search for as yet undiscovered particles such as WIMPs. The detection mechanism is based on the collection of visible or ultraviolet photons. These are emitted as Čerenkov radiation (*e. g.*, as in Kamiokande [135] and SNO [136]) or as scintillation photons. This chapter will focus on a likelihood-based method of position reconstruction for scintillator-based, unsegmented detectors, and the spatial resolutions that may be expected from the method.

5.1 The need for spatial reconstruction

Due to the extremely low interaction rates of neutrinos and their antiparticles (to say nothing of WIMPs and so forth), it is necessary for a detector to contain a large mass of scintillator with very low levels of internal radioactive contamination [55]. Ultra-pure materials are also used in order to screen radioactivity from materials surrounding the detector [55, 63]. Unfortunately, as already mentioned in Section 3.3, the photosensitive elements used to detect scintillation light are notorious for being among the main sources of radioactivity in an ultra-low-background detector.

It is therefore desirable to insert one or more layers of buffer material between the photosensitive elements and the scintillator to suppress radioactive background. Often the buffers are inactive, *i. e.*, not scintillating. An inactive buffer offers the advantage of minimizing the total trigger rate caused by the abundant radioactive decays generally produced within the photosensitive elements. Since the compositions of the scintillator and inactive buffer are different, a scintillator containment system analogous to the Borexino Inner Vessel is required to physically separate them. The containment system, being in direct contact with the scintillator, must satisfy stringent intrinsic radiopurity requirements.

For additional background prevention, the outer region of the scintillator volume can be used as an active buffer. This allows any residual radioactivity coming from the containment system, or passing through it, to be monitored and suppressed. A "fiducial volume" is commonly defined as a region at the center of the active volume of the detector in which radioactive background is expected to be at a minimum. The discrimination between events belonging to the fiducial and to the non-fiducial regions is performed by means of software implementation (reconstruction code) of an algorithm (reconstruction algorithm), which assigns to each single event a reconstructed position, either inside or outside the fiducial volume. The algorithm also provides a means of comparing the position of different events and is an important tool for the identification of several background sources. The designs of some planned detectors incorporate only a thin inactive buffer region or none at all, and in these cases, correct assignment of an event as belonging to the fiducial volume or the buffer region is even more important. The resolutions of detector reconstruction codes are generally studied with Monte Carlo methods. Event simulations allow close reproductions of the performance of these codes on real events. Typically, however, the reconstruction codes are fine tuned by calibrating the detector with the use of localized sources of radioactivity or light.

What seems lacking from the available literature is a comprehensive discussion of how the resolutions of detector reconstruction codes are related to some basic properties of the detector: the linear dimension, the time dispersion of the photon emission, the scintillator index of refraction, possible processes of absorption and re-emission and of scattering of the scintillator light, etc. Though some results are presented in [137], their form is complex due to the inclusion of various second-order effects such as light scattering and angular dependence of the photosensitive elements. This chapter presents an analytic study of the resolution for reconstruction in time and space of scintillation events. For simplicity, this study is restricted to the case of events at the center of the detector. Other analytical studies [137], full Monte Carlo simulations [80], and calibrations of existing experiments [138] show that the resolution of the reconstruction codes depends only mildly upon the location of the scintillation event.

This study also assumes that the optical properties of the media are uniform throughout the detector, that the indices of refraction of all materials between the active scintillator and the photodetectors are approximately the same, and that to a first approximation the scattering of light may be neglected.

5.2 Likelihood function derivation

The likelihood function is a standard statistical tool used to find parameters of a physical model. Suppose a set of N observations is composed of the independent values $\{t_i\}$ and dependent values $\{s_i\}$. For instance, $\{t_i\}$ could be a list of times at which a radioactive source is observed, and $\{s_i\}$ a list of observed activities at each time. We wish to model the data by some function f(s) with n free parameters, represented by the n-vector \mathbf{a} . In the example, the function would be a decaying exponential, and the parameters would be the initial activity and the half-life. The likelihood function \mathcal{L} over the parameters is defined as a probability distribution for obtaining the observed data, given that the parameters \mathbf{a} have the specific values \mathbf{a}_0 :

$$\mathcal{L}(\mathbf{a}_0; \{(t_i, s_i)\}) d^n s d^n t = P(\{(t_i, s_i)\} \text{ are observed } | \mathbf{a} = \mathbf{a}_0).$$
(5.1)

The infinitesimals on the left-hand side result from the definition of \mathcal{L} as a probability distribution function. The construct "P(A|B)" is a standard way to abbreviate "the probability of condition A occurring, given that condition B is true."

The difficult task is to calculate this probability based on the assumption that the data are correctly described by the model function f(s). Once this has been done, in order to calculate the most probable value of the parameters of the model, one simply finds the maximum of the likelihood function (or, as is usually computationally easier, the minimum of $-\log \mathcal{L}$) in the *n*-dimensional space defined by the free parameters **a**.

In the case of a scintillator-based detector, the parameters of interest are the position and time of an event in the detector, $\mathbf{a} = (\mathbf{x}_0, t_0)$. The observed data are the positions $\{\mathbf{x}_i\}$ of the photosensitive elements, usually PMTs (independent values), and the times $\{t_i\}$ at which each element detects a photon (dependent values); *i* ranges from 1 to *N*, with *N* being the number of detected photons. For now we assume that at most one photon is seen by each PMT, so all the \mathbf{x}_i 's are distinct, and *N* is also the number of PMTs that detect a photon. For conciseness, define the following possible conditions:

- A : an event occurs in the detector at position and time (\mathbf{x}_0, t_0)
- B : the positions and times at which PMTs detect photons are $\{(\mathbf{x}_i, t_i)\}$.

Then, Equation (5.1) becomes $\mathcal{L}(\mathbf{x}_0, t_0; \{(\mathbf{x}_i, t_i)\}) d^{3N} \mathbf{x} d^N t \equiv P(B|A).$

5.2.1 Factoring the detector likelihood function

Let us assume that the times at which photons are emitted by the scintillator are uncorrelated. Then the likelihood function will have one independent factor for the piece of data provided by each PMT.¹ Let the total number of working PMTs be T, so that N PMTs (labeled $1, \ldots, N$) have detected a photon, and T - N PMTs (labeled $N + 1, \ldots, T$) have not. If we further define

- C_i : A photon arrives at PMT *i* at an unspecified time
- D_i : PMT *i* detects a photon (*i. e.*, converts it to a photoelectron)
- E_i : PMT *i* detects a photon at the specific time t_i ,

then

$$P(B|A) = \prod_{i=1}^{N} P(E_i|A, C_i, D_i) P(D_i|A, C_i) P(C_i|A)$$
$$\times \prod_{j=N+1}^{T} [P(\neg D_j|A, C_j) P(C_j|A) + P(\neg C_j|A)]$$
(5.2)

(where \neg is the logical negation symbol). Of course, $P(D_i|A, C_i)$ is just the total efficiency q_i of PMT *i*, which for simplicity will be supposed independent of the original event position. This assumption is reasonable if, for instance, the PMTs are mounted at a distance from the fiducial volume of the detector, so arriving photons always have a small angle of incidence.

Now define a "per-PMT" likelihood function \mathcal{L}_i (we absorb the infinitesimals into the function definition for later convenience),

$$\mathcal{L}_{i}(\mathbf{x}_{0}, t_{0}; \mathbf{x}_{i}, t_{i}) = \begin{cases} q_{i} P(\mathbf{E}_{i} | \mathbf{A}, \mathbf{C}_{i}, \mathbf{D}_{i}) P(\mathbf{C}_{i} | \mathbf{A}), & i \leq N \\ (1 - q_{i}) P(\mathbf{C}_{i} | \mathbf{A}) + P(\neg \mathbf{C}_{i} | \mathbf{A}), & N < i \leq T. \end{cases}$$

$$(5.3)$$

¹Strictly speaking, this is not precisely true. We assume in this discussion that exactly N PMTs detected photons, instead of making the more basic assumption that exactly Γ scintillation photons were emitted, which would lead to the number of PMTs that see photons having a Poissonian distribution with some mean value $\bar{N}(\Gamma)$. With our assumption of N photons detected, the PMT hit data are in fact weakly correlated. For N reasonably large, though, the difference when $N = \bar{N}$ should be negligible. It would be interesting to compare results derived from the often-used Poisson and multinomial probabilistic models to the model put forth here.

The total likelihood function (times infinitesimals) is then the product of all per-PMT likelihood functions. Notice that the per-PMT likelihood function of a supposedly dead PMT ($q_i = 0$) that does not detect a photon reduces to 1, so does not influence the total likelihood function, just as expected.

5.2.2 Scintillator dispersion time at the emission point

The first non-trivial factor in the expression for the likelihood function of a PMT that detects a photon is based solely on timing information of a photon emitted by the scintillator. Scintillation photons are emitted, as discussed in Section 3.1, as a consequence of the ionization of the scintillator due to interacting particles or radioactive decays. The typical dispersion in the time of emission of organic liquid scintillators is on the order of a few nanoseconds, with a slower component that can reach hundreds of nanoseconds. Since electronically excited scintillator molecules retain no memory of the original direction of the ionizing particle, their emission of photons is isotropic. In this discussion we also assume that the time of emission of each photon, relative to the time of the event causing scintillation, is an independent random variable τ_e with probability density function $p(\tau_e)$.

We also refer to the function $p(\tau_e)$ as the scintillator response function. Its experimental determination will be described a bit in Section 7.5; in this chapter we take it as a given. Referring to the left half of Figure 5.1, one sees that at a specific time t, this function may also be regarded as an outgoing spherical photon probability wave, integrated over the solid angle 4π . In fact, the most important factor in Equation (5.2), the probability $P(E_i|A, C_i, D_i)$, is equal to it. Let τ_f^i be the time of flight from the origin \mathbf{x}_0 of the photon to the position \mathbf{x}_i of the i^{th} PMT. Then, with n being the scintillator index of refraction, we have:

$$\tau_f^i = \frac{|\mathbf{x}_i - \mathbf{x}_0| n}{c} \tag{5.4}$$

$$t_i = \tau_e + \tau_f^i + t_0. (5.5)$$

$$\mathcal{L}_i(\mathbf{x}_0, t_0; \mathbf{x}_i, t_i) \propto p(t_i - t_0 - \tau_f^i).$$
(5.6)

Of course, factors other than the dispersion time of the scintillator may also affect the probability distribution function of the recorded arrival times of photons at PMTs. The most important other effects are usually the effects of scattering and absorption and re-emission processes in the scintillator, as well as the finite time resolution of the PMTs themselves. The latter may in general be incorporated into the distribution $p(\tau_e)$ by convolution with the scintillator dispersion function. The former require a bit more care because scattering effects depend in general upon the light path length from the event to the PMT; a full treatment of these effects is beyond the scope of this chapter. Some description of them may be found in, for instance, references [75, 80].

5.2.3 Photon attenuation

As photons travel away from their origin, they are attenuated by the familiar inverse square law. This implies a formula for the probability $P(C_i|A)$ that a given PMT is hit by a scintillation photon. Suppose a PMT of infinitesimal area, at a distance $s_i \equiv |\mathbf{x}_i - \mathbf{x}_0|$ from the event, subtends a solid angle $d\Omega_i$ as seen from the event location. It will be struck by only a fraction $d\Omega_i/4\pi$ of all photons emitted. So if Γ photons were emitted, its probability of being hit by at least one of them is

$$P(C_i|A) = 1 - \left(1 - \frac{d\Omega_i}{4\pi}\right)^{\Gamma} \approx \Gamma \frac{d\Omega_i}{4\pi}.$$
(5.7)

If the i^{th} PMT has an area dA_i and is tilted away from the line of sight by an angle ψ_i , as shown on the right half of Figure 5.1, then $d\Omega_i = \cos \psi_i dA_i/s_i^2$, so the resulting factor in the likelihood function is given by

$$\mathcal{L}_{i}(\mathbf{x}_{0}, t_{0}; \mathbf{x}_{i}, t_{i}) \propto \Gamma \frac{\mathrm{d}\Omega_{i}}{4\pi} = \Gamma \frac{\mathrm{cos}\,\psi_{i}}{4\pi s_{i}^{2}}\,\mathrm{d}A_{i}.$$
(5.8)



Figure 5.1: Geometry of the likelihood function derivation. The concentric dotted lines, and the graph on the left, represent the probability function (an expanding spherical wave) of the emission time of a scintillation photon. The rectangle labeled dA_i represents a PMT of infinitesimal size with normal vector $\hat{\mathbf{n}}$, subtending a solid angle $d\Omega_i$ as seen from the position of the detector event. The PMT is tilted away from the direction of the event by an angle ψ_i . Note that we have not yet made any assumptions about the geometry of the detector.

As mentioned already, all constant factors in a likelihood function may be discarded with no effect on the location in parameter space of its maximum. To first order, this includes the efficiency q_i of each PMT. The per-PMT likelihood function for a PMT detecting a photon may thus be redefined as

$$\mathcal{L}_{i}(\mathbf{x}_{0}, t_{0}; \mathbf{x}_{i}, t_{i}) = p(t_{i} - t_{0} - \tau_{f}^{i}) \frac{\cos \psi_{i}}{s_{i}^{2}}.$$
(5.9)

Its logarithm is

$$\log \mathcal{L}_{i} = \log p(t_{i} - t_{0} - \tau_{f}^{i}) + \log \cos \psi_{i} - 2 \log s_{i}.$$
(5.10)

5.2.4 The PMTs not triggered

For completeness, we now consider the case of a PMT that does not detect a photon produced by an event in the detector. Its per-PMT likelihood function, from Equation (5.3), is given by

$$\mathcal{L}_{i}(\mathbf{x}_{0}, t_{0}) = (1 - q_{i}) \mathrm{P}(\mathrm{C}_{i} | \mathrm{A}) + \mathrm{P}(\neg \mathrm{C}_{i} | \mathrm{A})$$

$$= (1 - q_{i}) \left[1 - \left(1 - \frac{\mathrm{d}\Omega_{i}}{4\pi} \right)^{\Gamma} \right] + \left(1 - \frac{\mathrm{d}\Omega_{i}}{4\pi} \right)^{\Gamma}$$

$$= 1 - q_{i} + q_{i} \left(1 - \frac{\mathrm{d}\Omega_{i}}{4\pi} \right)^{\Gamma}$$

$$\approx 1 - q_{i} \Gamma \frac{\mathrm{d}\Omega_{i}}{4\pi}.$$
(5.11)

The logarithm of this per-PMT likelihood function is $\approx -q_i \Gamma d\Omega_i/4\pi$. This term, containing an infinitesimal, is negligible in size compared to the terms of Equation (5.10) coming from per-PMT likelihood functions for PMTs that have detected a photon. If PMTs are in fact very small compared to any other relevant dimensions of the detector, it may therefore be ignored.

5.2.5 Specialization to a spherical detector

As written, Equation (5.9) is applicable to any detector with pointlike PMTs forming the vertices of a convex polyhedron (so that light from an event at any point inside the detector may reach any one of the PMTs). Let us specialize to a spherical detector centered at the origin, having a uniform distribution of inward-facing PMTs at distance R from the detector center. The radius of the fiducial volume of the detector itself may be less than R; it is not an important quantity for this discussion. As above, we call the distance from an event to the i^{th} PMT $s_i \equiv |\mathbf{x}_i - \mathbf{x}_0|$. Let the distance from the center of the detector to the event be $a \equiv |\mathbf{x}_0|$, so we have the geometry of Figure 5.2.


Figure 5.2: Geometry of a spherical detector.

By dropping a perpendicular from point A to segment OB at point C, one readily sees that $s_i \cos \psi_i = R - a \cos \theta_i$, with θ_i being the angle between the event and the i^{th} PMT as seen from the origin. Hence the likelihood function becomes

$$\mathcal{L}(\mathbf{x}_0, t_0; \{(\mathbf{x}_i, t_i)\}) = \prod_{i=1}^N p\left(t_i - t_0 - \frac{s_i n}{c}\right) \frac{R - a\cos\theta_i}{s_i^3}$$
(5.12)

where s_i is given by the Law of Cosines,

$$s_i^2 = R^2 + a^2 - 2aR\cos\theta_i. \tag{5.13}$$

5.3 Analytical treatment of the likelihood function

It may be of interest to examine properties of the likelihood function in the particular case of a hypothetical event occurring at the center of a spherical detector. This allows the general nature of the problem of reconstruction to be understood analytically. For simplicity, let's assume that the distribution of the time emission of the photons is a Gaussian curve with width equal to the characteristic dispersion time of the scintillator:

$$p(\tau_e) = \frac{e^{-\tau_e^2/2\sigma^2}}{\sqrt{2\pi\sigma^2}}; \ \log p(\tau_e) = \text{const} - \frac{\tau_e^2}{2\sigma^2}.$$
 (5.14)

The same equation can also be used for the case when the original scintillation light is absorbed and then re-emitted by scintillation fluors in the immediate proximity of the energy deposition point, as with PPO in pseudocumene. In this case, the dispersion characteristic of the scintillator is effectively broadened by the absorption and re-emission process.

5.3.1 Taylor expansion of the likelihood function

For a point in the detector at a distance a from the center, in the direction of a particular unit vector $\hat{\mathbf{u}}$, the log likelihood function is

$$\log \mathcal{L}(a\hat{\mathbf{u}}, t_0) = \text{const} - \frac{1}{2\sigma^2} \sum_{i=1}^{N} \left(t_i - t_0 - \frac{s_i n}{c} \right)^2 + \sum_{i=1}^{N} \log \frac{R - a \cos \theta_i}{s_i^3}$$
(5.15)

where s_i and θ_i for each PMT are as shown in figure 5.2. We assume that the number of hit PMTs N is sufficiently large that we can, with little error, replace this expression by spatial and temporal averages over the expected angular and time distributions of the PMT hits. That is (discarding the constant term),

$$\log \mathcal{L}(a\hat{\mathbf{u}}, t_0) \approx -\frac{N}{2\sigma^2} \left\langle \left(t - t_0 - \frac{sn}{c}\right)^2 \right\rangle + N \left\langle \log \frac{R - a\cos\theta}{s^3} \right\rangle,$$
(5.16)

where t, s, θ are now continuous random variables with the expected distributions. We now calculate these averages for a point-like event located in the center $\mathbf{x}_0 = \mathbf{0}$ of the detector, occurring at time $t_0 = 0$.

First consider the time average. The time of flight of photons from the center to each PMT (assuming minimal scattering) is Rn/c, where n is the index of refraction and c is the velocity of light in vacuum. This means that the distribution curve of t is p(t - Rn/c). From the properties of a Gaussian distribution, the time averages of time-dependent quantities are

$$\langle t \rangle = \frac{Rn}{c} \tag{5.17}$$

$$\langle t^2 \rangle = \langle t \rangle^2 + \sigma_t^2 = \frac{R^2 n^2}{c^2} + \sigma^2.$$
 (5.18)

Likewise, since all PMTs are equidistant from an event at the center of a spherical detector, the distribution of PMT hits should be uniform over the solid angle. Hence the spatial averages over quantities dependent upon the event-to-PMT angle θ can be found using Equation (5.13) and taking the surface integral over the sphere of PMTs:

$$\langle s \rangle = \frac{1}{4\pi} \int d\phi \, d(\cos\theta) \, \sqrt{R^2 + a^2 - 2aR\cos\theta} = R + \frac{a^2}{3R} \tag{5.19}$$

$$\left\langle s^{2}\right\rangle = \frac{1}{4\pi} \int \mathrm{d}\phi \,\mathrm{d}\left(\cos\theta\right) \left(R^{2} + a^{2} - 2aR\cos\theta\right) = R^{2} + a^{2} \tag{5.20}$$

Finally, we observe that for a point-like event in the center of a uniform sphere of PMTs, there is no correlation between the expected spatial distribution of s and temporal distribution of t; that is, $\langle st \rangle = \langle s \rangle \langle t \rangle$. This and the above equations allow us to evaluate

$$\left\langle \left(t - t_0 - \frac{sn}{c}\right)^2 \right\rangle = \left\langle t^2 + t_0^2 + \frac{s^2 n^2}{c^2} - 2tt_0 - 2t\frac{sn}{c} + 2t_0 \frac{sn}{c} \right\rangle$$
$$= \frac{R^2 n^2}{c^2} + t_0^2 + (R^2 + a^2)\frac{n^2}{c^2} - 2\frac{Rn}{c}t_0$$
$$- 2\frac{Rn^2}{c^2}(R + \frac{a^2}{3R}) + 2t_0(R + \frac{a^2}{3R})\frac{n}{c}$$
$$= \operatorname{const} + t_0^2 + \frac{n^2}{3c^2}a^2 + \frac{2n}{3cR}a^2t_0$$
(5.21)

where the constant term contains whatever does not depend explicitly on t_0 and a.

The quantity averaged over in the last term of Equation (5.16), again substituting in Equation (5.13), becomes

$$\log \frac{R - a \cos \theta_i}{s_i^3} = \log \left(\frac{R - a \cos \theta_i}{(R^2 + a^2 - 2aR \cos \theta_i)^{3/2}} \right) = -2 \log R + \frac{2a}{R} \cos \theta_i + \frac{a^2}{2R^2} \left(5 \cos^2 \theta_i - 3 \right) + \dots$$
(5.22)

with the last equality above being the expansion into a Taylor series in a/R.

By once again averaging the expected distributions in s and θ over the solid angle, the result, obtained to second order in a/R, is determined to be

$$\left\langle \log \frac{R - a \cos \theta}{s^3} \right\rangle \approx \text{const} - \frac{2a^2}{3R^2}.$$
 (5.23)

The complete likelihood function for an event at the center of a spherical detector, to second order in a/R, is thus

$$\log \mathcal{L}(a\hat{\mathbf{u}}, t_0) \approx \text{const} - N \left[\frac{1}{2\sigma^2} \left(t_0^2 + \frac{n^2}{3c^2} a^2 + \frac{2n}{3cR} a^2 t_0 \right) + \frac{2}{3R^2} a^2 \right].$$
(5.24)

5.3.2 Likelihood function maximum and resolutions

Solving for the maximum of the likelihood function and requiring |a| < R gives the expected solutions:

$$\begin{cases} \frac{\partial}{\partial t_0} \log \mathcal{L} = 0 \\ \frac{\partial}{\partial a} \log \mathcal{L} = 0 \end{cases} \iff \begin{cases} t_0 = 0 \\ a = 0 \end{cases}$$
(5.25)

We next ask about the expected resolution of the detector. Notice that the information matrix is diagonal because the off-diagonal terms, $-\partial^2(\log \mathcal{L})/\partial a \partial t_0$, are zero when $a = t_0 = 0$. The theoretical resolutions of the detector in space and time are therefore given by reciprocals of the second derivatives of the likelihood function:

$$\begin{cases} \delta t_0 = \left(-\frac{\partial^2 \log \mathcal{L}}{\partial t_0^2}\right)^{-1/2} = \frac{\sigma}{\sqrt{N}} \\ \delta a = \left(-\frac{\partial^2 \log \mathcal{L}}{\partial a^2}\right)^{-1/2} = \left(\frac{Nn^2}{3c^2\sigma^2} + \frac{4N}{3R^2}\right)^{-1/2} \end{cases}$$
(5.26)

When the detector dimensions are much larger than the scintillator dispersion time, $R \gg c\sigma/n$, we can approximate $\delta a \approx \sqrt{\frac{3}{N}} \frac{c\sigma}{n}$. (It should be noted that this does not take into account scattering effects, which become increasingly important with larger detectors.) After the appropriate simplifications, this is in good agreement with Equation (64) of reference [137].

Because of the spherical symmetry of the problem, δa can be used as a stand-in for any of the three Cartesian spatial resolutions δx_0 , δy_0 , δz_0 . One may, for instance, make the substitution $a^2 = x_0^2 + y_0^2 + z_0^2$ in Equation (5.24) and obtain the same results for the resolution in each Cartesian coordinate.

5.3.3 Pattern matching

In case of use of a liquified noble gas as scintillator, as in the new generation of solar neutrino detectors [139, 140], Rayleigh scattering of the ultraviolet scintillation photons plays an important role. The photons are scattered intensely by the medium, such that they effectively diffuse out of the medium with a very long dispersion time; then $R \gg c\sigma/n$ is no longer valid. In this case, the information carried by the time of flight method about the original position of the events becomes less reliable. However, it is still possible to reconstruct the original position of the event by taking into account that the density of hits on the PMTs decreases with the inverse of the squared distance from the point where the energy is deposited [141].

Suppose that we have no timing information, so our only information about an event is the pattern of hit PMTs. In this case, the likelihood function simply determines the position of the event. It does not depend on time and cannot be used to reconstruct the time itself. We may set the function $p(\tau_e)$ to be constant and ignore it:

$$\log \mathcal{L}(a\hat{\mathbf{u}}) = \text{const} + \sum_{i=1}^{N} \log \frac{R - a \cos \theta_i}{s_i^3}.$$
(5.27)

By the same methods as above, we obtain

$$\log \mathcal{L}(a\hat{\mathbf{u}}) \approx \text{const} - \frac{2N}{3R^2}a^2 \tag{5.28}$$

for the second-order Taylor expansion in a/R of the likelihood function for an event at the detector center. In this case we find

$$\frac{\partial}{\partial a} \log \mathcal{L} = 0 \iff a = 0, \tag{5.29}$$

Detector	R	T	n	σ	ϵ	N	Pred.	Obs.
	[m]			[ns]	[pe]		$\delta a [{ m cm}]$	
Organic scintillator detectors								
CTF, ²¹⁴ Po α [138, 143]	3.3	100	1.8	5.1	225	90	12.0	12.3
Borexino, $1 \mathrm{MeV} \ e^- \ \mathrm{MC} \ [55]$	6.5	2240	1.5	5.1	400	366	8.8	8.0
Hypothetical ℓ Ne detector, 100 keV e^- MC [142]								
Spatial data only	3.0	1832	-	-	243	243	16.7	17.0
Timing included	"	"	1.2	10	162	155	15.0	13.6

Table 5.1: Comparison of the predicted resolutions of three liquid scintillator detectors with the values determined experimentally or by Monte Carlo (MC) methods. See the text for meanings of the columns and comments on values in *italics*.

and for the resolution,

$$\delta a = \left(-\frac{\partial^2 \log \mathcal{L}}{\partial a^2}\right)^{-1/2} = \sqrt{\frac{3}{N}} \frac{R}{2}.$$
(5.30)

Recall Equations (5.24) and (5.26) in the case where timing information is available:

$$\log \mathcal{L}(a\hat{\mathbf{u}}, t_0) \approx \operatorname{const} - \frac{2N}{3R^2}a^2 - \frac{N}{\sigma^2}\left(t_0^2 + \frac{n^2}{3c^2}a^2 + \frac{2n}{3cR}a^2t_0\right)$$
$$\delta a = \left(\frac{Nn^2}{3c^2\sigma^2} + \frac{4N}{3R^2}\right)^{-1/2} \approx \sqrt{\frac{3}{N}}\frac{c\sigma}{n}.$$

We see that use of timing information improves spatial resolution significantly when the scintillator dispersion time is much less than the travel time for light to cross the detector. In a liquid noble gas detector, the scintillator time dispersion is very broad due to the amount of internal Rayleigh scattering of scintillation light. Nevertheless, use of even the small amount of timing information available has been shown to improve the spatial resolution by a large fraction [142].

5.3.4 Comparison to observed resolutions

Experimentally, the position resolution of a detector can be determined in several ways. The simplest and most common is the use of a calibration source. In cases when the detector

has not yet been built, Monte Carlo methods are of course the only method that can be used. The detector resolutions obtained from experimental results for the Counting Test Facility (the Borexino 4-ton prototype), and from Monte Carlo simulations of Borexino and a hypothetical liquid neon dark matter detector [142], are shown in the last column of Table 5.1. For comparison, the physical attributes of the detectors and the predicted resolutions δa from Equation (5.26) are shown in the other columns of the table. As above, R is the detector radius, T the total number of PMTs, n the scintillator index of refraction, and σ the scintillator dispersion time. The average number of photoelectrons detected in each event from the source used is denoted by ϵ .

N is determined in most cases as follows. In detectors using a time-of-flight position reconstruction method, each PMT can measure the arrival time only of the first photon it detects. This difficulty will be discussed more thoroughly in Section 5.4. The immediate consequence is that N is a measure of the number of hit PMTs rather than the total number of detected photoelectrons. Basic probability tells us that given an event in which ϵ photoelectrons are detected, the expected number of hit PMTs is

$$\langle N \rangle = T \left[1 - \left(\frac{T-1}{T} \right)^{\epsilon} \right].$$
 (5.31)

Note, however, that for the spatial hit pattern, every photoelectron contributes to our knowledge, even for multiple hits on a single PMT. This implies that the term $4N/3R^2$ in the expression for δa in Equation (5.26) should in fact include ϵ , not N. In calculating the predicted values of δa in Table 5.1, we therefore use the modified expression

$$\delta a = \left(\frac{Nn^2}{3c^2\sigma^2} + \frac{4\epsilon}{3R^2}\right)^{-1/2}.$$
(5.32)

Some comments on idiosyncracies of the individual detectors are in order. The data used for the Counting Test Facility (CTF) are from the set of measurements taken by its first incarnation, known as CTF 1 (see Section 6.4.1). The value n = 1.8 tabulated for CTF 1 is an "effective index of refraction." The software historically used to analyze events in the CTF detector does not take into account the possibility of more than one photon detected per PMT; as will be explained in Section 7.1.1, the value of n used in the reconstruction must be increased in order to balance this effect. Additionally, note that the observed value of δa for the CTF takes into account only the spread in x and y coordinates; the CTF source had the shape of a cylinder, extended in z.

In the hypothetical liquid neon detector described in reference [142], events have a prompt component (relative intensity 2.0) and a delayed component (relative intensity 1.0) of scintillation light. For the Monte Carlo simulation taking into account only the spatial pattern of PMT hits ("spatial data only" row of Table 5.1), both components contribute useful data. In that case the photoelectron yield is 2428 pe/MeV, 1.5 times the prompt light yield of 1619 pe/MeV (10791.7 photons/MeV × 20% quantum efficiency × 75% geometric coverage) quoted in the reference. For the position reconstruction calculated from the spatial pattern only, we use $N = \epsilon_{total} \equiv \epsilon_{prompt} + \epsilon_{delayed}$ in Equation (5.30).

Calculation of the expected resolution in the liquid Ne detector is trickier when timing information is included ("timing included" row of table 5.1). The two terms contributing to δa in Equation (5.32) must be evaluated with different values for ϵ . The term $4\epsilon/3R^2$ comes from the spatial hit pattern and so uses $\epsilon_{total} = 243$, while the timing-dependent term $Nn^2/3c^2\sigma^2$ includes only the prompt component of scintillation light, and thus uses $\epsilon_{prompt} = 162$, with N = 155 derived from Equation (5.31).

The source of the largest potential errors in the predictions of Table 5.1 is the value of the scintillator dispersion, σ . The true scintillator dispersion function of a detector $p(\tau_e)$ is not actually a Gaussian, so the use of Equation (5.26) is only an approximation. The value of 5.1 ns used for σ in the CTF is obtained from the fit to a CTF 1 scintillator response function described in reference [80]. The parameters of this fit, shown in Figure 7.3, are given in Section 7.5; the function was sampled at 1 ns intervals and fit to a pure Gaussian in order to obtain a value for σ . (The same scintillator dispersion function was used in the Borexino Monte Carlo simulations.) Nevertheless, the predicted, observed and Monte Carlo values of the position resolution are in quite good agreement. For the liquid Ne detector,

 σ was estimated at 10 ns, based on Figure 7 of reference [142], as $\frac{1}{2}$ the difference between times with probability values equal to $e^{-1/2}$ times the value at the peak. One could plausibly estimate this value of σ to be anywhere in the range 5.5 to 15 ns, yielding estimates of δa from 12.6 to 15.9 cm. This range brackets the Monte Carlo simulation nicely.

5.4 Multiple PMT occupancy

So far it has largely been assumed that the occupancy of each PMT in the detector is at most one. If the detector has the capability to measure the time at which *every* detected photon hits a given PMT, or if the detector (as with some of the proposed noble gas detectors) has no timing capability at all, then the assumption may be lifted with no effect, except that some of the \mathbf{x}_i (and hence θ_i and s_i) will be identical in Equation (5.12). For a detector with timing capabilities, however, it is more likely that the detector only has the capability to measure the arrival time of the *first* photon detected by each PMT. The probability function of the first detected photon to reach a PMT is not the same as that of a random photon detected by the same PMT; it is biased toward earlier times. To account for this bias, the scintillator response function $p(\tau_e)$ must be corrected.

5.4.1 Correcting for timing bias

Let the probability function of the first photon detected by a PMT, out of the *n* photons seen by that PMT from an event, be represented by $p_n(\tau_e)$. (Naturally, $p_1(\tau_e) \equiv p(\tau_e)$ by definition.) This is known as the "first order statistic." In general, the corrected scintillator response function p_{corr} would then be some linear combination of the first order statistics,

$$p_{corr}(\tau_e) = \sum_{n=1}^{\infty} p_n(\tau_e) \times \mathcal{P}(n \text{ photons hit the PMT}),$$
(5.33)

and an *a priori* guess would have to be made for the probability that each possible number of photons was seen by the PMT. For simplicity, let us assume that the number of photons



Figure 5.3: A hypothetical Gaussian scintillator response functions $p(\tau_e)$ and its first order statistics for increasing values of n = 2, 3, 5, 10. Note how as n increases, the corrected response function narrows and shifts toward earlier times. The time axis is shown in units of the scintillator dispersion time σ .

detected by each PMT for an event is known (in Borexino, for instance, this is determined via ADC channels separate from the timing channels). We can then set p_{corr} equal to the function $p_n(\tau_e)$.

It remains only to calculate $p_n(\tau_e)$ given $p(\tau_e)$ and n. Label the emission times of the n photons detected by a given PMT in some random order unrelated to timing (for instance, in order of increasing photon energy or increasing z component of momentum) as τ_1, \ldots, τ_n . Also number them in order of increasing emission time as s_1, \ldots, s_n . Then $p_n(\tau_e)$ is the probability function of the randomly chosen emission time τ_1 given that $s_1 = \tau_1$:

$$p_{n}(\tau_{e}) d\tau_{e} = P(\tau_{1} \in [\tau, \tau + d\tau] | \tau_{1} = s_{1})$$

$$= P(\tau_{1} = s_{1} | \tau_{1} \in [\tau, \tau + d\tau]) \times \frac{P(\tau_{1} \in [\tau, \tau + d\tau])}{P(\tau_{1} = s_{1})}$$

$$= \frac{p(\tau_{e}) d\tau_{e}}{(1/n)} P(\tau_{1} = s_{1} | \tau_{1} \in [\tau, \tau + d\tau]), \qquad (5.34)$$

where the second equality is once again due to Bayes' Theorem. The probability in the last line above is just the probability that every other detected photon has a later arrival time than the randomly selected value τ_1 :

$$P(\tau_{1} = s_{1} | \tau_{1} \in [\tau, \tau + d\tau]) = \prod_{i=2}^{n} P(\tau_{i} > \tau_{1} | \tau_{1} \in [\tau, \tau + d\tau])$$

$$= P(\tau_{2} > \tau_{1} | \tau_{1} \in [\tau, \tau + d\tau])^{n-1}$$

$$= \left[\int_{\tau_{e}}^{\infty} p(\tau_{e}') d\tau_{e}' \right]^{n-1}.$$
 (5.35)

Hence (letting $F(\tau_e) \equiv \int_{-\infty}^{\tau_e} p(\tau'_e) d\tau'_e$ represent the cumulative distribution function of τ_e), the first order statistic of $p(\tau_e)$, if *n* photons are detected by a given PMT, is

$$p_n(\tau_e) = np(\tau_e) \left[1 - F(\tau_e)\right]^{n-1}.$$
 (5.36)

(This equation is derived independently in, for example, reference [144].)

Graphs of the first order statistics of a representative scintillator response function are shown in Figure 5.3 for values of n equal to 1, 2, 3, 5, and 10. (The specific response function shown is a Gaussian, Equation (5.14) offset by five units of σ from time zero.) Note how as n increases, the time distribution of the first PMT hit narrows and shifts toward earlier times.

5.4.2 Effects on detector resolution

One may ask about the effect of this correction on the likelihood function and spatial resolution. Consider again the case of a Gaussian scintillator time response function. We have

$$\log p_n(\tau_e) = \text{const} + \log p(\tau_e) + (n-1)\log[1 - F(\tau_e)].$$
(5.37)

Substituting in $F(\tau_e) = \frac{1}{2} [1 + \operatorname{erf}(\tau_e/\sigma\sqrt{2})]$, the Taylor expansion to second order in τ_e becomes

$$\log p_n(\tau_e) = \text{const} - (n-1)\sqrt{\frac{2}{\pi}}\frac{\tau_e}{\sigma} - \left(\frac{1}{2} + \frac{n-1}{\pi}\right)\frac{\tau_e^2}{\sigma^2} + O(\tau_e^3).$$
(5.38)

That is, the first photon detected by each PMT contributes to the log of the likelihood function in the amount of $-\tau_e^2/2\sigma^2$, but each additional photon seen contributes only in the amount of $-\tau_e^2/\pi\sigma^2$ (plus a term linear in τ_e which has relatively little effect on the resolution for a large detector); compare to Equation (5.14). The resolution is better than if the corrected scintillator response function were not used, but still poorer than if the time of arrival of every detected photon were known.

Suppose that the total number of photons detected is ϵ , by N PMTs, and in particular that the i^{th} PMT sees n_i photons. Denoting the emission time by $\tau_e^i \equiv t_i - t_0 - s_i n/c$, the general likelihood function is then

$$\log \mathcal{L}(a\hat{\mathbf{u}}, t_0) = \operatorname{const} - \frac{1}{\sigma^2} \sum_{i=1}^{N} \left(\frac{1}{2} + \frac{n_i - 1}{\pi} \right) \left(\tau_e^i \right)^2 - \frac{1}{\sigma} \sqrt{\frac{2}{\pi}} \sum_{i=1}^{N} (n_i - 1) \tau_e^i + \sum_{j=1}^{\epsilon} \log \frac{R - a \cos \theta_j}{s_j^3}.$$
(5.39)

Define the excess photoelectron multiplicity as $\delta \equiv (\epsilon - N)/N$. The likelihood function in the limit of homogeneous PMT coverage as $N \to \infty$, for an event at the detector center, becomes

$$\log \mathcal{L}(a\hat{\mathbf{u}}, t_0) = \operatorname{const} - \frac{N}{\sigma^2} \left(\frac{1}{2} + \frac{\delta}{\pi}\right) \left\langle \left(\tau_e^i\right)^2 \right\rangle \\ - \frac{N\delta}{\sigma} \sqrt{\frac{2}{\pi}} \left\langle \tau_e^i \right\rangle + N(\delta + 1) \left\langle \log \frac{R - a\cos\theta_j}{s_j^3} \right\rangle.$$

Running through calculations analogous to those of Section 5.3.1, we finally obtain the explicit function

$$\log \mathcal{L}(a\hat{\mathbf{u}}, t_0) = \operatorname{const} - \frac{N}{\sigma^2} \left(\frac{1}{2} + \frac{\delta}{\pi} \right) \left(t_0^2 + \frac{n^2}{3c^2} a^2 + \frac{2n}{3cR} a^2 t_0 \right) - \frac{N\delta}{\sigma} \sqrt{\frac{2}{\pi}} \left(t_0 + \frac{n}{3cR} a^2 \right) - N(\delta + 1) \frac{2}{3R^2} a^2.$$
(5.40)

In the limit $c\sigma/R \to 0$ (that is, for a very large detector compared to the width of the scintillator response function), it can be shown that the spatial resolution at the center of a detector is given by

$$\delta a = \frac{c\sigma}{n} \sqrt{\frac{3}{\epsilon} \frac{\pi(1+\delta)}{\pi+2\delta}}.$$
(5.41)

Compare with the approximation $\frac{c\sigma}{n}\sqrt{\frac{3}{N}}$ for Equation (5.26). Hence, holding ϵ fixed, the resolution of an event with an average photoelectron multiplicity of $\delta = 0.5$ excess photoelectrons per PMT is a factor of $\sqrt{\pi(1+\delta)/(\pi+2\delta)}$, or 6.7%, worse than if PMTs could measure the arrival time of every photon seen. With $\delta = 1$ excess photoelectron per PMT (every hit PMT seeing an average of 2 photons), the resolution is 10.5% worse. In the limit of large δ (for instance with a high-energy event), the resolution reaches an asymptote of $\sqrt{\pi/2}$ times (about 25.3% worse) that of an ideal detector observing an event of equal energy.

Realistically, construction of an ideal detector, one that measures the time of arrival for every detected photon, would be non-trivial. One may on the other hand ask, given a detector capable of measuring time of arrival only for the first photon seen by each PMT, how the use of the statistically corrected scintillator dispersion function improves the results over the use of an uncorrected function. This comparison is equivalent to fixing N while (for the uncorrected dispersion function) setting δ to zero. In this case, the use of the corrected dispersion function is an improvement by the factor $\sqrt{\pi}/\sqrt{\pi + 2\delta}$ (recall that smaller resolutions are better). For $\delta = 0.5$, the reciprocal of the improvement factor is 1.15, and for $\delta = 1$, it is 1.28; for large δ , it would theoretically improve without bound. This comparison even leaves aside the fact that for events offset from the center of the detector, use of the uncorrected scintillator dispersion function will produce a statistically biased position estimate; refer to Section 7.1 for an example.

5.5 Expected spatial distribution of reconstructed events

In this section, we leave off thinking about the magnitude of the likely error in position for a single reconstructed event, and instead consider how this error affects the spatial distribution of many events observed in a detector. We first clarify the distinction between the true and observed position distributions. If the spatial resolution of the detector were zero (infinitely good), the two would be identical. Only one distribution would need to be considered, the function $\phi(\mathbf{x})$, where the probability for an event to occur within the volume element $d^3\mathbf{x}$ of the detector would be given by $\phi(\mathbf{x}) d^3\mathbf{x}$. In the real world, however, the reconstructed position of an event will always differ somewhat from the true position. The true and observed positions of an event will henceforth be labeled \mathbf{x}_r and \mathbf{x}_d , respectively ("r" for "real" and "d" for "detected"), and the spatial distributions of the true and observed positions of events will be labeled $\phi_r(\mathbf{x})$ and $\phi_d(\mathbf{x})$, respectively.

5.5.1 Calculating the expected observed event distribution

We can define a "resolution function" of the detector $R(\mathbf{x}_r, \mathbf{x}_d) d^3 \mathbf{x}_d$ to be the probability that, if the true event is at position \mathbf{x}_r , the reconstructed position of the event will be within the volume element $d^3 \mathbf{x}_d$ centered about \mathbf{x}_d . More formally,

$$R(\mathbf{x}_0, \mathbf{x}_d) \,\mathrm{d}^3 \mathbf{x}_d \equiv \mathrm{P}(\mathbf{x}_d \in \mathrm{d}^3 \mathbf{x}_d \,|\, \mathbf{x}_r = \mathbf{x}_0).$$

We expect that $R(\mathbf{x}_r, \mathbf{x}_d)$ can be approximated by a Gaussian in \mathbf{x}_d centered at \mathbf{x}_r :

$$R(\mathbf{x}_r, \mathbf{x}_d) = (2\pi\sigma^2)^{-3/2} e^{-\frac{(\mathbf{x}_d - \mathbf{x}_r)^2}{2\sigma^2}}.$$
(5.42)

The constant factor $(2\pi\sigma^2)^{-3/2}$ is needed because the integral over all space $\int R(\mathbf{x}_r, \mathbf{x}_d) d^3 \mathbf{x}_d$ must be 1. The value σ is the detector's spatial resolution, which has until now been labeled δa . It is in general a function of the true event position and number of photoelectrons, $\sigma = \sigma(\mathbf{x}_r, \epsilon)$. For now we assume mono-energetic events and do not explicitly write out the energy dependence of σ .

It should be noted that this Gaussian form of the resolution function is only a simple model. It assumes (necessarily, but not sufficiently) that (1) the reconstruction is not biased, *i. e.*, has no systematic sources of error; and (2) the reconstruction error is isotropic: $\sigma_x(\mathbf{x}_r) = \sigma_y(\mathbf{x}_r) = \sigma_z(\mathbf{x}_r) = \sigma(\mathbf{x}_r)$ at each point \mathbf{x}_r in the detector. We emphasize that the model is *not* a logical consequence of the likelihood function analyses in previous sections of this chapter.

We now ask how to obtain $\phi_d(\mathbf{x})$ given $\phi_r(\mathbf{x})$. The probability that the reconstructed event will be observed within the volume element $d^3\mathbf{x}$ is given by

$$\phi_d(\mathbf{x}) d^3 \mathbf{x} = \int P(\mathbf{x} \in d^3 \mathbf{x} | \mathbf{x}' \in d^3 \mathbf{x}') P(\mathbf{x}' \in d^3 \mathbf{x}')$$
$$= \int R(\mathbf{x}', \mathbf{x}) d^3 \mathbf{x} \phi_r(\mathbf{x}') d^3 \mathbf{x}'$$

yielding, with the insertion of equation (5.42),

$$\phi_d(\mathbf{x}) = \int \phi_r(\mathbf{x}') \frac{e^{-\frac{(\mathbf{x}-\mathbf{x}')^2}{2\sigma^2}}}{(2\pi\sigma^2)^{3/2}} d^3 \mathbf{x}'.$$
(5.43)

The integral in \mathbf{x}' is of course taken only over the detector volume of active scintillator. (Or, equivalently, $\phi_r(\mathbf{x}')$ is defined to be zero when \mathbf{x}' is outside that volume.)

5.5.2 Radially symmetric event distributions

In order to obtain some understanding of this relationship, we for now make the simplifying assumptions that

- the detector scintillator volume and set of PMTs are spherical and centered on the origin of coordinates;
- the detector resolution is radially symmetric, and therefore is a function only of an event's distance from the origin;
- the physical distribution of events in the scintillator is also radially symmetric.

We now consider how to obtain the distribution of the reconstructed positions of events for three cases: events produced by impurities distributed homogeneously throughout the scintillator volume ("internal events"); alpha and beta particles emitted from the surface of a vessel containing the scintillator ("surface events"); and gamma rays coming into the scintillator from that surface or from outside ("external events").

Use of radial symmetry

When the true event distribution ϕ_r is radially symmetric, the observed event distribution ϕ_d must also be radially symmetric. Without loss of generality, we may therefore work in the coordinate system in which $\mathbf{x} = |\mathbf{x}|\hat{\mathbf{z}}$, so the angle between the vectors \mathbf{x} and \mathbf{x}' is θ , the co-latitude of \mathbf{x}' . Defining $r \equiv |\mathbf{x}|$, we can expand the argument of the exponential in equation (5.43) as

$$(\mathbf{x} - \mathbf{x}')^2 = r^2 + r'^2 - 2\mathbf{x} \cdot \mathbf{x}' = r^2 + r'^2 - 2rr' \cos \theta.$$

Converting to spherical coordinates, equation (5.43) becomes

$$\begin{split} \phi_d(r) &= \frac{1}{(2\pi)^{3/2}} \int d\Omega' \, dr' \, r'^2 \, \frac{\phi_r}{\sigma^3} \, e^{-\frac{r^2 + r'^2}{2\sigma^2}} \, e^{-\frac{rr'\cos\theta'}{\sigma^2}} \\ &= \frac{2\pi}{(2\pi)^{3/2}} \int_0^R dr' \, r'^2 \, \frac{\phi_r}{\sigma^3} \, e^{-\frac{r^2 + r'^2}{2\sigma^2}} \int_{-1}^1 d(\cos\theta') \, e^{-\frac{rr'\cos\theta'}{\sigma^2}} \\ &= \frac{1}{\sqrt{2\pi}} \int_0^R dr' \, r'^2 \, \frac{\phi_r}{\sigma^3} \, e^{-\frac{r^2 + r'^2}{2\sigma^2}} \frac{2\sigma^2}{rr'} \sinh \frac{rr'}{\sigma^2} \\ &= \sqrt{\frac{2}{\pi}} \frac{1}{r} \int_0^R dr' \, \frac{\phi_r}{\sigma} \, r' \, e^{-\frac{r^2 + r'^2}{2\sigma^2}} \sinh \frac{rr'}{\sigma^2} \end{split}$$
(5.44)

with R defined here as the radius of the volume of active scintillator, not the distance of the PMTs from the center as in previous sections.

If ϕ_r and σ are even functions of r' (they can always be made so, since they are not physically defined for r' < 0 anyway), this may be simplified slightly by expanding the sinh function and taking $r' \to -r'$ in the second term:

$$\phi_d(r) = \frac{1}{\sqrt{2\pi} r} \int_0^R dr' \frac{\phi_r}{\sigma} r' \left[e^{-\frac{(r-r')^2}{2\sigma^2}} - e^{-\frac{(r+r')^2}{2\sigma^2}} \right]$$
$$= \frac{1}{\sqrt{2\pi} r} \int_{-R}^R dr' \frac{\phi_r}{\sigma} r' e^{-\frac{(r-r')^2}{2\sigma^2}}.$$
(5.45)

Note the change to the bounds of integration in the final step. We always require that r itself is non-negative.

Internal event distribution function

In the case of a uniform radioactive contamination throughout the active scintillator volume V, the true distribution of internal events $\phi_r(r)$ is constant and equal to $1/V = 3/(4\pi R^3)$. Therefore, substituting into equation (5.45),

$$\phi_d(r) = \frac{3}{4\pi R^3} \frac{1}{\sqrt{2\pi} r} \int_{-R}^{R} dr' \frac{r'}{\sigma} e^{-\frac{(r-r')^2}{2\sigma^2}}.$$
(5.46)

The relevant quantity to which experimental data should be fit is not $\phi_d(r)$, but instead is the distribution function in r of reconstructed events,

$$f_d(r) = r^2 \int d\Omega \phi_d(r) = 4\pi r^2 \phi_d(r)$$
 (5.47)

$$= \frac{3}{R^3} \frac{r}{\sqrt{2\pi}} \int_{-R}^{R} \mathrm{d}r' \frac{r'}{\sigma} e^{-\frac{(r-r')^2}{2\sigma^2}}.$$
 (5.48)

(This function, by construction, is the derivative of the function that gives the fraction of events having a reconstructed radial position less than r.)

The detector resolution σ may be determined experimentally in a number of different ways, or estimated as in Section 5.3.4. The functions $\phi_d(r)$ and $f_d(r)$ are plotted in Figure 5.4 for a volume of scintillator with a radius of 1 meter, assuming a constant value of σ for the three cases $\sigma = 10$, 15, and 20 cm. These are comparable to the values for Borexino's Counting Test Facility.

Event distribution functions for surface α particles and electrons

Since the average scintillator penetration length of α particles in an organic scintillator such as pseudocumene is < 1 mm, and that of 1 MeV electrons is ~ 1 cm, the travel distances of these particles originating from the scintillator containment vessel are negligible compared to σ . Then ϕ_r may be approximated by the delta function

$$\phi_r(r) = \frac{\delta(r - R + \epsilon)}{4\pi R^2} \tag{5.49}$$

(where we let $\epsilon \to 0$ after any integration).

Substituting this into equation (5.44), the reconstructed event density and radial position become

$$\phi_d(r) = \sqrt{\frac{2}{\pi}} \frac{e^{-\frac{R^2}{2\sigma^2}}}{4\pi R} \frac{e^{-\frac{r^2}{2\sigma^2}}}{\sigma r} \sinh \frac{rR}{\sigma^2}$$
(5.50)



Figure 5.4: Internal distribution functions $\phi_d(r)$, top, and observed event radial distribution functions $f_d(r)$, bottom, for constant $\sigma = 10$, 15 and 20 cm.

$$f_d(r) = \sqrt{\frac{2}{\pi}} \frac{r}{R} \frac{e^{-\frac{R^2 + r^2}{2\sigma^2}}}{\sigma} \sinh \frac{rR}{\sigma^2} = \frac{1}{\sqrt{2\pi}\sigma} \frac{r}{R} \left[e^{-\frac{(r-R)^2}{2\sigma^2}} - e^{-\frac{(r+R)^2}{2\sigma^2}} \right]$$
(5.51)

where σ is the value of $\sigma(r)$ evaluated at r = R. These functions are graphed for various values of $\sigma(R)$ in Figure 5.5.

Event distribution functions for external γ rays

Gamma rays, unlike α and β particles, can travel tens of centimeters in a liquid scintillator before being completely absorbed. A Monte Carlo simulation would be necessary to evaluate their true event position distribution accurately, but a crude estimation can be made analytically by assuming that all the energy of each γ ray is absorbed at once (no scattering). The results of a Monte Carlo simulation will be examined in Section 9.4.

Let the average scintillator penetration length of a γ ray of specific energy be λ , and assume $\lambda \ll R$, the radius of the scintillator volume. If s is the total distance traveled by an emitted particle, the distribution of s is of course given by

$$g(s) ds = (1/\lambda)e^{-s/\lambda} ds.$$
(5.52)

We first calculate the true distribution $\phi_r(r)$ of events produced by absorption of γ rays originating on the scintillator containment vessel. Consider the geometry shown in Figure 5.6, in which a γ ray produced at point A penetrates a distance s into the volume of scintillator before being absorbed at point B. The angle ψ is that between the γ ray path and a normal vector into the scintillator. The Law of Cosines gives us $r^2 = R^2 + s^2 - 2Rs \cos \psi$. The question is then, what is the distribution function $f_r(r)$ for the distance from the center at which a γ ray originating on the vessel is absorbed?

We know that $\cos \psi$ is a uniformly distributed random variable on the interval (0, 1] (recall that a value of $\cos \psi \leq 0$ would imply an event outside the scintillator, which cannot be detected). Hence, for a fixed value of s, r^2 is a uniform random variable on $[(R-s)^2, R^2 + s^2)$.



Figure 5.5: Surface distribution functions $\phi_d(r)$, top, and observed surface event radial distribution functions $f_d(r)$, bottom, for $\sigma = 10$, 15 and 20 cm.



Figure 5.6: Geometry of surface event distribution calculation.

The change of variables theorem for a random variable tells us that therefore r (for a fixed value s) has the distribution

$$P(r \in dr | s \in ds) = \begin{cases} r dr/2Rs, & |R-s| \le r < \sqrt{R^2 + s^2} \\ 0 & \text{otherwise.} \end{cases}$$
(5.53)

(The factor of $\frac{1}{2}$ comes about because we reject the non-physical case r < 0.)

If it is assumed that $0 \le r < R$, then the above function is non-zero only when $R - r \le s \le R + r$. Therefore the distribution of r without the constraint of fixed s becomes

$$P(r \in dr) \equiv f_r(r) dr = \int_{s=R-r}^{R+r} P(r \in dr | s \in ds) P(s \in ds)$$
$$= \int_{s=R-r}^{R+r} \frac{r dr}{2Rs} \frac{e^{-s/\lambda}}{\lambda} ds$$
$$= \frac{r dr}{2R\lambda} \left[ei \left(-\frac{R+r}{\lambda} \right) - ei \left(\frac{r-R}{\lambda} \right) \right]$$
(5.54)

where ei is the exponential integral function $ei(x) \equiv -\int_{-x}^{\infty} (e^{-t}/t) dt$. Therefore, in the case of a uniform distribution of γ rays produced at the scintillator containment vessel, we obtain

$$\phi_r(r) = \frac{f_r(r)}{4\pi r^2} = \frac{1}{8\pi R r \lambda} \Theta(R-r) \left[\operatorname{ei}\left(-\frac{R+r}{\lambda}\right) - \operatorname{ei}\left(\frac{r-R}{\lambda}\right) \right]$$
(5.55)

 $(\Theta(x))$ being the Heaviside function: one for x > 0, zero otherwise).

As an aside, note that the previous two equations are not properly normalized. Their respective integrals over the volume of the vessel and over the range $0 \le r < R$ do not evaluate to 1, but instead to the value

$$\int_{V} \phi_{r}(r) \,\mathrm{d}V = \int_{0}^{R} f_{r}(r) \,\mathrm{d}r = \frac{1}{2} \left[1 - \frac{\lambda}{2R} \left(1 - e^{-2R/\lambda} \right) \right].$$
(5.56)

This should not come as a surprise. The quantity above represents the fraction of γ rays emitted from the containment vessel surface that are actually absorbed within the scintillator. The factor of $\frac{1}{2}$ in front represents the fact that a γ ray produced at the surface has only a 50% chance of being directed inward instead of outward. As the ratio λ/R increases, the chance for a random γ ray to be absorbed within the scintillator becomes ever smaller.

The true event distributions $\phi_r(r)$ and the observed radial distribution functions $f_d(r)$ for varying values of λ are plotted in Figure 5.7. The latter is of course obtained from the former by substituting into Equation (5.44) and multiplying by $4\pi r^2$:

$$f_d(r) = \frac{1}{\sqrt{2\pi}} \frac{r}{R\lambda} \int_0^R \mathrm{d}r' \left[\mathrm{ei}\left(-\frac{R+r'}{\lambda}\right) - \mathrm{ei}\left(\frac{r'-R}{\lambda}\right) \right] \frac{1}{\sigma} e^{-\frac{r^2+r'^2}{2\sigma^2}} \sinh\frac{rr'}{\sigma^2}.$$
 (5.57)

In fact, most γ rays will originate outside the scintillator containment vessel, for instance from the PMTs. The event distribution at a fixed radius r within the scintillator for a fixed energy will then be the integral of Equation (5.55) over all external sources:

$$\phi_r(r) = \frac{1}{8\pi r\lambda} \int_R^\infty \rho(r') \frac{\mathrm{d}r'}{r'} \Theta(r'-r) \left[\operatorname{ei}\left(-\frac{r+r'}{\lambda}\right) - \operatorname{ei}\left(\frac{r-r'}{\lambda}\right) \right]$$
(5.58)

where $\rho(r')$ is the normalized density of γ ray emissions at the given energy within the spherical shell of radius r'. This formula assumes spherical symmetry. As before, this function may be substituted into equation (5.44) and multiplied by $4\pi r^2$ to obtain the observed event distribution in r, $f_d(r)$.

It must be noted again that this derivation is not completely realistic. A better model of external γ ray behavior will be considered in Section 9.4.1.



Figure 5.7: True distribution functions $\phi_r(r)$, top, and observed radial distribution functions $f_d(r)$, bottom, for γ rays with mean scattering distances of $\lambda = 5$, 15 and 25 cm, and (in the second plot) constant detector resolution $\sigma = 10$ cm. The vertical asymptotes at r = 1 m in the upper graph result from the non-physical assumption that the vessel film thickness is infinitesimal yet emits γ rays at a finite rate. In these graphs, the factor of $\frac{1}{2}$ resulting from Equation (5.53) has been suppressed.

5.5.3 Previously used radial distribution functions

Above, we have followed the steps of assuming a true spatial distribution of events $\phi_r(r)$, deriving from this function the observed spatial distribution $\phi_d(r)$, and multiplying by $4\pi r^2$ to obtain a reconstructed radial distribution function:

$$\phi_r(r) \xrightarrow{\text{convolve with}} \phi_d(r) \xrightarrow{\times 4\pi r^2} f_d(r).$$
(5.59)

The radial distribution functions $f_d^*(r)$ used in previous CTF analysis work resulted from the assumption that the true radial distribution of events $f_r(r) \equiv 4\pi r^2 \phi_r(r)$ can be convolved with a one-dimensional Gaussian, $e^{-\frac{(r-r')^2}{2\sigma^2}}/\sqrt{2\pi}\sigma$, to produce the radial distribution of reconstructed events. That is,

$$\phi_r(r) \xrightarrow{\times 4\pi r^2} f_r(r) \xrightarrow{\text{convolve with}} f_d^*(r).$$
(5.60)

For a particular event distribution $\phi_r(r)$, the function $f_d^*(r)$ is not necessarily equal to the function $f_d(r)$ defined in equation (5.47).

Given a constant internal distribution function $\phi_r = 1/V = 3/(4\pi R^3)$, the true radial event distribution is

$$f_r(r) = 4\pi r^2 \phi_r = \frac{3r^2}{R^3}.$$
(5.61)

The convolution (assuming constant σ) yields

$$f_d^*(r) = \frac{k_b}{\sqrt{2\pi}\sigma} \int_0^R f_r(r') e^{-\frac{(r-r')^2}{2\sigma^2}} dr'$$
(5.62)

$$= \frac{3}{R^3} \frac{k_b}{\sqrt{2\pi}\sigma} \int_0^R \mathrm{d}r' \, r'^2 \, e^{-\frac{(r-r')^2}{2\sigma^2}},\tag{5.63}$$

with k_b a normalization constant chosen such that $\int_0^\infty f_d^*(r) dr = 1$. Specifically,

$$k_b = \left[\frac{1}{2} \left(1 + \operatorname{erf} \frac{R}{\sqrt{2}\sigma} \right) + \frac{1}{\sqrt{2\pi}} \frac{\sigma}{R} e^{-\frac{R^2}{2\sigma^2}} - \sqrt{\frac{2}{\pi}} \left(\frac{\sigma}{R} \right)^3 \left(1 - e^{-\frac{R^2}{2\sigma^2}} \right) \right]^{-1}.$$
 (5.64)

In the limit $\sigma/R \ll 1$, we find $k_b \to 1$. (With R = 1 m and $\sigma = 15$ cm, approximating k_b as 1 is already accurate to within 0.3 percent. In the CTF 3, σ is more on the order of 10 to 12 cm.)

We now consider the surface event distribution. Using the spatial distribution of equation (5.49), we obtain $f_r(r) = r^2 \delta(r - R + \epsilon)/R^2$, and the convolution with a Gaussian in r results in

$$f_d^*(r) = \frac{k_s}{\sqrt{2\pi}\sigma R^2} \int_0^R dr' r'^2 \,\delta(r' - R + \epsilon) \, e^{-\frac{(r - r')^2}{2\sigma^2}} = \frac{k_s}{\sqrt{2\pi}\sigma} \, e^{-\frac{(r - R)^2}{2\sigma^2}} \tag{5.65}$$

(taking $\epsilon \to 0$ in the last step) with the normalization requiring

$$k_s = 2\left(1 + \operatorname{erf} \frac{R}{\sqrt{2}\sigma}\right)^{-1}.$$
(5.66)

In the limit $\sigma/R \ll 1$, we have $k_s \to 1$ with an error of < 0.1 ppb at $\sigma = 16$ cm.

Figure 5.8 shows a comparison of the reconstructed radial distribution functions developed in equations (5.48) and (5.51) with the functions previously used in CTF analysis [equations (5.63) and (5.65)], using the limiting values of $k_b = k_s = 1$. Notable differences are that

- the previously used functions do not quite go to zero at r = 0;
- the previously used functions have peaks at smaller values of r than the corresponding functions developed above.

It is not simple to determine how the use of these functions instead of the "correct" ones discussed earlier in this section would affect analyses of internal and surface events in the CTF. The differences between the two sets of functions presumably permit the "correct" ones to fit the observed data slightly better. Still, the two sets of functions are not really that different. In addition, other effects such as the radial dependence of σ and the presence of γ events from external sources would tend to obscure the differences, so their effects upon real experimental analyses should be small.



Figure 5.8: Comparison of the above-developed internal (top) and surface (bottom) radial distribution functions, $f_d(r)$, with those previously used in CTF analysis, $f_d^*(r)$, assuming $\sigma = 12$ cm and setting $k_b = k_s = 1$. The previously used functions are those with peaks at smaller values of r.

5.5.4 Cylindrical geometries

Until now this chapter has considered event distributions in a spherically symmetric detector. Consider, though, the problem of a spherical detector which has two "hot spots" on opposite sides; call these the north and south poles. Suppose we want to make a fit of the observed radial distribution of events to a weighted sum of the internal, surface, and external distribution functions $f_d(r)$ defined above. A naïve treatment using these functions, ignoring the θ and ϕ coordinates of events, will yield overly pessimistic conclusions about the average number of surface events per unit area, and the internal and external event distributions may be swamped by the hot spots. Instead, the hot spots should somehow be excluded from the data. This can be done in two obvious ways:

- Cut out wedges (cones) around the poles. Doing so will exclude equal ratios of scintillator containment vessel surface and scintillator fluid volume from the analysis.
- Consider only a relatively thin slice through the equator of the scintillator containment vessel, and analyze it assuming a cylindrical geometry.

The latter solution requires us to develop a new set of reconstructed event distributions for a cylindrical geometry. These may also be useful in other circumstances, since most particle detectors at beam colliders are cylindrical as well.

Without loss of generality, aside from the assumptions pointed out in Section 5.5.1, we may start again with Equation (5.43):

$$\phi_d(\mathbf{x}) = \int \phi_r(\mathbf{x}') \frac{e^{-\frac{(\mathbf{x}-\mathbf{x}')^2}{2\sigma^2}}}{(2\pi\sigma^2)^{3/2}} d^3\mathbf{x}'.$$

If we assume that the cylinder is infinitely long, or at least of length $L \gg \sigma$; and that the true event distributions are not z-dependent (the z-axis being the axis of the cylinder); then the problem is essentially two-dimensional due to its translational invariance. The above



Figure 5.9: Internal and surface radial event distributions, $f_d(r)$, for a cylindrical geometry detector. They are shown for constant values of σ at 10, 15, and 20 cm.

equation may be converted into cylindrical coordinates and rewritten as

$$\phi_d(r,\theta) = \frac{1}{(2\pi)^{3/2}} \int d\theta' \, dz' \, dr' \, r' \, \frac{\phi_r}{\sigma^3} \, e^{-\frac{(\mathbf{x}-\mathbf{x}')^2}{2\sigma^2}}.$$
(5.67)

(Note that r is now the coordinate perpendicular to the z axis, not the spherical radial coordinate.)

With true event distributions that are also rotationally invariant around the z axis (no θ dependence), we may perform the integration in a primed coordinate system specified such that $\theta = 0$ and z = 0. That is,

$$(\mathbf{x} - \mathbf{x}')^2 = r^2 + r'^2 + z'^2 - 2\mathbf{x} \cdot \mathbf{x}' = r^2 + r'^2 + z'^2 - 2rr'\cos\theta'.$$
 (5.68)

Conveniently, the integral in z' simply evaluates to $\sqrt{2\pi\sigma^2}$, leaving

$$\phi_d(r) = \frac{1}{2\pi} \int_0^R dr' r' \frac{\phi_r}{\sigma^2} e^{-\frac{r^2 + r'^2}{2\sigma^2}} \int_0^{2\pi} d\theta' e^{\frac{rr'}{\sigma^2} \cos\theta'} = \int_0^R dr' r' \frac{\phi_r}{\sigma^2} I_0\left(\frac{rr'}{\sigma^2}\right) e^{-\frac{r^2 + r'^2}{2\sigma^2}},$$
(5.69)

where $I_0(x)$ is the zeroth-order modified Bessel function of the first kind.

Using $dV = r dr dz d\theta \rightarrow 2\pi r dr \Delta z$, and the respective true internal and surface distributions $1/(\pi R^2 \Delta z)$ and $\delta(r - R + \epsilon)/(2\pi R \Delta z)$, we immediately have the reconstructed distributions of internal and surface events,

$$f_d(r) \text{ [internal]} = \frac{2r}{R^2} \int_0^R dr' \frac{r'}{\sigma^2} I_0\left(\frac{rr'}{\sigma^2}\right) e^{-\frac{r^2 + r'^2}{2\sigma^2}}$$
(5.70)

$$f_d(r) \text{ [surface]} = \frac{r}{\sigma^2} I_0 \left(\frac{rR}{\sigma^2}\right) e^{-\frac{r^2 + R^2}{2\sigma^2}}.$$
(5.71)

The ϕ_d functions look very similar to those for spherical geometries. The radial distribution functions $f_d(r)$ are shown in Figure 5.9. Notice how the internal reconstructed distribution function obeys $f_d(r) \sim r$ for small r, as opposed to the spherical internal distribution function that goes like r^2 . Although largely irrelevant to a thin slice through a spherical detector, but perhaps useful in (e. g.) a wire drift chamber, it is not hard to show that the assumption of complete isotropy in σ is overkill in the geometry of an infinitely long cylinder. It is sufficient to have $\sigma_x = \sigma_y$, with a different value for σ_z , without affecting the results above; one would start from this modified form of Equation (5.42):

$$R(\mathbf{x}_{r}, \mathbf{x}_{d}) = \frac{1}{2\pi\sigma_{\perp}^{2}} e^{-\frac{(\mathbf{x}_{d} - \mathbf{x}_{r})_{\perp}^{2}}{2\sigma_{\perp}^{2}}} \frac{1}{\sqrt{2\pi\sigma_{z}^{2}}} e^{-\frac{(z_{d} - z_{r})^{2}}{2\sigma_{z}^{2}}}.$$
(5.72)

Chapter 6

The Borexino Counting Test Facility

In planning the design of the Borexino experiment, it was always clear that a prototype, or miniature version of Borexino, must first be constructed and tested for the experiment to have any chance of success. This prototype, the "Counting Test Facility" (CTF), would be used for a number of purposes: to test design, purification, and data acquisition techniques; to facilitate research and development of materials to be used in Borexino; to measure the radioactive backgrounds caused by these materials in the scintillator; to test the optical properties of the scintillator on a large scale in three dimensions; and to ensure that the scintillator used in Borexino would be sufficiently low in ¹⁴C and U- and Th-chain isotopes. (The sensitivity of mass spectroscopy is on the order of 10^{-12} g/g, whereas the concentration of U and Th required for Borexino is around 10^{-16} g/g.) The CTF was also able, although it was not originally planned, to perform useful science in its own right, setting new upper limits on various exotic subatomic processes [145, 146, 147, 148]. Most importantly, however, the main goal of the CTF was to demonstrate the feasibility of a few-ton scale, ultra-low background scintillator detector.

The history of the CTF consists of three main phases, which will be described in more detail later in this chapter. The CTF originally took data from April 1995 – July 1997.

This campaign first demonstrated the potential for a Borexino-like low-energy solar neutrino detector [80, 138, 149]. The CTF hardware was given a major upgrade, and in its second phase, the detector (now CTF 2) acquired data with a different scintillator during summer 2000. Essentially the same hardware was used in the CTF 3, a third phase of the detector (sometimes also called the CTF 2b in older references), which acted as a test-bed for various methods of purifying pseudocumene. An unintentional spill of about 50ℓ of pseudocumene in August 2002 resulted in the partial shutdown of the Gran Sasso labs for many months, which as a side effect permitted the observation of ²¹⁰Po decay in undisturbed scintillator. The final operations currently foreseen for the CTF are a test of the scintillator distillation process to be used for Borexino, and last-minute checks of the radiopurity of pseudocumene just before it is inserted into the Borexino detector.

6.1 CTF physical characteristics

Like Borexino, the CTF design is based on the principle of graded shielding; see Figure 6.1. The active scintillation liquid in CTF is a 4 ton mass of pseudocumene and fluor enclosed in a transparent nylon sphere, the CTF vessel. Outside that is a volume of ultra-pure water which is enclosed in a second nylon sphere, the CTF shroud. A set of inward-facing PMTs is arrayed outside the shroud. The entire apparatus, surrounded by another volume of water, is contained in a cylindrical tank. The bottom surface of the tank holds 16 upward-facing PMTs used to tag the tracks of muons passing through the detector.

Many improvements were made to the CTF design in the upgrade from the original version of the detector. The design described in this section applies to the second and third versions of the detector (CTF 2 and CTF 3) except where otherwise noted.



CTF-2 BASIC DIMENSIONS all dimensions referenced to floor of blue tank all dimensions in cm

Figure 6.1: Side view of the design of the CTF 2 and CTF 3. The vessel (labeled R100 in this drawing) and shroud (R200) are shown, as well as the six rings of PMTs, the cylindrical tank, and the tubes used for filling and draining the vessel. The point PNT0 is the nominal center of the sphere of PMTs and of the CTF vessel. Water-filled volumes are shaded blue, while scintillator-filled volumes are shaded yellow. Dimensions are given in cm.



Figure 6.2: A photograph of the CTF 3 viewed from below. At this point in its history, the CTF vessel, shroud, and water tank were all filled with water. The vessel, assembled from 16 discrete panels of film, is at center. The 16 nylon strings holding it down are visible below it. The CTF shroud and three upper rings of PMTs with attached light concentrators are visible in the background.

6.1.1 Details of the CTF design

The scintillator fluid used in CTF 1 and CTF 3 is the same material that will be used in Borexino—a solution of 2,5-diphenyloxazole (PPO) in pseudocumene at a concentration of $1.5 \text{ g/}\ell$. (For CTF 2, see Section 6.4.2.) The volume of scintillator fluid in CTF 3 is 4.2 m^3 , or a mass of 3.7 tons. The scintillator is contained in a transparent spherical vessel 1 m in radius. The vessel is composed of C38F nylon-6 film (*cf.* Section 4.1.4) with a thickness of 500 μ m. It was assembled from 16 lens-shaped panels, joined at the edges with a resorcinol-based glue similar to that used in constructing the Borexino vessels. Because the pseudocumene-filled CTF vessel is suspended in water ($\Delta \rho = 0.124 \text{ g/cm}^3$), it must withstand a buoyant load of 460 kg force. Hence the CTF vessel is much thicker than the



Figure 6.3: A closer view of the design of the vessel and shroud in the CTF 3. The vessel (filled with pseudocumene-based scintillator) is shown in yellow. The shroud (blue), surrounding it, is shown in cross-section; it is filled with water, as is the volume outside it. Pipes for fluid filling and draining operations leave the vessel at top and bottom. They pass through the two cover plates of the shroud at the poles. The nylon strings holding down the vessel against buoyant force are not shown.
$125 \,\mu\text{m}$ film used in the Borexino vessels, which have pseudocumene solutions with nearly identical densities on both sides.

The CTF 3 vessel is also held down against the buoyant load by 16 lengths of monofilament nylon line that cross over the top of the vessel and are fixed in place underneath. Fishing line was used to minimize mass (and therefore intrinsic radioactivity). The line used was identical to that in CTF 1. In CTF 2, the hold-down system instead comprised a set of Vectran ropes; Vectran was being tested as a candidate material for the Borexino ropes. Unfortunately, Vectran (which was otherwise a prime candidate, exhibiting very little creep) turned out to be unacceptably high in 40 K [150]; see Section 6.4.2. The nylon strings and the discrete panels making up the CTF 3 vessel may be seen clearly in Figure 6.2.

Outside the vessel, as shown in Figure 6.3, is a region of about 30 tons of ultra-pure water, contained by a second nylon vessel of radius 2 m, made from nylon-6 Capran DF400 film produced commercially by Allied Signal. This vessel, the CTF shroud, is functionally equivalent to the Borexino Outer Vessel. It prevents radon emanating from the PMTs, outer tank, and nitrogen buffer gas from diffusing into the CTF vessel. Since it separates two volumes of water, it need have a thickness of only $100 \,\mu$ m. Unlike the Borexino Outer Vessel, the shroud was not designed to maintain a spherical shape, nor is it leak-tight. Still, it provided a valuable demonstration that radon in the vessel could be reduced significantly with an additional layer of nylon film. During the CTF 2 campaign, the shroud collapsed, so in the CTF 3 design, ropes were added to hold the shroud away from the vessel. The bottom of the shroud is attached to the rim of a circular nylon cover plate with a diameter of 103 cm.

The shroud and the PMT support structure surrounding it are completely enclosed within a cylindrical steel tank, radius 5.5 m and height 10 m (total volume 950 m^3). This tank is composed of 8 mm thick carbon steel with an interior lining of black Permatex epoxy resin. The tank is filled with water which has been measured to have a radon activity of $20-30 \text{ mBq/m}^3$; one major source is the Permatex lining, which emanates a considerable

amount of radon, 25 mBq/m^2 [63]. Thin nylon pipes through which the CTF vessel may be filled or drained are attached to the vessel film at two small nylon collars. The pipes connect the vessel top and bottom with the top and bottom of the tank. The north pipe is 3.0" in inner diameter with $\frac{1}{4}$ "-thick walls, and 4.5 m long. It runs from the top of the CTF vessel, through the north end-cap of the shroud, to a stainless steel pipe where purified scintillator can be inserted in a clean room atop the water tank. The much thinner south pipe is $\frac{7}{8}$ " in inner diameter (with $\frac{1}{16}$ " walls) and runs 1.8 m from the bottom of the CTF vessel through a hole in the cover plate beneath the shroud. Below the cover plate, it has three 90° elbows that place its exit route on a horizontal along the bottom of the steel tank.

At the top of the cylindrical tank, the north pipe is attached to a bellows in order to give it some vertical freedom of motion. Since the support ropes stretched by a length that was unknown in advance, keeping the vessel centered on the origin of the CTF coordinate system requires the ability to adjust its position vertically. Once positioned properly, the vessel is maintained in place with a set of pulleys and a tank of water used as a counterbalance.

6.1.2 The CTF photomultiplier tubes

Because most of the set of CTF 1 photomultiplier tubes failed, the upgrade from CTF 1 to CTF 2 included installation of a completely new set of PMTs. These PMTs, the same model which is used in Borexino, are the 8"-diameter ETL 9351 tubes manufactured by Electron Tubes Limited and already described in Section 3.3. The CTF 2 tubes were saved and also used in CTF 3. As in Borexino, one cable per PMT supplies the high voltage for PMT operation, and transmits the PMT signal back to the data acquisition system. The means of sealing these cables against water is described in detail in reference [151].

The CTF contains 100 of these PMTs mounted in an approximately isotropic fashion on six horizontal rings. The rings are located at latitudes $\pm 65^{\circ}$ (Rings ± 3), $\pm 40^{\circ}$ (Rings ± 2), and $\pm 12^{\circ}$ (Rings ± 1). They are arrayed with 8 PMTs on each of the polar rings, 18 on each



Figure 6.4: The layout of the CTF photomultiplier tubes, courtesy of M. Göger-Neff. The upper hemisphere is shown at top, and the lower hemisphere at bottom. Both hemispheres are seen from above. The CTF coordinate system has x increasing toward the right (roughly true north) and y increasing toward the top (roughly true west) of the figures. Smaller green numbered circles indicate the PMTs; the larger blue circles indicate the Group 1 TDC and all Group 2 electronics channels (channels shown as halfway between two PMTs are shared by those PMTs). See Sections 6.2.1 through 6.2.3 and Table 6.1 for details.



Figure 6.5: A light concentrator mounted on one of the CTF PMTs.

of the middle rings, and 24 on each of the equatorial rings, as shown in Figure 6.4. Each PMT photocathode is nominally at a distance of 3.3 m from the center of the CTF vessel. Three visible-light digital cameras are also mounted on the PMT scaffolding; one looks at the vessel from slightly above the equator, while the other two peer up at it from below (Figure 6.2). The cameras may be used to photograph the CTF vessel while the PMT high voltage is turned off.

Each PMT, as in Borexino, is surrounded by a so-called "light concentrator," a curved metal surface roughly parabolic in cross-section. Unlike those in Borexino, the CTF light concentrators are metal-coated plastic, not solid aluminum. One of the CTF light concentrators is shown in Figure 6.5. They are designed according to principles detailed in [95, 96] such that almost all light entering the aperture of the light concentrator facing the vessel will be directed, with at most one reflection, into the PMT photocathode. The length of each light concentrator in the CTF is 57 cm, while the maximum radius (at the aperture away from the PMT) is 25.3 cm. The amount of scintillation light entering each PMT is therefore increased by a geometric factor ranging from 8.5–11, depending upon the exact



Figure 6.6: History of the number of surviving PMTs in the CTF 3. Various periods of purification and other CTF operations are represented by vertical lines. During the period December 2001 – March 2004, the number of working PMTs steadily decreased from 96 to 75. Since then, in an attempt to conserve PMTs, the internal CTF cameras have been used less often, and the rate of PMT failures is now dramatically reduced. Most of the data analysis in this dissertation covers the period April 2003 – June 2005, for a maximum of 81 and a minimum of 74 working PMTs. See also Table 6.10. Figure taken from reference [42].

event position. The PMT optical coverage for an event at the center of the CTF is 21%, compared to the 2.3% coverage that would be obtained without the light concentrators. (The cost to implement this level of optical coverage without light concentrators would have been around $10 \,\mathrm{M} \oplus$ [96].) In addition, the light concentrators reject light originating outside the CTF vessel (as seen from each PMT) with high efficiency.

It should be noted that the number of live CTF PMTs has been steadily decreasing over time (Figure 6.6). At the beginning of CTF 2 data taking in 2000, all 100 PMTs were functional; as of this writing, 74 are operational. There appears to be some correlation between PMT deaths and turning the high voltage on or off, for instance in order to take a photograph of the vessel.

6.1.3 The muon veto system

Although muons are not as much of a problem as for Borexino, the CTF still needs a way to exclude muon tracks from its data set. During the construction of CTF 2, the CTF water tank had 16 photomultiplier tubes installed on its floor, facing upward into the tank. They are arrayed in two concentric rings of eight PMTs each, with radii of 2.4 and 4.8 m. These also are the ETL model 9351 8"-diameter tubes described in Section 3.3. The muon veto PMTs are completely enclosed in a stainless steel housing, filled with mineral oil, with a window at the top. This design, which is similar to that used in the Borexino Outer Muon Veto System, is an alternate method of sealing that doesn't require the PMT to be in contact with epoxy. These PMTs have no light concentrators.

As in Borexino, these PMTs are used to detect the Čerenkov light produced by muons traveling through the water buffer fluid. The trigger threshold is 4 photoelectrons. For a muon which causes the inner detector PMTs to detect an amount of light corresponding to an event in the neutrino energy window (250–800 keV), the probability that the muon veto system fails to trigger is only 0.8% [72].

Because only transparent fluids and nylon film separate the active scintillator volume from the muon veto PMTs (unlike the case of Borexino, where the Stainless Steel Sphere comes between the scintillator and muon veto system), these PMTs are also able to see scintillation light. However, the CTF muon veto system was designed to have only a negligible probability of registering point-like events in the scintillator. The probability $\eta(E)$ of the muon veto system flagging such an event at E = 1.9 MeV is only 1% [146], and in the neutrino window it is < 0.5% [72]. Furthermore, since the CTF muon veto system does nothing more than flag events in software, there is no possibility of data loss.

As with the main set of PMTs, not all muon veto PMTs are still operational. The number of functioning muon veto PMTs as of this writing is 12 out of 16, or 75%—about the same as the current fraction of working PMTs in the main data acquisition system.

6.2 CTF data acquisition

The beginning of an event in CTF occurs when a number of PMTs meeting or exceeding a programmable threshold (typically six, corresponding to an event of about 20 keV) are triggered within a 30 ns time period. This causes a majority logic unit to generate an event trigger. This trigger starts several data acquisition units:

- A clock with about 0.1 ns resolution is started on each of the time-to-digital conversion (TDC) channels. A TDC clock stops when it receives a digital signal on its channel. (The stop signal is delayed by 150 ns to ensure that it arrives after the event trigger.)
- A gate on each analog-to-digital conversion (ADC) channel is opened for a period of 500 ns. The ADC channels measure the integrated charge collected on their respective photomultiplier tubes during that time. (Again, the charge signal is delayed by 150 ns to ensure it arrives after the event trigger.)
- Several other data acquisition subsystems are fed ADC channel values, as described in Section 6.2.4.

The CTF electronics and data acquisition system are older and more primitive than those of Borexino: data on the ADCs and TDCs are saved to disk by a μ VAX reading from a CAMAC data acquisition crate. A full description and schematic of the CTF event trigger logic, of which the above is only a summary, may be found in reference [143]. A more upto-date schematic (though with less supporting information) is present in reference [151].

6.2.1 Group 1 and Group 2 events

In addition to the above signals, when an event trigger is generated, a separate TDC clock with a range of 8 ms is started. If another event occurs within the time scale of this clock, it is passed through a separate, parallel set of TDCs and ADCs, the "Group 2" electronics. This

system counteracts the dead time during which the first set of TDCs and ADCs recovers from the first event. Borexino, in contrast, does not have such a dedicated second data acquisition system; its electronics instead have the capability (after a $\sim 10 \,\mu$ s dead time) to acquire a second event while still reporting on the first event.

Events processed by the first set of TDCs and ADCs are referred to as "Group 1" events. The "Group 2" events processed by the second set are generated mainly by decay products of whatever triggered the first event: for instance, capture of a neutron ($\tau = 214 \,\mu s$ [42]) produced by a muon, or the decay of ²¹⁴Po ($\tau = 237 \,\mu s$) immediately following that of ²¹⁴Bi. A Group 1 event not followed by a Group 2 event is called a "single." Group 2 events, because they are so useful in tagging various types of "coincidence event," are vitally important in studying radioactive and muon-induced backgrounds in CTF and extrapolating them to Borexino.

6.2.2 Event energy: the ADC channels

The data acquisition system of the CTF has 100 ADC channels for Group 1 events, one for each PMT. For financial reasons, there are only 64 ADC channels dedicated to Group 2 events. (Before the upgrade to CTF 2, this was also the case for Group 1 ADC channels.) Thirty-six of the Group 2 ADC channels are therefore shared by two PMTs, while 28 are connected to a single PMT. The shared channels report the total charge collected at both PMTs to which they are connected. The two mappings from PMT to ADC channel are reported in Table 6.1. The 16 muon veto PMTs also each have one ADC channel.

6.2.3 Event timing: the TDC channels

Each of the Group 1 and Group 2 electronics systems contains 64 TDC channels. As with the Group 2 ADC system, some PMTs are forced to share a TDC channel for reasons of cost. In this case, the time at which a photoelectron is recorded for a shared channel is the

PMT	G1	G2	ϕ [rad]	PMT	G1	G2	ϕ [rad]	PMT	G1	G2	ϕ [rad]	PMT	G1	G2	ϕ [rad]
Ring	$3: \theta =$	0.43	6 rad	26	40	34	5.829	Ring -	-1: θ	= 1.7	780 rad	76	15	15	0.070
1	1	1	-0.532	Ring 1	$l: \theta =$	1.36	1 rad	51	87	58	-0.131	77	89	59	0.454
2	2	2	0.253	27	63	46	-0.131	52	65	47	0.131	78	90	59	0.733
3	3	3	1.039	28	41	35	0.131	53	66	47	0.393	79	16	16	1.117
4	4	4	1.823	29	42	35	0.393	54	67	48	0.655	80	91	60	1.501
5	5	5	2.609	30	43	36	0.655	55	68	48	0.916	81	92	60	1.780
6	6	6	3.395	31	44	36	0.916	56	69	49	1.178	82	17	17	2.164
7	7	$\overline{7}$	4.180	32	45	37	1.178	57	70	49	1.440	83	93	61	2.548
8	8	8	4.966	33	46	37	1.440	58	72	50	1.702	84	94	61	2.827
Ring	$2: \theta =$	0.87	3 rad	34	47	38	1.702	59	71	50	1.964	85	18	18	3.211
9	9	9	-0.070	35	48	38	1.964	60	73	51	2.225	86	95	62	3.595
10	29	29	0.314	36	49	39	2.225	61	74	51	2.487	87	96	62	3.875
11	30	29	0.593	37	50	39	2.487	62	75	52	2.749	88	19	19	4.259
12	10	10	0.977	38	51	40	2.749	63	76	52	3.011	89	97	63	4.643
13	31	30	1.361	39	52	40	3.011	64	77	53	3.273	90	98	63	4.922
14	32	30	1.641	40	53	41	3.273	65	78	53	3.534	91	20	20	5.306
15	11	11	2.025	41	54	41	3.534	66	79	54	3.796	92	100	64	5.690
16	33	31	2.409	42	55	42	3.796	67	80	54	4.058	Ring -	<i>-3:</i> θ =	= 2.7	05 rad
17	34	31	2.688	43	56	42	4.058	68	81	55	4.320	93	21	21	-0.532
18	12	12	3.072	44	57	43	4.320	69	82	55	4.582	94	22	22	0.253
19	35	32	3.456	45	58	43	4.582	70	83	56	4.843	95	23	23	1.039
20	36	32	3.735	46	59	44	4.843	71	84	56	5.105	96	24	24	1.824
21	13	13	4.119	47	60	44	5.105	72	85	57	5.367	97	25	25	2.609
22	37	33	4.503	48	61	45	5.367	73	86	57	5.629	98	26	26	3.395
23	38	33	4.782	49	62	45	5.629	74	88	58	5.891	99	27	27	4.180
24	14	14	5.166	50	64	46	5.891	Ring -	$-2: \theta$	= 2.2	269 rad	100	28	28	4.966
25	39	34	5.550					75	99	64	-0.314				

Table 6.1: Mapping between PMTs and DAQ channels in the CTF 2 and 3. Three numbering systems are used for the 100 PMTs (also illustrated in Figure 6.4), for the 100 Group 1 ADC channels (G1), and for the 64 Group 2 ADC channels and the TDC channels in both electronics groups (G2). The co-latitude θ and longitude ϕ of each PMT are given in radians.

first time at which either PMT attached to the channel records it. In order to ensure good position resolution for events near the top and bottom of the detector, the 16 PMTs on the top and bottom rings (Rings ± 3) are each connected to single TDC channels. (In CTF 1, the top and bottom sets of PMTs shared TDC channels in pairs, and the position resolution in z was poor.) The mapping from PMT to TDC channel is reported in Table 6.1. The 16 muon veto PMTs each have one corresponding TDC channel.

6.2.4 Other data channels

The tail to total ratio

The analog sum of the charge signals from all PMTs is split off and sent to four ADCs, with different time offsets: 0 ns, 16 ns, 32 ns, and 48 ns. On each of these ADCs, the charge is integrated for 500 ns. The resulting values are referred to, respectively, as ADC 3000, ADC 3016, ADC 3032, and ADC 3048. Because scintillation light produced by a β particle or γ ray dies off more quickly than that produced by an α , the shapes of their respective pulses as functions of time are different. The ratio between the values ADC 3032 and ADC 3000, called the tail to total ratio, can therefore be used as a means of particle identification. (Conceivably the ratio between any other pair of these values could also be used.) These values are available for both Group 1 and Group 2 events.

The Digital Pulse Shape Analyzer board

An alternative approach to particle identification is provided by the DPSA board [81] installed in the upgrade to CTF 2. This board acquires the sum of PMT charges as a function of time and feeds it to an analog integrator. The output of the integrator is sampled at 120 MHz by an interleaved pair of 10-bit flash ADCs that each run at 60 MHz. Time bins in the output are therefore 8.3 ns wide; data for 150 bins are stored. Because the PMT



Figure 6.7: Outputs of the DPSA board in (part a, top) the laboratory and (part b, bottom) the CTF. One would expect the functional shape to rise monotonically to a plateau, as shown at top, since it is the time integral of PMT charge. In fact this is not the case in the CTF; a capacitive connection between the PMT signal and the front-end board causes a slow decrease during the tail of the event, as shown at bottom, due to reversed polarity of the signal at the DPSA input. In both graphs, time in ns is shown on the horizontal axis and the normalized DPSA output on the vertical. Taken from references [81] and [152].

sum gate width is 500 ns, only the first 60 of these time bins are of interest. Like the main ADC and TDC signals, the DPSA board receives the PMT charge signal with some delay after the event trigger; thus the first 10 or so bins provide a pedestal value which may be subtracted off.

Since the charge signal should always have the same polarity, the expected output of the DPSA board, which is the time integral of the charge signal, is a monotonically increasing function. This was seen in laboratory tests of the board, shown in Figure 6.7a. The output of the DPSA board in the CTF (Figure 6.7b) does not meet this expectation. This is a result of a capacitive connection between the PMT output signal and the front-end board, which causes an undershoot in the signal at the DPSA board input [152]. Once the polarity is reversed, the integral of the signal slowly decreases. This effect is thought not to be harmful, as the descent rate should be different for α and β events.

The stored data from the DPSA board is converted into a single parameter used for α/β discrimination by applying the Gatti optimum filter [153] during reconstruction. This method will be discussed in more detail in Section 8.1. It should be noted that the DPSA board is optimized for events in the neutrino window; events that appear to have an energy greater than about 1.5 MeV will saturate the board. This is not a big handicap since α particles in the CTF appear (due to α quenching) to have energies in the range 400–800 keV.

During recent periods of data taking, apparent fluctuations in the DPSA electronics have made the data difficult to interpret for α/β discrimination of events in CTF 3. In this case the tail to total ratio is used as a fallback method of particle identification.

The Time Transient Recorder

The Time Transient Recorder (TTR) is an ADC that samples the output of the PMT analog sum at 200 MHz, and therefore produces 100 data points during the 500 ns window. A total of 200 values are saved in 8-bit format. This measure of the pulse shape arriving at the PMTs was primarily used for detection of muons in CTF 1, before a dedicated muon veto system was installed.

The flash ADCs

The flash ADCs make up a separate data acquisition subsystem which was installed during the CTF 2 runs. Their output is not stored into the standard raw data files, but is separated out into a different file. The purpose of the flash ADCs is to determine accurately the energies of events of 1 MeV or more, which saturate at least one of the main set of ADC channels. This is important in observing, for instance, supernova neutrinos. The CTF flash ADCs, described in detail in references [154] and [155] (both in French), are prototypes for those installed in Borexino.

6.3 File structure of events

It is necessary to save the data from each event in the CTF into a format easily parsed by computer. For several reasons, two different formats are used to represent event data: a "raw" format, and a "reconstructed" format. The raw format is intended for quick writing by the data acquisition computer programs, with minimal processing, so it consists of binary data organized byte-by-byte into simple data structures. The reconstructed format has a more hierarchical structure that can be browsed easily by users of the ROOT data analysis software [102]. More importantly, the basic data stored in the raw event (charge detected on ADC channels; time of PMT hits from TDC channels) is converted into information about the position of the event, whether it was more like an α or β particle, and so on, before being stored as a reconstructed event. This permits very high-level analysis. The process of converting an event from a raw format to a reconstructed format is of course called "reconstruction." For practical reasons, CTF data acquisition is run for a few hours to days at a time. All the events acquired during this time period, or "run," are saved in raw format into a single file. Each file consisting of raw events also contains some header information describing the calibration of the ADC and TDC channels.

6.3.1 Note on different numerical formats used

Different computers store numerical data in different formats. The CTF raw data is acquired by a VAX, while reconstruction programs are typically run under the GNU/Linux and Microsoft Windows operating systems on Intel-compatible processors, or under the Mac OS X operating system on IBM's PowerPC processor. Therefore it is important to be aware of the formats used for numerical data, and how to convert between them.

32-bit integers may be broken up into four bytes of eight bits each. For simplicity, we consider only unsigned integers. These bytes may be stored in memory in either "little-endian" format, which has the least significant byte (LSB) coming first, or "big-endian" format, with the LSB coming last. (This specification of byte order generalizes to 16-bit and 64-bit integers as well.) The 32-bit integer with hexadecimal representation 6789ABCD₁₆ would be represented in little-endian format as $CD_{16} AB_{16} 89_{16} 67_{16}$. In big-endian format, it would instead be represented in memory as $67_{16} 89_{16} AB_{16} CD_{16}$. Intel-compatible processors and the VAX use little-endian format, while PowerPC processors use big-endian. Big-endian is also the default format used by computers sending data over a network. Converting between the two formats requires reversing the order of all the bytes. Formats with the bytes in an order other than these two are called "middle-endian," and are generally obsolete for use with integers.

Things become more complicated when considering floating-point numbers. A floating-point number such as -5.3×10^{-8} can be decomposed into three parts: a sign (-), an exponent (-8), and a mantissa (5.3). In standard form, the mantissa is in the interval [1,10). Of

course, computers use binary instead of decimal. For 32-bit floating-point numbers, one bit is used to represent the sign of the number $(0 \Rightarrow \text{positive}, 1 \Rightarrow \text{negative})$; eight bits for the exponent; and 23 bits for the mantissa. There are two complications. First, a binary mantissa in standard form will always begin with the digit 1 (since it is between 1 and $10_2 = 2$). The computer representation of the mantissa therefore omits this redundant digit for 24 bits of precision. Second, to avoid a negative value of the exponent, it is stored with some "bias": the computer representation of the exponent is given by the bias plus the real value of the exponent, $E_{\text{stored}} = E_{\text{bias}} + E_{\text{true}}$.

To summarize all this, let a 32-bit floating point number have sign bit s, stored exponent $E = E_{\text{bias}} + E_{\text{true}}$, and mantissa M (where $1 \leq M < 2$). Let the bits of E be $e_1e_2 \dots e_8$ (with e_1 being the most-significant bit and e_8 the least). Let the binary representation of the mantissa be $1.m_1m_2 \dots m_{23}$. Then the number may be represented in big-endian format as

$$(-1)^s \times 2^{E_{\text{true}}} \times M \iff \underbrace{s \, e_1 \dots e_7}_{byte \ 1} \underbrace{e_8 \, m_1 \dots m_7}_{byte \ 2} \underbrace{m_8 \dots m_{15}}_{byte \ 3} \underbrace{m_{16} \dots m_{23}}_{byte \ 4}.$$

In little-endian format, the order of the bytes is reversed. The VAX instead uses a middleendian format: big-endian bytes labeled 1,2,3,4 above are arranged in the order 2,1,4,3. Furthermore, the bias of the exponent on most systems is $E_{\text{bias}} = 127$, conforming to the IEEE 754 standard [156], while on the VAX it is 129. Hence to convert a floatingpoint number from VAX format to a standard little-endian or big-endian format, one must re-order the bytes appropriately, followed by subtracting one from big-endian byte 1 (littleendian byte 4) considered as an unsigned 8-bit integer (in order to decrease the exponent by two, since e_7 is the two's digit of E). Additional complications occur for numbers near the extremes of the range of 32-bit floating point numbers, but these need not be considered for the reasonably sized values stored in CTF data files. The number zero is stored as a bit pattern consisting entirely of zeros in all cases.

6.3.2 Structure of raw data files

As shown in Table 6.2, a CTF 2 or CTF 3 raw data file consists of a 276-byte header, a set of three ADC calibration tables, a set of three TDC calibration tables, an unspecified number of raw events, and finally a run footer. Each of these data structures begins with a 32-bit magic number which can be used as a consistency check. The individual data structures are summarized in tables below. In all of the following tables in the "Type" columns, 16-bit unsigned integers are abbreviated as "u16"; an array of 15 8-bit unsigned integers is abbreviated as "u8[15]"; etc. All integer data are stored in the raw data files in little-endian format. Floating point data are stored in the VAX 32-bit format.

A breakdown of the contents of each data structure is given in Tables 6.3–6.7. Only the raw event data structure will be described in detail; others will be mentioned in passing or are self-explanatory. A raw event in CTF 2 and 3 consists of a number of different fields, listed in Table 6.7. (In CTF 1, a slightly different event structure was used which will not be described here.) Those that are important for reconstruction are as follows.

- Magic number: a 32-bit unsigned integer, used to make sure that the data file is not corrupt at this point. For the event data structure, this number must be 16776960 (in hexadecimal, FFFF00₁₆).
- **TDC raw channel data:** 64 16-bit unsigned integers. Each integer is the raw value acquired by the TDC channel during the event. These values are in units of 0.05 ns. To make the channels comparable, each raw value T must be converted according to the formula $t = (0.05 \text{ ns})(T - T_0)$. Here, t is in nanoseconds and T_0 is an offset specific to each channel, typically about 700 (corresponding to 35 ns). Values of T_0 are read from the TDC calibration table for the group appropriate to the event. (In principle this is also the case for the conversion factor, but it is always equal to 0.05 ns.) A TDC that receives no signal during the 500 ns-long window after the beginning of an event outputs a raw value around 2700 or greater.

Start	Size	Type	Purpose
[bytes]	[bytes]		
0	276	Run header	Data about the run
276	1640	ADC calibration table	Group 1 ADC calibration
1916	1640	ADC calibration table	Group 2 ADC calibration
3556	1640	ADC calibration table	Muon veto ADC calibration
5196	1032	TDC calibration table	Group 1 TDC calibration
6228	1032	TDC calibration table	Group 2 TDC calibration
7260	1032	TDC calibration table	Muon veto TDC calibration
8292	1002	Event	Raw event data
		:	÷
	1002	Event	Raw event data
	28	Run footer	Data about the run elapsed time

Table 6.2: Organization of data within a raw CTF 2 or CTF 3 data file.

Start	Size	Type	Purpose
[bytes]	[bytes]		
0	4	u32	Magic number for consistency check: 16776961
4	8	u64	Run start time (units in 10^{-7} s)
12	4	u32	Run number
16	4	u32	Run trigger
20	128	string	Run comments
148	128	string	Run shifters

Table 6.3: Organization of data within a run header.

Start	Size	Type	Purpose
[bytes]	[bytes]		
0	4	u32	Magic number for consistency check: 16776962
4	8	u64	Run end time (units in 10^{-7} s)
12	8	u64	Duration of run (units in $(1/256) \times 10^{-6}$ s)
20	8	u64	Live time of run (units in $(1/256) \times 10^{-6}$ s)

Table 6.4: Organization of data within a run footer. Starting positions are relative to the beginning of the data structure, not the file.

Start	Size	Type	Purpose
[bytes]	[bytes]		
0	4	u32	Magic number for consistency check: 16776965
4	4	u32	Calibration run number (usually not updated)
8	400	float[100]	Per-channel ADC offsets from zero, Q_0
408	4	float	ADC 3048 offset
412	4	float	ADC 3032 offset
416	4	float	ADC 3016 offset
420	4	float	ADC 3000 offset
424	400	float[100]	Per-channel ADC single photoelectron peaks, $Q_{\rm pe}$
824	400	float[100]	Per-channel ADC offset deviations
1224	4	float	ADC 3048 offset deviation
1228	4	float	ADC 3032 offset deviation
1232	4	float	ADC 3016 offset deviation
1236	4	float	ADC 3000 offset deviation
1240	400	float[100]	Per-channel ADC photoelectron peak widths

Table 6.5: Organization of data within an ADC calibration table. Starting positions are relative to the beginning of the data structure, not the file.

Start	Size	Type	Purpose
[bytes]	[bytes]	-J F -	
0	4	u32	Magic number for consistency check: 16776966
4	4	u32	Calibration run number (usually not updated)
8	256	float[64]	Per-channel TDC relative time zeros, T_0
264	256	float[64]	Per-channel TDC time zero deviations
520	256	float[64]	Per-channel TDC conversions to nanoseconds (always 0.05)
776	256	float[64]	Per-channel TDC offset (unused)

Table 6.6: Organization of data within a TDC calibration table. Starting positions are relative to the beginning of the data structure, not the file.

Start	Size	Type	Purpose
[bytes]	[bytes]		
0	4	u32	Magic number for consistency check: 16776960
4	128	u16[64]	TDC raw channel data
132	200	u16[100]	ADC raw channel data
332	16		(unused)
348	2	u16	ADC 3048
350	4		(unused)
354	2	u16	ADC 3032
356	4		(unused)
360	2	u16	ADC 3016
362	4		(unused)
366	2	u16	ADC 3000
368	2		(unused)
370	2	u16	ADC noise
372	32	u16[16]	Muon veto TDC raw channel data
404	32	u16[16]	Muon veto ADC raw channel data
436	200	u8[200]	TTR buffer
636	300	u16[150]	DPSA buffer
936	4	u32	Muon flag
940	4	u32	Mult. flag
944	4	u32	Muon coincidence
948	12		(unused)
960	4	u32	Coincidence time, ns (for Group 2 events)
964	12		(unused)
976	8	u64	Event time (units in 10^{-7} s)
984	4	u32	Event number
988	4	u32	Event status
992	4	u32	Event words
996	4	u32	Run number (should match that in run header)
1000	2	u16	Event group

Table 6.7: Organization of data within a raw event. Starting positions are relative to the beginning of the data structure, not the file.

- **ADC raw channel data:** 100 16-bit unsigned integers. (As there are only 64 Group 2 ADC channels, only 64 of these integers are filled by Group 2 events.) Each integer is the raw value acquired by the ADC channel during the event. To convert to the number of photoelectrons q seen by that channel, each raw value Q must be converted according to the formula $q = (Q - Q_0)/Q_{pe}$. Q_0 is a channel-specific offset, and Q_{pe} is the channel-specific value for the position of the single photoelectron peak. Typically the peak Q_{pe} is at a raw ADC value of about 50, with a width of about 30, and the offset is in the range 20–70. These parameters are read from the ADC calibration table for the group appropriate to the event. The channels have the capability to measure charges only up to about 20 photoelectrons, so a raw ADC value greater than about 1100 means that the channel was saturated.
- **ADC 3048–3000:** Four 16-bit unsigned integers, separated from each other by 32 unused bits. Each value ADC 30xx is the raw value corresponding to an integral of accumulated charge over all ADC channels from the time xx ns after the beginning of an event to time 500 + xx ns. These values can be used for α/β discrimination, as described later in Section 8.1. For Group 1 events, they are stored as shown in Table 6.7. For Group 2 events, the values are stored immediately after the end of the 64 ADC data channels (so ADC 3048 starts at byte 276 instead of byte 348); this repositioning does not affect positions of any other values in the data structure. The same is true of the corresponding offsets stored in the ADC calibration table (*cf.* Table 6.5).
- Muon veto raw channel data: Analogous to the TDC and ADC channel data for the inward-pointing photomultiplier tubes, but obtained from the 16 PMTs in the CTF muon veto system, pointing upward from the bottom of the water tank (Section 6.1.3). Only the first 16 elements in any of the per-channel calibration tables are filled for muon veto data.
- **DPSA buffer:** 150 16-bit unsigned integer values obtained from the DPSA board; refer to Section 6.2.4. These are used in an alternate approach to α/β discrimination.

- **Muon flag:** A 32-bit unsigned integer. It takes the value one if the muon veto system has been triggered for this event, and zero if not.
- **Coincidence time:** A 32-bit unsigned integer, only relevant to Group 2 events; the time (in nanoseconds) elapsed since the corresponding Group 1 event.
- **Event time:** A 64-bit unsigned integer value. This is the time at which the event occurred, relative to the CTF "time zero" at approximately 20 April 2000. Units are in tenths of microseconds.
- **Event number:** A 32-bit unsigned integer that takes on consecutive values, starting at 1, for each event in the run.
- **Run number:** A 32-bit unsigned integer value that should match the run number occurring in the event header for the same run. This is provided as a consistency check.
- **Event group:** A 16-bit unsigned integer that should only take the values one (for a Group 1 event) or two (for a Group 2 event).

6.3.3 Reconstructed events

Reconstruction is the black art of using raw event data from the data acquisition system to draw conclusions about the true position, time, energy, and nature of each event that caused scintillation light in the detector. Though these characteristics are important, it is not always trivial to determine them from the raw data.

The position and time of an event can in principle be found by applying a likelihood fit to the observed times output by TDC channels, in the manner described in Chapter 5. The CTF detector presents a large number of complications for any spatial reconstruction algorithm, which will be described in the following chapter. Energy of an event may be determined from the total integrated charge over all PMT ADC channels. The calibration, of course, must be worked out in advance, for instance by using the known energy spectrum

Name	Purpose
run	Run number
nevent	Event number
group	Group 1 or Group 2
coinc	Coincidence time (ns) for Group 2 events
Muonflag	Whether this event was associated with a muon
Time	Absolute time of the event, in seconds since the CTF time zero
Timediff	Elapsed time in seconds since the previous event
х	Reconstructed x position of the event (m)
У	Reconstructed y position of the event (m)
Z	Reconstructed z position of the event (m)
\mathbf{t}	Reconstructed time of the event (ns), relative to the first PMT hit
npe	Number of photoelectrons
ene	Estimated energy, keV
$r48_tot$	Ratio of (pedestal-subtracted) ADC3048 to ADC3000
g_ab	Gatti α/β -discrimination parameter

Table 6.8: Some of the more important values stored for reconstructed CTF events. These are all in 32-bit floating-point format. The endianness is irrelevant since the files are intended to be read with ROOT, which determines it automatically. Note that the names are case-sensitive.

of ¹⁴C decays. And although it is seldom that one can determine exactly what caused a particular radioactive decay in the scintillator (the exceptions usually involving Group 2 coincidences), it is often possible to decide whether the emitted particle was an α instead of a β or a γ ray; methods of particle identification will be described in Section 8.1.

A file containing reconstructed CTF events is a ROOT file containing a tree of histograms produced by the reconstruction. An additional object in the tree is called h777. Inside h777 are arrays of stored floating-point values for each event. Some of the more important values are summarized in Table 6.8. These are not set in stone; a new reconstruction program could store reconstructed data in any format desired; however, preserving backwards compatibility is usually a good idea.



Figure 6.8: Geometry of the two-arm design with nylon vessel, one of the three CTF designs proposed to the National Science Foundation in 1992. PMTs were to be located only at top and bottom, with all light being reflected into them by a cylindrical mirror. From reference [158].

6.4 History of the CTF

The original 1991 CTF proposal [157] suggested a volume of scintillator (at the time, still composed mostly of the trimethylborate, $B(OCH_3)_3$, that was to be used in Borexino) enclosed in a cube-shaped acrylic vessel 2 m on a side. The design was "six-armed," with a long support structure extending to each side of the cube, holding PMTs facing into the cube. The region outside the vessel and arms was to be filled with water used as shielding. Even at the time, this CTF design was quite different from the proposed Borexino geometry.

The 1992 CTF proposal submitted by Princeton University to the National Science Foundation [158] contained three suggested designs. Two of these were "two-armed," featuring a 1 m radius vessel supported only on top and bottom by two support structures. The two designs differed mainly in that one incorporated a thick-walled acrylic tank to hold the scintillator, while the other used a thin $(500 \,\mu\text{m})$ nylon vessel. By this time it was realized that PMTs would be a large source of radioactive background, so in these designs, they were moved to the far ends of the arms. An inward-facing cylindrical mirror inside the arms surrounded the vessel and would reflect all scintillation light upwards or downwards into the PMTs; see Figure 6.8.

These designs did not permit spatial reconstruction of event locations, which was soon realized to be an important capability for discriminating between events due to radioactivity on the surface of the vessel, and those due to internal scintillator contamination. As a result, a third design called the CTF Upgrade, which surrounded a nylon vessel with PMTs on all sides, was put forward. To keep the PMTs far from the vessel while maintaining good optical coverage, this design incorporated the light concentrators first developed in [95]. The new proposed design was very similar to the design eventually used in CTF 1, and finally resembled a miniature version of the Borexino detector.

6.4.1 CTF 1: The test of feasibility

Installation of the CTF began in 1993. The original CTF 1 vessel was built with a radius of 1.05 m and therefore had a capacity of 4.8 m^3 . The hardware is described in detail in reference [143]. The cylindrical tank and vessel were filled with water in January 1995, and scintillator (pseudocumene/PPO solution) was placed in the vessel between February and April 1995. The scintillator fluid remained pseudocumene with PPO until August 1996, when it was drained and replaced with water. Unlike later CTF incarnations, the CTF 1 had no muon veto system, no shroud to reduce radon diffusion, no DPSA board for α/β discrimination, and only 64 Group 1 ADC electronics channels. Since all the PMTs on the top and bottom ring shared electronics channels, the spatial resolution in the z direction was poor. CTF 1 was further handicapped by a rapid loss rate of its PMTs due to poor sealing between the base and glass. As water leaked in, short circuits occurred in most PMTs over the two-year operational period.

In October 1996, the CTF vessel was refilled with a different scintillator, the phenyl-oxylylethane (PXE) mixture later used in CTF 2; see the next section. Data taking continued for nine more months, including tests of silica gel column purification and filtration [72]. However, the quality of the PXE data was poor due to the limited number of surviving PMTs (about 30 remained when PXE data-taking began).

Despite its limitations, CTF 1 had a very successful data-taking campaign. The different data-taking periods of CTF 1 are described in reference [128] and designated alphabetically as phases "Alpha" through "Sierra." The periods of most interest are phase "Foxtrot," Runs 82–99, a 35-day period immediately following the introduction of unpurified scintillator; phase "Lima," Runs 220–229, a 6.4-day period following scintillator purification with water extraction and nitrogen stripping; and phase "Oscar," Runs 290–296, a 22-day period following distillation of the scintillator. The main results of the detector are discussed thoroughly in references [80, 128, 138, 149]. They will only briefly be summarized here.

The isotope ¹⁴C is by far the most troublesome radioimpurity in an organic scintillator at energies below 200 keV. The actual endpoint of its β spectrum is 156 keV, but the finite energy resolution in a detector causes ¹⁴C contributions to the radioactive background well above the endpoint. "Pile-up," the occurrence of two discrete events during a time frame short enough that they are detected as a single event, also contributes to higher-energy backgrounds. In CTF 1, the spectral shape of event energies between 60–250 keV was fit to a function consisting of the theoretical ¹⁴C energy spectrum (leaving some uncertain parameters free) convoluted with the Gaussian energy resolution of the detector. The result was a ratio of ¹⁴C to ¹²C atoms of $(1.94 \pm 0.09) \times 10^{-18}$ [149], or $(1.45 \pm 0.07) \times 10^{-18}$ grams ¹⁴C per gram scintillator. Lowering the concentration of PPO in the scintillator showed

that the ¹⁴C contamination in the $2 g/\ell$ PPO solute was not a major contributor to that in the scintillator. It must be emphasized that at the time, this was the most sensitive measurement of ¹⁴C isotopic abundance ever performed, and it was the first measurement ever made of ¹⁴C in a petroleum-derived organic material.

Contamination of the scintillator with the ²³²Th decay chain was measured by taking advantage of the short mean life of the isotope ²¹²Po ($\tau = 431 \text{ ns}$). The β -decay of ²¹²Bi is followed immediately by α -decay of ²¹²Po; these so-called ²¹²BiPo coincidences are easy to tag. During the 35-day phase "Foxtrot" after the scintillator was loaded into the detector, the calculated ²³²Th contamination, assuming secular equilibrium, was $4.4^{+1.5}_{-1.2} \times 10^{-16}$ grams thorium per gram of scintillator [138]. During phase "Lima," a 6.4-day period after water extraction and nitrogen stripping of the scintillator, no candidate ²¹²BiPo coincidences were observed, yielding an upper limit on the ²³²Th contamination of $< 4.1 \times 10^{-16} \text{ g/g}$ at 90% confidence level [138]. From these data it is not possible to draw conclusions about the effectiveness of water extraction as a purification method. A third measurement during phase "Oscar," a 22-day period after the scintillator was distilled, yielded $8.0^{+2.6+0}_{-2.1-0.8} \times 10^{-16} \text{ g/g}$ [138]. The second set of errors in this value are systematics due to a problem with the TTR data acquisition system, used in CTF 1 for muon rejection. Again, not much can be said about the efficiency of distillation as a purification method to remove ²³²Th and daughters.

The radon concentration in the scintillator was measured similarly by observing ²¹⁴BiPo coincidences; τ for ²¹⁴Po is 237 μ s. For the 35 days of phase "Foxtrot" following introduction of the scintillator, the detector was left alone so that the newly introduced radon would decay away. A decaying exponential plus constant term was fit to the internal ²¹⁴BiPo coincidence rate as a function of time. Phase "Foxtrot" was the only period in which the scintillator was left undisturbed long enough to observe a constant component of the coincidence rate; the constant term was found to be 1.5 ± 0.6 counts/day [138]. (The decay time, if left a free parameter of the fit, was consistent with the ²²²Rn mean life of 5.516 days.) Provided that the assumption of secular equilibrium in the ²³⁸U decay chain held, this number would correspond to $(3.5 \pm 1.3) \times 10^{-16}$ g/g of ²³⁸U contamination [138].

However, let us suppose this constant term of the ²¹⁴BiPo event rate is due not to ²³⁸U in the scintillator, but solely to radon diffusion from the external water buffer through the CTF vessel nylon film ($d = 500 \,\mu\text{m}$). From Equation (4.30) and the values in Tables 4.12 and 4.13, this supposition implies a radon concentration in the external water of $25 \pm 10 \,\text{mBq/m}^3$, consistent with the known value. We conclude from this analysis that the ²³⁸U contamination in the scintillator of $(3.5 \pm 1.3) \times 10^{-16} \,\text{g/g}$ given above is only an upper limit.

On the other hand, suppose that the event rate is due only to radon emanation from the nylon film. Using Equation (4.37) and the diffusion length for wet nylon in Table 4.12, we determine an implied 222 Rn production rate in the C38F film of only $6.7 \pm 2.7 \,\mathrm{mBq/m^3}$, significantly less than the value of $18.2 \,\mathrm{mBq/m^3}$ measured for Sniamid film. This low a production rate of radon in C38F is a bit difficult to believe. The fact that there does not seem to be room in the extrapolated 214 BiPo constant rate for a realistic radon emanation from the nylon film (let alone both emanation from film and diffusion from outside) suggests that the assumption of complete mixing in the scintillator is overly pessimistic. The same conclusion may be drawn from Figure 6.14 of reference [12] in which a similar analysis was performed.

Some interesting results were also obtained from studies of the CTF 1 singles event spectrum. A histogram of the event rate as a function of distance from the center of the vessel proved enlightening. This histogram was assumed to be made up of three components: a constant internal distribution with an r^2 dependence convoluted with a Gaussian representing the detector spatial resolution; a Gaussian around the vessel surface to represent impurities in the CTF nylon vessel; and a distribution obtained by flooding the external water buffer with radon, yielding mostly external γ rays. To a first approximation, these correspond to the functions $f_d(r)$ shown in the respective bottom halves of Figures 5.4, 5.5, and 5.7. After the scintillator was distilled (phase "Oscar")¹, fitting the actual event

¹The results for the earlier phase "Lima" immediately following water extraction on the scintillator were roughly consistent with phase "Oscar" [128], but with much poorer statistics due to the shorter data acquisition time, 6.4 days versus 22 days.

distribution for events in the neutrino window to the sum of the three components gave

- 21 ± 47 counts/day (consistent with zero) for internal events;
- 170±60 counts/day for surface events;
- 2050 ± 100 counts/day for external γ events [138].

This result was in one way encouraging, since the internal background before the distillation was $470\pm90 \text{ counts/day [138]}$. On the other hand, it underscored the importance of keeping any possible sources of γ rays as far from the scintillator volume as possible.

The CTF 1 was decommissioned in July 1997. In terms of demonstrating the feasibility of Borexino, which requires scintillator contaminations on the order of a few times 10^{-16} g/g of U and Th or less, it was a great success. Furthermore, many of the observed deficiencies in the CTF 1 design led the way to improvements in future CTF campaigns and in Borexino, of which a dedicated muon veto system and a second nylon barrier (the CTF shroud and Borexino Outer Vessel) to minimize radon diffusion and external γ rays are only two examples.

6.4.2 CTF 2: Tests with a new scintillator

The upgrade to CTF 2, as mentioned already, was a major improvement to much of the detector hardware and data acquisition electronics. A new nylon vessel was installed, as well as a second thinner vessel, the "shroud," surrounding it to act as a radon barrier. The CTF 2 campaign started inauspiciously: in December 1999, the new vessel, after being kept in a dry nitrogen atmosphere for several months, cracked in two places. Both cracks were about 4 cm long, and a third leak developed in the south polar cap [151]. The leak was sealed, and the two cracks were fixed by sandwiching them between nylon plates. As discussed in Chapter 4, the two cracks most probably developed due to brittle failure of nylon film along creases; the vessel thickness of $500 \,\mu$ m was significantly greater than the



Figure 6.9: The absorption and emission spectra of PXE (left), pTP and bis-MSB (right). In part (a) at left, the absorption spectrum of PXE is curve 1 and the emission spectrum is curve 2. In part (b) at right, absorption spectra are shown as solid lines and emission spectra as dashed lines; curves labeled 1 are for pTP and those labeled 2 are for bis-MSB. Taken from reference [73]. Compare with the PC/PPO spectra and the PMT quantum efficiency as functions of wavelength in Figure 3.2.

critical thickness of about $350 \,\mu\text{m}$ found for dry C38F nylon in Section 4.3.5. The lesson was learned, and CTF and Borexino nylon vessels were always kept under moist conditions afterwards.

CTF 2 scintillator

The scintillator used in the last several months of the CTF 1 campaign, and throughout the CTF 2 tests, was primarily phenyl-*o*-xylylethane (PXE), a hydrocarbon with two benzene rings:



The bridge carbon between the rings, having four different groups attached to it, is chiral; therefore PXE has two different enantiomers—forms which are non-superimposable mirror images. The PXE used in CTF, according to the manufacturer (Koch Chemical Company), was produced using a reaction scheme not favoring either enantiomer, so it was presumably a racemic (50%/50%) mixture of the two [159, 160]. But no tests were performed to confirm this [160].

PXE was under consideration as a candidate for the Borexino scintillator fluid partially because its high density of 0.988 g/cm^3 meant that water could conceivably be used as a buffer fluid with little stress on the nylon vessels. (Even so, the density difference of 1% might still have been problematic.) However, the mismatch between its index of refraction n = 1.565 and that of water (1.333) would have caused difficulties for spatial reconstruction algorithms in Borexino similar to those seen in CTF and described in Chapter 7.

A two-fluor system was used to shift the peak wavelength of PXE scintillation light from 290 nm to 430 nm in order to avoid absorption bands due to optical impurities at 360 and 380 nm [73]. At 430 nm, the PMT quantum efficiency is still about 80% of its peak value; *cf.* Figure 3.2. The primary fluor was 1,4-diphenylbenzene, also known as para-terphenyl (pTP), in a concentration of $2.0 \text{ g/}\ell$. The secondary fluor was 1,4-bis-(2-methylstyryl)-benzene (bis-MSB), in a concentration of $20 \text{ mg/}\ell$. Their respective molecular structures are as follows:



The absorption and emission spectra of PXE, pTP, and bis-MSB are shown in Figure 6.9. Physical properties of these compounds are tabulated in Table 3.1. More details about the scintillation properties of the mixture are available in reference [73].



Figure 6.10: The CTF 2 "column." At left is a vertical cross-section through the CTF 2, showing the reconstructed positions of 214 Po decays within 10 cm of the *xz*-plane during Runs 777–789. The spatial reconstruction method used was the standard reconstruction code developed at Milan. The CTF 2 vessel is represented by the circle at center. The color of each point represents the number of 214 Po decays in a 5 cm × 5 cm bin about that point. The image at right shows a similar cross-section through the vessel equator (the *xy*-plane). A vertical column of radon daughter events shows up in both plots.

CTF 2 chronology and results

Following the repair of the CTF 2 vessel, May 2000 (Runs 706–716) was a period of data acquisition during which all parts of the detector were filled with water, mainly to explore the behavior of the new muon veto system. The CTF vessel was filled with PXE scintillator in four batches of one ton each during June and July (Runs 717–757). From mid-July to early September 2000 (Runs 758–789), data were acquired in a stable mode (no operations) of the CTF. Finally, a set of calibration measurements (Runs 790–810) was made using a ²²²Rn point source, in order to check the resolution and accuracy with which event positions could be determined. These calibrations will be described further in Section 7.7.

The isotopic abundance of ¹⁴C in the CTF 2 scintillator was reported as the ratio ¹⁴C/¹²C = $(9.1 \pm 0.3 \pm 0.3) \times 10^{-18}$ [73]. This is equivalent to $(6.9 \pm 0.2 \pm 0.2) \times 10^{-19}$ grams of ¹⁴C

per gram of scintillator. For comparison, this is roughly five times the level of ¹⁴C observed in CTF 1.

During two one-week periods of data taking, a total of two candidate ²¹²BiPo coincidences were observed, corresponding to a ²³²Th concentration in the scintillator of $(4.0\pm3.9\pm0.4) \times 10^{-16}$ g/g. Since this value is, with the errors, consistent with zero, it may be reported as an upper limit on thorium of $< 1.7 \times 10^{-15}$ g/g (90% C.L.) [73]. This value is comparable to that obtained for the CTF 1.

For radon in the scintillator, a measurement of ²¹⁴BiPo coincidences during the last one-week period of data taking was reported as $99 \,\mu \text{Bq/m}^3$ [73], equivalent to 36 counts/day over the whole detector. (The rate during the second-to-last week of data taking was 37 counts/day, so it is assumed that the rate of radon decays had reached a constant value.) However, as seen in Figure 6.10, most of these events occurred in a narrow vertical column centered on the z-axis of the detector. The source of this column is unknown; it is hypothesized to be a result of impurities dripping down from the filling tube on top of the vessel.

When ²¹⁴BiPo coincidences reconstructed to be within 60 cm of the z-axis were excluded, the result was a ²¹⁴BiPo activity of $23 \pm 5 \,\mu$ Bq/m³ [73], equivalent to 8.3 ± 1.8 counts/day over the whole detector. This is almost six times the constant ²¹⁴BiPo background rate seen in CTF 1. It is therefore reasonable to suppose that most of this event rate *is* due to the ²³⁸U decay chain in secular equilibrium in the scintillator, rather than to radon from film emanation or diffusion. The calculated concentration of ²³⁸U in the PXE is then $(1.9 \pm 0.4) \times 10^{-15}$ g/g.

A high rate in the singles events near an energy of 1.4 MeV did not seem correlated to any of the heavy isotope decay chains; this was found through Monte Carlo analysis to be a result of 40 K γ rays produced in the Vectran ropes holding down the CTF vessel. These ropes were later analyzed and found to have a potassium content of 45 ppm [150]. This example underscores the importance of the CTF as a testing facility for materials planned to be used in Borexino.

6.4.3 CTF 3: Purification and background tests

With the PXE exhibiting disappointing results for radiopurity, the Borexino collaboration decided to perform a third campaign with the CTF based again on the original pseudocumene scintillator. For various reasons, it was not possible to fabricate new materials for use in the vessel; leftover film from the original CTF 1 vessel was used for the CTF 3 vessel, and the same monofilament fishing line used in the CTF 1 hold-down ropes was also used in CTF 3. Additionally, some of the light concentrators needed to be replaced; the plastic had come into contact with PXE and been damaged. Other than these changes, and minor modifications to the pipes used for filling the vessel, the hardware and electronics in CTF 3 were almost identical to that used in CTF 2.

The history of the CTF 3 may be divided into several periods. The new vessel was installed in February 2001, and the detector was filled with water during May. Some test runs were performed to check the muon veto performance with Čerenkov light. The detector was filled with scintillator, pseudocumene with added PPO in the usual concentration of $1.5 \text{ g/}\ell$, in late November 2001. Several batches of pure pseudocumene were added at different times in order to determine how the amount of ¹⁴C present varied from sample to sample. (Each batch was added after an equal amount of scintillator was drained, in order to keep the vessel from becoming overinflated.) The PPO concentration has of course been decreased with each addition. Since Run 2180 (June 3, 2002), the most recent addition of pseudocumene, the PPO concentration is $1.03 \text{ g/}\ell$ [42], assuming it was not affected by any of the scintillator purification tests.

A number of purification tests were performed. The scintillator was passed through a silica gel column at two separate times; first, in a continuous loop, then in a batch mode. Water extraction was a second tested method of purification, which was performed both in a continuous loop, and in a "stop-and-go" mode. None of these methods attained the expected improvement factor in radioactive background, and some actually worsened it. Methods and results are detailed in reference [67]. Each period of no activity between these

Date	Runs	Operation	comments			
2001 Feb		CTF 3 vessel installed				
2001 May		filling with water				
2001 Jul 21	2000	begin data taking (water)				
2001 Nov 20–28	2040 - 2046	filling with scintillator				
2002 Jan 10–14	2074 - 2075	add 2nd batch scintillator				
$2002 { m \ Feb} 5-6$	2092 - 2093	blank loop				
$2002 \ \text{Feb} \ 611$	2094 - 2103	silica gel column test	continuous loop mode			
2002 Mar 20–28	2123 - 2130	water extraction test	continuous loop mode			
$2002 {\rm ~Apr} {\rm ~24-26}$	2150 - 2152	blank loop				
2002 Apr 26–28	2153	add 3rd batch scintillator				
2002 May 9–23	2162 - 2173	water extraction test	stop-and-go mode			
$2002 \ \mathrm{Jun} \ 3$	2180	add 4th batch scintillator				
2002 Jun 9–20	2184 - 2187	silica gel column test	batch mode			
2002 Aug 16	2216	-	pseudocumene spill			
2003 May 29	2315	-	Hall C under sequestration			
$2004~{\rm Sep}~1$	2444	-	construction in Hall C			
2005 Mar 17–18	2532 - 2540	calibration with Rn source				
$2005 \ {\rm Jun} \ 22$	2563	-	last data analyzed here			
$2005 {\rm \ Dec} {\rm \ 9}$	2596	-	last data with old scint.			
$2005 \ \mathrm{Dec} \ 15$		refill vessel with water				
2006 Feb 23	2598 +	tests with distilled scint.				

Table 6.9: Timetable of the various CTF 3 operations. Adapted from reference [42].

#	Run	End date	#	Run	End date	#	Run	End date
96	2048	2001 Nov 29	88	2169	2002 May 17	80	2323	2003 Jul 14
95	2064	$2001 \ \mathrm{Dec}\ 22$	87	2219	$2002~{\rm Aug}~26$	79	2336	$2003~{\rm Sep}~03$
94	2092	$2002 \ {\rm Feb} \ 05$	86	2231	$2002~{\rm Sep}~30$	78	2344	$2003 \ {\rm Oct} \ 10$
93	2112	$2002~{\rm Mar}~02$	85	2254	$2002 \ \mathrm{Nov} \ 24$	77	2355	$2003~{\rm Nov}~27$
92	2128	$2002~{\rm Mar}~27$	84	2273	2003Jan 04	76	2391	$2004~{\rm Mar}~08$
91	2141	$2002 { m Apr} 09$	83	2277	2003 Jan 17	75	2552	$2005 {\rm \ May\ } 23$
90	2166	$2002 { m May} 14$	82	2293	$2003~{\rm Mar}~25$			
89	2170	$2002 {\rm \ May\ } 20$	81	2319	$2003 \ {\rm Jun} \ 15$			

Table 6.10: Table of the CTF 3 PMT failures. The table lists the last run during which the number of operational PMTs, as determined by the Princeton reconstruction software, was equal to each given number. (The number of working PMTs was not strictly a monotonically decreasing function, because occasionally a PMT would seem not to work due to a broken electronics channel that could be fixed.) See also Figure 6.6.

tests and the additions of new scintillator should be analyzed for radioactive background independently.

The longest period of data taking in a quiescent mode began in June 2002. Following a minor spill (~ 50 ℓ) of pseudocumene into the environment on August 16, 2002, the detector and the surrounding Hall C were largely shut down for about two years. As an unintended side effect, this forced halt to operations permitted observation of the decay of ²¹⁰Po ($\tau = 200$ days) in the scintillator and on the vessel. Construction in Hall C began in September 2004, in order to make the hall leak-tight against further spills, and may be blamed for an increase in radon background seen since then. No major detector operations happened after the spill until a position calibration of the detector with a radon source was performed in March 2005. After the calibration, normal data taking continued. Results of the calibration will be discussed in Section 7.8, and the radioactive backgrounds observed during normal data taking will be discussed in Chapters 8 and 9. The last data analyzed in this work are in Run 2563 (June 2005), and the most recent available data are in Run 2596 (December 2005).

As of this writing, the CTF is being used for a series of final purification tests for distillation of the pseudocumene, following a procedure very similar to that which will be used in filling Borexino. For that reason, on December 16, 2005, the scintillator was drained from the CTF and replaced with water. Final results of the tests are not currently available. Once the Borexino detector is filled, the future of the CTF is not clear—it may be used as a facility to measure impurities in new pseudocumene to be introduced to Borexino, or it may be reworked into, for instance, a detector for neutrinoless double beta decay (see Section 1.4.3) or for dark matter.

Chapter 7

Writing a Position Reconstruction Code for the CTF

In principle, writing a position reconstruction code for a spherically symmetric detector should not be difficult. Following the methods derived in Chapter 5, one would first determine the probability distribution function (PDF) of the scintillator as a function of time elapsed since the event, p(t). If we use the same coordinate system as in Chapter 5, with the center of the detector defined as the origin, then the likelihood function, from Equation (5.12), is

$$\mathcal{L}(\mathbf{x}_{0}, t_{0}; \{(\mathbf{x}_{i}, t_{i})\}) = \prod_{i=1}^{N} p_{n_{i}} \left(t_{i} - t_{0} - \tau_{f}^{i}(\mathbf{x}_{0}) \right) \left(\frac{R - \mathbf{x}_{i} \cdot \mathbf{x}_{0} / |\mathbf{x}_{i}|}{|\mathbf{x}_{i} - \mathbf{x}_{0}|^{3}} \right)^{n_{i}}.$$
(7.1)

Here, \mathbf{x}_i and t_i are the position of the i^{th} PMT and the time at which it records its first photoelectron; n_i is the photoelectron multiplicity at the i^{th} PMT; p_{n_i} is the PDF corrected for the multiplicity as in Equation (5.36); R is the distance of each PMT from the center of the detector; and \mathbf{x}_0, t_0 are the nominal position and time of the actual event. The function $\tau_f^i(\mathbf{x}_0)$ is the time of flight for a photon traveling from the nominal event position to the i^{th} PMT. We are summing over PMTs, not photoelectrons; since every photoelectron detected by a PMT contributes equally to our knowledge of the hit pattern, the spatial factor is raised to the power n_i . Maximizing \mathcal{L} with respect to (\mathbf{x}_0, t_0) gives the "most likely" actual event
position and time. This equation is a simplification that assumes PMTs of infinitesimal size and no photon scattering in the scintillator.

Because the Counting Test Facility, described in Section 6.1, contains fluids with two different indices of refraction, developing a position reconstruction code for it is not an easy task. The superstructure holding PMTs in the CTF design (see Figure 6.1) is open, as opposed to the case in Borexino. (A proposal to build an analog to Borexino's Stainless Steel Sphere in the CTF was rejected.) Hence, the entire volume of the cylindrical tank outside the CTF vessel—950 m³—must be filled with a single material. This is roughly the same volume as the Borexino Inner and Outer Buffers. If the fluid in question were the same as the buffer fluid to be used in Borexino, it would have to be similarly pure. This was deemed not cost-effective for the CTF, as it is merely a prototype. Instead, ultra-pure water was used as the buffer fluid. Neither pseudocumene (n = 1.504) nor PXE (n = 1.565) has an index of refraction near that of water (n = 1.333), and we must deal with the index of refraction mismatch in the CTF as best we can.

In principle, two modifications have to be made to Equation (7.1). The time of flight $\tau_f^i(\mathbf{x}_0)$, which in Chapter 5 was simply $(n/c)|\mathbf{x}_i - \mathbf{x}_0|$, should be changed to reflect the true travel time between an event in the scintillator and a PMT in the water buffer, taking into account refraction at the interface. The refraction at the CTF vessel also causes the image of a PMT seen from a point inside the vessel to subtend a different solid angle d Ω than it would in a detector with a single index of refraction. Hence the factor $(R - \mathbf{x}_i \cdot \mathbf{x}_0/|\mathbf{x}_i|)/|\mathbf{x}_i - \mathbf{x}_0|^3$ in the likelihood function should likewise be adjusted. However, this spatial factor is less sharply peaked, and therefore provides less information, than the timing component of the likelihood function. Furthermore, its modification for two indices of refraction would be truly difficult. This latter task was not undertaken in the present work. Nor were corrections considered for other factors that make the CTF a non-ideal detector—finite PMT size, existence of scattering, etc. Nevertheless, the position reconstruction method developed in this chapter yields acceptably accurate results.

7.1 Previous approaches to CTF position reconstruction

7.1.1 An effective index of refraction

The most commonly used extant reconstruction software for the CTF [161], a code developed at the University of Milan, attempts to avoid all these complications by subsuming them into one tunable parameter. The code models the CTF as being filled with only one liquid, so this "effective index of refraction" averages over the effects of having two real indices of refraction. This seems reasonable in principle. It will be shown later in this chapter that the CTF may be viewed as a perturbation of a detector with a single index of refraction (that of the scintillator) and a time offset. The offset corresponds to the difference between times of flight from vessel surface radially outward to a PMT in the cases that the region between is filled with scintillator, and that it is filled with the actual buffer material used.

However, independently of the reconstruction problems due to having two indices of refraction, the specific software in question does not deal correctly with the case described in Section 5.4 that more than one photon is detected at a single photomultiplier tube. (This is somewhat surprising, as the correct probability density function to use was described in 1993 in [144]. In that paper it was derived for the purpose of α/β discrimination, though, so perhaps the applicability to position reconstruction was not recognized.) The software therefore has energy-dependent systematic errors as well. Using Equation (5.31) and assuming a detector light yield of about 4 photoelectrons per MeV per PMT, the photoelectron multiplicity for a hit PMT in the CTF for an event in the neutrino window (250–800 keV) ranges between 1.6 and 3.3. In order to correct the statistical bias caused by this effect, the effective index of refraction historically used in CTF analysis has ranged from 1.75 to 1.9.

The need to increase the effective index of refraction to make up for ignoring photoelectron multiplicity effects can be explained using a qualitative argument. Consider Figure 7.1. The true time distribution of photoelectron arrivals, with the time of flight subtracted out, would look like one of the more sharply-peaked functions. By ignoring the correction required to



Figure 7.1: This figure shows the actual scintillator response function (blue) used in CTF 2 and 3, and its first order statistics: the corrections to the function required when more than one photoelectron is detected by each PMT. The first order statistics shown correspond to n = 2, 3, 5, and 10. The time axis is in nanoseconds.

the scintillator PDF when more than one photon is detected at a PMT, an attempt is made to fit this time distribution with the broader lowest-order (blue) function for single photoelectron occupancy. Doing so increases the variance in the calculated times of flight. A greater spread in the time of flight values means one of two things: either the event will be reconstructed farther from the origin than it really is, or else (to balance out the effect) photons must be assumed to be slower than they really are—meaning that a value of ngreater than the true value must be used.

7.1.2 Monte Carlo techniques

An experimental code developed at Munich [72] took the approach of running a Monte Carlo generator for scintillation photons generated at grid points on a 20-cm cubical mesh within the scintillator volume. The simulated photons were tracked in their passage through the scintillator and buffer fluids to the PMTs, taking into account the scintillator response function, Snell's Law at the interface between the two liquids, light scattering, absorption and re-emission, the number of photons detected at each phototube, and other effects. This Monte Carlo simulation yielded a likelihood function into which could be plugged actual PMT timing and occupancy data for an event, giving a likelihood value at each discrete grid point.

The maximum value of the likelihood function for an event was then found by checking the values at each grid point. To determine the values of the likelihood function between grid points, quadratic interpolations were made between nearest neighbors. Although computationally expensive, this approach in principle should work quite well. However, when tested against the CTF 2 source runs, it yielded inaccurate positions for events near the surface of the nylon vessel, as tabulated later in Section 7.7. The reasons for this misbehavior were not investigated in depth. The Munich software was unfortunately not available to be tested against the source runs performed in CTF 3.

7.2 Outline of the CTF reconstruction software

Before getting into details, we will describe the general layout of the position reconstruction software developed at Princeton. The software is written partly in C (mainly for the input reading functions), partly in C++ (mainly for the reconstruction and output functions), and a bit in FORTRAN (one canned routine to calculate cubic spline interpolations, taken from the CERN Library [162]). It is designed for the GNU/Linux operating system, although it will probably compile and run correctly on most UNIX-like architectures. Since the build process checks the machine endianness (see Section 6.3.1) at compile time, the reconstruction software works equally well running on any 32-bit PowerPC or Intel chip. Porting the software to 64-bit chips has not yet been attempted. The reconstruction software has three main tasks:

- Input: It must read in a raw data file of the format described in Section 6.3.2.
- **Reconstruction:** It must use this raw data to estimate the position (and other characteristics) of each event.
- **Output:** It must output the reconstructed data to a file in a format that the ROOT analysis software can read.

7.2.1 Input of raw data

The input portion of the software, before reading in raw event data, first reads the headers and calibration tables at the beginning of a raw data file. It decides whether the data file is a CTF 2 or CTF 3 run¹, based on the run number, and chooses the appropriate index of refraction for the scintillator. It also reads three configuration files giving it various parameters about the CTF. These are searched for in the user's home directory under the subdirectory .ctf; if not found, the software tries to look in /etc/ctf instead. The files in question are named pmtmap2.inp, which describes the mapping of PMTs to electronic channels found in Table 6.1, and scintresponse.inp, which tabulates the scintillator PDF values reported in Table 7.3. A third file, eventaccept.inp, gives the user the possibility to only reconstruct certain events, restricted by energy, muon flag, and/or coincidence time. This is useful when, for instance, only high-energy or coincidence events are of interest; then reconstruction of the large number of ¹⁴C events and muons may be skipped to save time.

The software's next task is to determine which PMTs are dead or untrustworthy. This is accomplished by looping over every Group 1 event in the raw data file and checking which ADC channels were triggered in each event. If the fraction of events in which a PMT was

¹In CTF 1, the detector geometry, the mapping from PMTs to electronic channels, and the format of raw data files all differed slightly from those in CTF 2 and 3. Hence, reconstruction of CTF 1 data is not currently supported by this reconstruction software.

triggered is less than 10%, the PMT is assumed to be dead or dying. If the fraction is greater than 80%, the PMT is assumed to be noisy. In either case, it is flagged to be ignored by the reconstruction algorithm.

At this point, the software is ready to read raw events from the data file. Each event is read in sequence, with a few checks to make sure that all values are meaningful: the stored event number must match the software's current event index; the group number must be either one or two; etc. The ADC and TDC raw data are translated into numbers of photoelectrons and timing in nanoseconds, respectively, as described in Section 6.3.2. The information about each ADC and TDC is merged into a data structure describing the corresponding PMT. Finally, several ancillary pieces of data (for instance, status of the muon veto; coincidence time, where relevant; tail-to-total ratio) are stored together with an array of the PMT data structures. The event information is now passed into the code dedicated to position reconstruction.

7.2.2 Reconstruction code

The reconstruction code calculates the negative of the natural logarithm of the likelihood function $(-\log \mathcal{L})$ as a function of position \mathbf{x}_0 and event time t_0 . This is done numerically. Each PMT whose corresponding ADC and TDC channels have both triggered contributes one term to the function value, consisting of minus the logarithm of one factor of Equation (7.1), with further caveats made later in this chapter. TDCs triggered more than 15 ns after the first TDC to be triggered are ignored; any photons arriving after that time contribute little or no position information. (Since some PMTs share a TDC channel, this algorithm must be modified somewhat as described in Section 7.3.) The goal is to minimize $-\log \mathcal{L}$ over the variables \mathbf{x}_0 and t_0 .

The function $-\log \mathcal{L}(\mathbf{x}_0, t_0)$ is not calculated everywhere inside the CTF vessel; that would require a tremendous amount of computing time. It and its gradient are only explicitly evaluated at points required by the minimization routine. The minimization routine used comes from the MINUIT package [163], a venerable piece of CERN software. Once the routine has calculated the most likely values of \mathbf{x}_0 and t_0 , these and their statistical errors are stored in memory along with the other event information, and passed along to the code for data output.

7.2.3 Output of reconstructed data

Two types of output are produced. The standard output method is to save the reconstructed event positions and other values as n-tuples in a file in ROOT format. In addition, if the reconstruction software is compiled with the flag -DDEBUG, it will print copious amounts of text about the run and about each reconstructed event to the terminal. The information printed includes the raw and calibrated values of each ADC and TDC channel, the approximate event energy, and the reconstructed position of the event as calculated with several different time of flight functions, to be described in Section 7.6. This debugging output is very useful for getting immediate feedback about the operation of the software. It is also in a format sufficiently organized that ROOT scripts may be written to extract data directly from it even while the reconstruction code is still running.

Several auxiliary programs have been developed to complement the main reconstruction software. One program, plot_test, is used to create contour plots of the different time of flight functions under consideration. Example plots may be seen in Figures 7.9, 7.12a, and 7.14a. A second program, named evdisp (EVent DISPlay), will print out the same information generated by the reconstruction software in DEBUG mode, but only for a single user-specified event of a raw data file. Such "random access" event reconstruction is very useful for testing purposes. evdisp also creates contour plots of likelihood values in the CTF xz-plane, seen for example in Figures 7.10, 7.12b, and 7.14b.

7.3 Shared electronics channels

Recall that some of the TDC channels are shared by two PMTs. In addition, for Group 2 events, the same sets of PMTs also share ADC channels. How should the likelihood function, Equation (7.1), be modified to handle this? In these cases, the data from the two PMTs are not independent. The first step is to rewrite the likelihood function as a product of likelihood functions of individual ADC and TDC channels instead of likelihood functions of PMTs:

$$\mathcal{L}(\mathbf{x}_0, t_0) = \prod_{i \in \{\text{hit TDCs}\}} \mathcal{L}_i^{\text{TDC}}(\mathbf{x}_0, t_0) \prod_{j \in \{\text{hit ADCs}\}} \mathcal{L}_j^{\text{ADC}}(\mathbf{x}_0).$$
(7.2)

When the i^{th} ADC or TDC channel serves exactly one PMT, labeled A, the individual likelihood functions have the values taken directly from Equation (7.1),

$$\mathcal{L}_{i}^{\text{TDC}}(\mathbf{x}_{0}, t_{0}) = p_{n_{\text{ADC}(A)}}\left(t_{i} - t_{0} - \tau_{f}^{A}(\mathbf{x}_{0})\right)$$

$$(7.3)$$

$$\mathcal{L}_{i}^{\text{ADC}}(\mathbf{x}_{0}) = \left(\frac{R - \mathbf{x}_{A} \cdot \mathbf{x}_{0} / |\mathbf{x}_{A}|}{|\mathbf{x}_{A} - \mathbf{x}_{0}|^{3}}\right)^{n_{i}}$$
(7.4)

(in the first equation above, ADC(A) is the number of the ADC channel corresponding to PMT A). We now examine the cases where TDC and/or ADC channels are shared by two PMTs.

7.3.1 Shared channels in Group 1 events

Label two PMTs sharing a TDC channel *i* as PMT A and PMT B. Suppose that in an event, A observes n_A photoelectrons, while B sees n_B . For Group 1 events, these values are separately known to us (with some error) from the ADC channels. The ADC likelihood factors in Equation (7.2) are independent of the timing information and do not need to be modified. We also assume that both $n_{A,B} \ge 1$; otherwise it is known which PMT triggered the TDC, and Equation (7.3) applies with no ambiguity.

We now want to determine the probability that the hit time of the first photon detected by either PMT A or PMT B, for an event occurring at the test point (\mathbf{x}_0, t_0) , is the observed value t_i . The time distribution of the scintillator emission time for the first photoelectron detected at PMT A is given by $p_{n_A}(t_i - t_0 - \tau_f^A(\mathbf{x}_0))$, and likewise for PMT B. Since the first detected photon must be at one or the other of the PMTs, the likelihood function for this TDC channel has two terms:

$$\mathcal{L}_{i}^{\text{TDC}}(\mathbf{x}_{0}, t_{0}) = w \, p_{n_{A}}(t_{i} - t_{0} - \tau_{f}^{A}) + (1 - w) \, p_{n_{B}}(t_{i} - t_{0} - \tau_{f}^{B}).$$
(7.5)

The weight w is a function of relevant parameters n_A, n_B, τ_f^A , and $\Delta \tau_f \equiv \tau_f^B - \tau_f^A$. Let t_A be the time (not necessarily known) at which the first photoelectron is detected by PMT A, and likewise for PMT B. Then w is the probability that $t_A < t_B$, given our knowledge that $t_{A,B} \geq t_i$.

We know that t_j (with j = A, B) is a random variable with distribution $p_{n_j}(t_j - t_0 - \tau_f^j)$. Given random variables $T_{A,B}$ with distributions $f_{A,B}(t)$, the probability that $T_A < T_B$ when both random variables are known to be greater than or equal to a specific value t_i is

$$P(T_A < T_B | T_A \ge t_i, T_B \ge t_i) = \frac{\int_{t_i}^{\infty} dt_B \int_{t_i}^{t_B} dt_A f_A(t_A) f_B(t_B)}{[1 - F_A(t_i)] [1 - F_B(t_i)]} \\ = \frac{\int_{t_i}^{\infty} dt F_A(t) f_B(t)}{[1 - F_A(t_i)] [1 - F_B(t_i)]} - \frac{F_A(t_i)}{1 - F_A(t_i)}$$
(7.6)

where $F_{A,B}(t) \equiv \int_{-\infty}^{t} \mathrm{d}t' f_{A,B}(t').$

Recall Equation (5.36):

$$p_n(t) = np(t) [1 - F(t)]^{n-1},$$

with $F(t) \equiv \int_{-\infty}^{t} dt' p(t')$. Notice that the cumulative distribution function for $p_n(t)$ is given by $F_n(t) = 1 - [1 - F(t)]^n$ (this is easily shown by checking that $F'_n(t) \equiv p_n(t)$ and observing that $F_n(\infty) = 1$). Using these facts and substituting in the known distributions of $t_{A,B}$ for $f_{A,B}$ in Equation (7.6), we have

$$w(n_A, n_B, \tau_f^A, \Delta \tau_f) = \frac{n_B \int_{\tau_e^A - \Delta \tau_f}^{\infty} d\tau_e \left\{ 1 - \left[1 - F(\tau_e + \Delta \tau_f) \right]^{n_A} \right\} \left[1 - F(\tau_e) \right]^{n_B - 1} p(\tau_e)}{\left[1 - F(\tau_e^A) \right]^{n_A} \left[1 - F(\tau_e^A - \Delta \tau_f) \right]^{n_B}} - \frac{1 - \left[1 - F(\tau_e^A) \right]^{n_A}}{\left[1 - F(\tau_e^A) \right]^{n_A}},$$
(7.7)

where we have defined $\tau_e^A \equiv t_i - t_0 - \tau_f^A$, and also changed the variable of integration to $\tau_e \equiv t - t_0 - \tau_f^B$.

When $\Delta \tau_f = 0$, this expression can be shown to equal $n_A/(n_A + n_B)$, as we would expect intuitively. We wish to find the magnitude of the change when $\Delta \tau_f$ is moved away from zero. Holding τ_f^A constant, the partial derivative of the expression with respect to $\Delta \tau_f$, evaluated at $\Delta \tau_f = 0$, is

$$\frac{\partial w}{\partial \Delta \tau_f} \bigg|_0 = n_A n_B \left[\frac{\int_{\tau_e^A}^{\infty} \mathrm{d}\tau \left[1 - F(\tau) \right]^{n_A + n_B - 2} \left[p(\tau) \right]^2}{\left[1 - F(\tau_e^A) \right]^{n_A + n_B}} - \frac{1}{n_A + n_B} \frac{p(\tau_e^A)}{1 - F(\tau_e^A)} \right].$$
(7.8)

The integral in this expression cannot be evaluated in closed form. However, since $[p(\tau)]^2 \leq p_{\max}p(\tau)$ (where p_{\max} is the maximum value attained by the scintillator PDF), it has an upper bound given by

$$p_{\max} \int_{\tau_e^A}^{\infty} \mathrm{d}\tau [1 - F(\tau)]^{n_A + n_B - 2} p(\tau) = p_{\max} \frac{[1 - F(\tau_e^A)]^{n_A + n_B - 1}}{n_A + n_B - 1}.$$

Note that the value of the integral is positive. Therefore,

$$\left. \frac{\partial w}{\partial \Delta \tau_f} \right|_0 > -\frac{n_A n_B}{n_A + n_B} \frac{p(\tau_e^A)}{1 - F(\tau_e^A)} \tag{7.9}$$

$$\left. \frac{\partial w}{\partial \Delta \tau_f} \right|_0 < \frac{n_A n_B}{n_A + n_B} \frac{p_{\max}}{1 - F(\tau_e^A)} \left[1 - \frac{p(\tau_e^A)}{p_{\max}} + \frac{1}{n_A + n_B - 1} \right].$$
(7.10)

In all cases, two PMTs that share a TDC channel are adjacent, with the centers of their photocathodes separated by at most 86 cm. The maximum difference in the distances of adjacent PMTs from a point in the CTF vessel is on the order of 30 cm. Hence $|\Delta \tau_f|$ is at

most about 1.5 ns. When the likelihood function is near a maximum, τ_e^A should be near the value at which $p(\tau_e^A)$ is a maximum. For the particular PDF of the CTF scintillator, the value of p(t) at this point is about 0.12 ns^{-1} , and that of F(t) is about 0.25. The upper bound of $\partial w/\partial \Delta \tau_f$ in this situation comes out to roughly 0.08 ns^{-1} when $n_A, n_B \ge 1$. The lower bound comes to about $-0.08(n_A + n_B) \text{ ns}^{-1}$. That is, for an event in which two PMTs sharing a TDC channel each detect two photoelectrons, the maximum expected changes in the weight of $p_{n_A}(\tau_e^A)$ in Equation (7.5) from its mean value of $n_A/(n_A + n_B)$ are -0.24 or +0.12.

7.3.2 Shared channels in Group 2 events

For Group 2 events, some PMTs share both an ADC and TDC channel. It is then unknown which PMT was hit first, as well as the individual numbers of photoelectrons seen at each PMT. In this case, to obtain a factor of the likelihood function for this TDC channel, ideally we would sum Equation (7.5) over all possible values of n_A . The sum $n_A + n_B$ is fixed to the observed ADC channel value n_{AB} , and each term multiplied by the probability that PMT A saw n_A photoelectrons. So we would have

$$\mathcal{L}_{i}^{\text{TDC}}(\mathbf{x}_{0}, t_{0}) = \sum_{n_{A}=0}^{n_{AB}} \binom{n_{AB}}{n_{A}} \left[w(n_{A}, n_{AB} - n_{A}, \tau_{f}^{A}, \Delta \tau_{f}) p_{n_{A}}(t_{i} - t_{0} - \tau_{f}^{A}) + \ldots \right].$$
(7.11)

(The ellipses contain the analogous term for PMT B whose weight is 1 - w.) In reality, an integral should be taken, rather than a discrete sum over possible integer values of n_A , since the observed channel values n_{AB} are continuous due to fluctuations in PMT response.

In practice, we approximate by setting $n_A = n_B = n_{AB}/2$ and using these values directly in Equation (7.5). This is not expected to cause significant errors, as the binomial coefficient function $\binom{n_{AB}}{n_A}$ is sharply peaked about $n_A = n_{AB}/2$. Similarly, rather than obtaining the ADC factor of the likelihood function from the ADC channel by summing over all possible values of n_A given the channel value n_i , we instead use

$$\mathcal{L}_{i}^{\text{ADC}}(\mathbf{x}_{0}) = \left(\frac{R - \mathbf{x}_{A} \cdot \mathbf{x}_{0} / |\mathbf{x}_{A}|}{|\mathbf{x}_{A} - x_{0}|^{3}} \frac{R - \mathbf{x}_{B} \cdot \mathbf{x}_{0} / |\mathbf{x}_{B}|}{|\mathbf{x}_{B} - x_{0}|^{3}}\right)^{n_{i}/2}.$$
(7.12)

7.3.3 Results using the approximate weight function

From the results at the end of Section 7.3.1, one might think it is not safe to neglect the change in relative weights of two PMTs sharing a TDC as their relative distances from a tentative event location vary. This can be tested empirically, using data provided by the CTF 2 source calibration runs. The source runs will be described in more detail in Section 7.7; here it is sufficient to mention that they consist of data taken while a radon source was held inside the CTF vessel at a known position. The most useful events in the radon decay chain are the easily recognizable coincidences $^{214}\text{Bi} \rightarrow ^{214}\text{Po} \rightarrow ^{210}\text{Pb}$. Since ^{214}Po decays by α emission, it has the added advantage of providing monoenergetic events. Position reconstruction of ^{214}Po events from a radon source therefore provides feedback about the performance of reconstruction software at a single point for a single energy. As ^{214}Po events are obtained using the Group 2 electronics system, they also provide a test of the approximations described in Section 7.3.2.

The difference in performance of the reconstruction code using the exact Equation (7.7) and the approximation $w = n_A/n_{AB}$, in both cases with $n_A = n_B = n_{AB}/2$, was tested using CTF 2 source runs 795 and 798. In Run 795, a radon source was placed at the north pole of the CTF vessel. The difference in the average position of ²¹⁴Po events reconstructed with wset equal to the full form of Equation (7.7), versus the average position of the same events reconstructed using the approximation for w, was less than 0.2 cm.² Furthermore, when Equation (7.7) was used, the position reconstruction failed to converge for 1.7% of events, compared with complete success when using the approximation for w. This is probably due to the method of numerical evaluation used for the integral in the equation.

Similar results were obtained with CTF 2 source run 798, in which the radon source was located in the equatorial plane of the CTF vessel, 50 cm from the center. Here the difference in average positions was 0.3 cm, and the fraction of events failing to converge when using the

 $^{^{2}}$ Note that this comparison of positions is between two methods of reconstruction. Comparisons between the reconstructed positions and the actual positions of CTF 2 sources will be discussed later, in Section 7.7.

exact expression for the weight was 0.4%. Additionally, in both runs the position resolution was marginally better when the approximate weight function was used. (In all cases, the time of flight function used for these analyses was the truncated Taylor series described in Section 7.6.) For these reasons, only the approximation $w = n_A/(n_A + n_B)$ is applied in the present CTF reconstruction software. Finally, skipping the calculation of the value of Equation (7.7) also increases the speed of the software by almost a factor of ten.

7.4 Performance of spatial pattern recognition alone

It is interesting to ask how accurately CTF events may be reconstructed when only the spatial pattern of hits is used. That is, we use the likelihood function of Equation (7.2), but all the functions $\mathcal{L}_i^{\text{TDC}}$ are set to be identically one. Only the ADC data are used. In this case, using Equation (5.30), we predict a resolution of $\sigma = 286 \text{ cm}/\sqrt{\epsilon}$, with ϵ being the total number of photoelectrons detected in an event. Again, data from the CTF 2 source calibration runs permit testing this prediction.

Two different sources were used in CTF 2; each consisted of a small quartz vial filled with PXE scintillator enriched in radon. Because it was not possible to remove oxygen from the scintillator without also removing radon, the scintillator was quenched, causing events to produce fewer scintillation photons than would be expected from the event energy. Typically a ²¹⁴Po event occurring in PXE would yield about 300 ± 25 photoelectrons detected in CTF 2. The first source, used in Runs 791–795, yielded roughly 200 ± 16 photoelectrons, while the second source used in the remainder of source runs was more highly quenched $(\bar{\epsilon} = 100 \pm 12$ photoelectrons). We therefore expect position resolutions on the order of 20 cm for reconstructed ²¹⁴Po events from the first source, and 30 cm for the second source. The actual results are reported in Table 7.1.

The table certainly shows a correlation between the reconstructed positions and the nominal position of the sources. With the major exception of Run 795, and a few other lesser

Run	Non	ninal pos	ition	Recons	tructed p	osition	Spatial resolution		
	$x [\mathrm{cm}]$	$y [\mathrm{cm}]$	$z [{ m cm}]$	$x [\mathrm{cm}]$	$y [\mathrm{cm}]$	$z [{ m cm}]$	$\sigma_x [\mathrm{cm}]$	$\sigma_y [\mathrm{cm}]$	$\sigma_z [{\rm cm}]$
791	0	0	32	-0.3	14.8	22.4	27.6	31.2	24.0
793	0	0	0	-3.6	4.2	5.8	28.1	33.7	24.3
795	0	0	100	7.5	26.0	-86.7	34.7	42.4	84.1
797	0	48	0	12.5	94.4	-22.4	39.3	46.3	39.4
798	-48	0	0	-47.4	19.6	-6.2	40.9	53.8	35.8
799	27	-40	0	16.6	-71.4	-10.8	37.1	53.8	37.8
800	40	-27	-30	46.8	-32.2	-40.5	41.6	53.9	38.4
801	-40	27	-30	-37.9	61.7	-43.4	39.5	48.8	34.7
802	0	0	-78	8.7	8.1	-82.5	39.8	46.9	49.3
803	0	0	-83	8.9	-13.6	-78.7	44.3	41.4	50.8
804	0	0	-75	-8.4	-15.7	-60.9	44.7	45.6	51.1
805	0	0	-62	-4.7	-0.9	-58.2	44.9	45.8	45.4
807	-40	27	-56	-63.0	47.0	-49.0	42.3	45.7	36.4
809	-40	-27	-56	-60.5	-13.6	-48.8	40.8	48.0	40.4
810	27	-40	11	24.7	-70.9	6.8	42.0	52.0	37.4

Table 7.1: Performance of the reconstruction software, using only the ADC data (spatial pattern recognition), in CTF 2 source runs. These data come from ²¹⁴Po α decays. Positions are given in the usual CTF-centered coordinate system.

problems, sources are reconstructed to be within 30 cm of their nominal positions. The reconstructed positions of ²¹⁴Bi decays, which are usually in the few-hundred photoelectron range, show similar but slightly more accurate results. The observed position resolutions for ²¹⁴Po events are about 1.5–1.9 times larger than predicted.

Both the poor position resolutions and the problems in Run 795 (shown in Figure 7.2) may be due to light reflection. Scintillation light from an event that is reflected from the inside surface of the CTF tank may end up at any PMT with some non-negligible probability. As a result, part of the PMT hit pattern is invalid data that nevertheless is considered in maximizing the likelihood function. This tends to worsen the resolution. In the special case of Run 795, photons may have been reflected from the nearby north end-cap of the vessel and directed mainly downwards, making events appear to have occurred near the south pole of the vessel instead of the north pole. (No such problems are observed in Runs 802–804 near the south pole, which has a less obtrusive end-cap.) Such reflections from objects near



Figure 7.2: Histogram of the z positions of 214 Po events, reconstructed using only ADC data, in Run 795, which had the source nominally at the north pole (z = +100 cm) of the CTF vessel. The spatial pattern recognition algorithm fails dismally with these data, although for most CTF 2 source runs, the algorithm performed decently considering its limitations.

the event are less of a problem with timing data; the photon times of flight would still be roughly consistent with events at the north pole.

Due to these problems, it was decided to use only timing (TDC) data in the CTF reconstruction software developed at Princeton; that is, to use Equation (7.2) for the likelihood function with all the functions $\mathcal{L}_i^{\text{ADC}}$ set to be identically one. All further results reported in the present work thus use the following likelihood function:

$$\mathcal{L}(\mathbf{x}_0, t_0) = \prod_{i \in \{\text{hit TDCs}\}} \mathcal{L}_i^{\text{TDC}}(\mathbf{x}_0, t_0)$$
(7.13)

where

$$\mathcal{L}_{i}^{\text{TDC}}(\mathbf{x}_{0}, t_{0}) = \begin{cases} p_{n_{i}}(t_{i} - t_{0} - \tau_{f}^{i}(\mathbf{x}_{0}) & \text{[one PMT per TDC]} \\ w \, p_{n_{A}}(t_{i} - t_{0} - \tau_{f}^{A}(\mathbf{x}_{0})) \\ + (1 - w) \, p_{n_{B}}(t_{i} - t_{0} - \tau_{f}^{B}(\mathbf{x}_{0})) & \text{[two PMTs per TDC]} \end{cases}$$
(7.14)

Scintillator	Event	$ au_1$	$ au_2$	$ au_3$	$ au_4$	q_1	q_2	q_3	q_4
	type	[ns]	[ns]	[ns]	[ns]				
PC	α	3.3	13.5	56.0	279.1	0.630	0.178	0.119	0.073
\mathbf{PC}	eta	3.6	17.6	59.5	-	0.895	0.063	0.042	0
PXE	β	3.8	15.9	63.7	243.0	0.832	0.114	0.041	0.013
PXE^*	α	3.1	13.4	56.2	231.6	0.588	0.180	0.157	0.075
PXE*	eta	3.1	12.4	57.1	185.0	0.788	0.117	0.070	0.025

Table 7.2: Time constants and weights of the components of scintillation events caused by α and β particles in the pseudocumene (PC) + PPO ($1.5 \text{ g/}\ell$) and the PXE + pTP ($2.0 \text{ g/}\ell$) + bis-MSB ($20 \text{ mg/}\ell$) mixtures used in the CTF. The mixture denoted PXE* instead had a concentration of pTP of $3.0 \text{ g/}\ell$. These values were measured in the laboratory on a small scale. Taken from references [73, 80].

with $w = n_A/(n_A + n_B)$ for Group 1 events (one PMT per ADC) and $w = \frac{1}{2}$, $n_A = n_B = n_{AB}/2$ for Group 2 events (two PMTs per ADC).

7.5 The scintillator PDF

At the heart of any position reconstruction code in a scintillation-based optical detector is the scintillator response function or PDF, p(t). It is not Gaussian, as was assumed in Chapter 5 for simplicity. The emission of scintillation light, described in Section 3.1, is a multi-stage process; hence its time evolution after an event at t = 0 can be approximated as a sum of decaying exponentials:

$$S(t) = \Theta(t) \sum_{i=1}^{N} \frac{q_i}{\tau_i} e^{-t/\tau_i}.$$
(7.15)

Usually N is only three or four; other components are negligible. The sum of the q_i is arbitrarily set to one so that the time integral $\int_0^\infty S(t) dt$ is normalized. These components for the pseudocumene + PPO and PXE-based solutions that have been used in the CTF have been measured on a small scale in the laboratory; the results are shown in Table 7.2.

Two issues prevent Equation (7.15) from being used as the scintillator PDF in CTF event reconstruction. First, the PMT array has a non-negligible "jitter" that manifests as a



Figure 7.3: Scintillator response function observed in CTF 1. This histogram was created from the observed time distribution of scintillation photons caused by β particles, and fit (solid line) to Equation (7.16). The distribution is shifted slightly toward negative values ($t_0 = -5.22 \text{ ns}$); this is merely a calibration effect which does not affect the position reconstruction. The histogram is not normalized. From reference [80]. Compare with the blue curve in Figure 7.1 obtained from Monte Carlo methods.

Gaussian error in the reported time with $\sigma \approx 1$ ns. Second, and more importantly, photons traveling through a medium will occasionally be scattered, absorbed, and/or re-emitted. This increases their average travel time and also broadens their time distribution. An additional point to consider is that a photon arriving at a PMT a long time after the triggering event (tens or hundreds of ns) may have been reflected one or more times from the inner wall of the steel cylindrical tank or other CTF superstructure instead of truly being emitted in the tail of the scintillator PDF. Such photons carry little or no useful timing information.

t	$10^{5}p(t)$								
[ns]	$[\mathrm{ns}^{-1}]$								
0.00	0	15.25	2999	20.75	3528	25.25	1593	37.25	136
9.25	1	16.25	5189	21.25	3234	26.25	1356	42.25	93
10.25	2	17.25	5716	21.75	2995	27.25	1107	47.25	58
11.25	5	18.25	5323	22.25	2785	28.25	863	52.25	52
12.25	8	19.25	4565	22.75	2462	29.25	716	57.25	43
13.25	108	19.75	4199	23.25	2312	30.25	576	62.25	25
14.25	809	20.25	3861	24.25	1966	32.25	371		

Table 7.3: Values of the scintillator response function used in CTF 3 data analysis, obtained via a Monte Carlo simulation that included the effects of scattering and absorption. The values p(t) shown have been multiplied by 10^5 to make them easier to read. These values must be normalized, making the function's time integral equal to one, before being used as a PDF. This is the same function plotted as the blue curve in Figure 7.1. Entries 2–4 did not come from the simulation; they were included to force the cubic spline interpolation of the data always to be non-negative.

Fitting the observed photon arrival time distribution for ${}^{14}C$ events near the center of the CTF permits these effects to be quantified. The time distribution observed in CTF 1, shown in Figure 7.3, was fit to a convolution of Equation (7.15) with a Gaussian:

$$p(t) = \int_{-\infty}^{\infty} S(t') \, e^{-\frac{(t-t'-t_0)^2}{2\sigma^2}} \, \mathrm{d}t'.$$
(7.16)

In this case only two terms in S(t) with weights q_1, q_2 and time constants τ_1, τ_2 were considered. The broadening effects were all incorporated into a single Gaussian factor with time constant σ and offset t_0 . The values from the fit were $\sigma = 3.17 \text{ ns}$, $t_0 = -5.22 \text{ ns}$, $\tau_1 = 4.75 \text{ ns}$, and $\tau_2 = 26.3 \text{ ns}$ [80]. The weights of the two decaying exponentials are reported in [80] as 296.4 and 26.3, but there is reason to believe that these numbers are the respective values of q_1/τ_1 and q_2/τ_2 . After making this correction, the normalized weights are given by $q_1 = 0.712$ and $q_2 = 0.288$.

The increase in the time constants τ_1, τ_2 from the values in the second row of Table 7.2 are expected results of the absorption and re-emission processes occurring in large-scale volumes of scintillator. The shift in weights toward the longer time constant— q_2 is larger than the expected sum of the long-term component weights by about a factor of 3—is still not well explained. It may be due to light trapping from total internal reflection or other geometric effects in the detector.

In the past, both observed time distributions such as this one, and functions created by Monte Carlo methods, have been used as scintillator PDFs for use in analyzing CTF data. The PDF currently used was derived from a Monte Carlo model that includes the effects of photon absorption and re-emission. It is shown in blue in Figure 7.1. In the relevant computer algorithms, it is modeled as a discrete set of points (t, p(t)); values between the known points are interpolated with a cubic spline. Note that the function in Figure 7.3 is offset slightly in the negative direction, while that in Figure 7.1 is offset about 17 ns in the positive direction. It should also be noted that the CTF 3 function is significantly narrower (~35%) than the CTF 1 function. Perhaps not all effects that could delay photon arrival times at the PMTs were taken into account in the Monte Carlo.

No such effort in observing or modeling a scintillator PDF was undertaken for the PXE scintillator of CTF 2. The PDF used in CTF 2 analysis is the same one used for CTF 3, modeled with a Monte Carlo for a pseudocumene scintillator. It is believed that the systematic errors thereby introduced in CTF 2 position reconstruction are not large.

7.6 A semi-analytical ray tracing approach

The approach taken in the Princeton Borexino group to the problem of position reconstruction in the CTF has been the development of an algorithm to calculate the travel time for light between an arbitrary point in the scintillator fluid and an arbitrary PMT, taking into account refraction at the nylon vessel interface between scintillator fluid and water.

7.6.1 The path from event to PMT

Assume for now that the index of refraction mismatch at the boundary between water and scintillator is the only complication in spatial reconstruction. We now determine the time that a photon would take to travel from a particular test point A in the scintillator volume to a specific PMT at point B. The radius of the CTF vessel is r = 1.0 m, and the distance of each PMT from the vessel center is R = 3.3 m. Consider a coordinate system in which the center of the vessel is at the origin O(0,0,0), and the PMT under consideration is at position B(-R, 0, 0). By rotational symmetry, we can always choose a coordinate system such that the coordinates of the test point are A($x_0, y_0, 0$), with $y_0 \ge 0$. (This coordinate system is unique unless $y_0 = 0$.) As in Chapter 5, denote the vector from O to A as \mathbf{x}_0 , and the vector from O to B as \mathbf{x}_i . We know $|\mathbf{x}_i| = R$; also define $a \equiv |\mathbf{x}_0|$.

Assuming no scattering or absorption and re-emission, the path that would be traversed by a photon between the test point and the PMT has two parts, each a straight line segment: the part from PMT to vessel surface, distance ℓ_1 , and the part from vessel surface to test point, distance ℓ_2 . These distances are bounded by $\ell_1 \in [R - r, \sqrt{R^2 - r^2}]$ and $\ell_2 \in [0, 2r]$. The travel time along the path is given by

$$\tau_f = (n_1 \ell_1 + n_2 \ell_2)/c. \tag{7.17}$$

 n_1 is the index of refraction of water, about 1.33, and n_2 is that of pseudocumene or PXE. For these liquids, $n_2 > n_1$, leading to problems with total internal reflection that will be described in Section 7.6.3.

Define a point C where the photon path is refracted at the vessel surface. (The nylon film composing the vessel has a thickness of 0.5 mm, negligible compared to other dimensions in the detector, and an index of refraction very similar to that of pseudocumene. It can therefore be treated as a surface of zero thickness.) With an ideal spherical vessel, any normal to the vessel surface is also a radius. This and the fact that refraction is a planar phenomenon let us conclude that C is in the *xy*-plane along with A, B, and the origin.



Figure 7.4: Geometry used to determine the path length of a photon traveling from a given point in the CTF scintillator to a specific PMT. The diagram is drawn in the plane common to the test point A, PMT B, and center O of the CTF vessel. The point C where the photon is refracted is also in the same plane if the CTF vessel is perfectly spherical. See text.

As shown in Figure 7.4, let P be the point on line \overrightarrow{OC} such that \overrightarrow{PB} is a perpendicular to \overrightarrow{OC} ; let Q be the point on \overrightarrow{OC} such that \overrightarrow{QA} is a perpendicular to \overrightarrow{OC} . Define $\alpha = m\angle BOC$, $\beta = m\angle BOA$, $\theta_1 = m\angle PCB$, and $\theta_2 = m\angle ACO$. Since the paths of a refracted light ray on either side of the refracting interface are on opposite sides of the normal, $\angle ACO$ and $\angle PCB$ must be on opposite sides of \overrightarrow{OC} . Therefore $\alpha \leq \beta$, as drawn in the figure. Equality holds only when the test point is along the line \overrightarrow{OB} through PMT and origin. Note that we know β a priori in a coordinate-independent way. It is the angle between test point and PMT from the origin, so

$$\cos\beta = \frac{\mathbf{x}_i \cdot \mathbf{x}_0}{R|\mathbf{x}_0|}.\tag{7.18}$$

We can find four equations relating the unknown variables $\ell_1, \ell_2, \theta_1, \theta_2$, and α . The length \overline{BP} is a side of the right triangles ΔBPC and ΔBPO , with respective hypotenuses ℓ_1 and

R, so $\ell_1 \sin \theta_1 = R \sin \alpha$. Length \overline{CP} is a side of ΔBPC , and is r less than a side of ΔBPO , so $\ell_1 \cos \theta_1 = R \cos \alpha - r$. These and analogous considerations for ΔAQC and ΔAQO yield the system

$$\ell_1 \cos \theta_1 = R \cos \alpha - r \tag{7.19}$$

$$\ell_1 \sin \theta_1 = R \sin \alpha \tag{7.20}$$

$$\ell_2 \cos \theta_2 = r - a \cos \left(\beta - \alpha\right) \tag{7.21}$$

$$\ell_2 \sin \theta_2 = a \sin (\beta - \alpha). \tag{7.22}$$

This system can be reduced to an equation relating any two unknowns, so we can write ℓ_1 and ℓ_2 solely in terms of α . Squaring the sides of Equations (7.19) and (7.20) and adding gives

$$\ell_1^2 = R^2 + r^2 - 2Rr\cos\alpha, \tag{7.23}$$

while the same for Equations (7.21) and (7.22) yields

$$\ell_2^2 = a^2 + r^2 - 2ar\cos(\beta - \alpha). \tag{7.24}$$

We do know that $\cos \alpha$ cannot be less than r/R. When $\cos \alpha = r/R \Rightarrow \alpha \approx 72.4^{\circ}$, the path ℓ_1 is tangent to the CTF vessel at C. If α were any greater, this segment would intersect the vessel in two places. Let α_c represent the critical value $\cos^{-1}(r/R)$, so that $\alpha \in [0, \min \{\beta, \alpha_c\}]$. Still, the actual value of α (equivalently, the location of point C along the rim of the CTF vessel) is not yet completely determined.

7.6.2 Enforcing Snell's law

One more equation is required to determine a unique value of α . This equation is provided by Snell's Law,

$$n_1 \sin \theta_1 = n_2 \sin \theta_2, \tag{7.25}$$



Figure 7.5: The angle α between PMT and point of refraction, as a function of β (the angle between PMT and event) and a (the test point radial coordinate) in CTF 3; refer to Figure 7.4. β is given on the horizontal axis (in degrees) and α on the vertical axis (in degrees). Values of α greater than $\alpha_c \approx 72.4^{\circ}$ are not physically meaningful since they would imply a light path with an angle of refraction $\theta_1 > 90^{\circ}$. Curves are shown for (from bottom to top) a = 10, 30, 50, 70, 85, 90 and 100 cm. Notice that the top two curves have values of β for which α is not defined in the range [0, 72.4°]; these "dark zones" are due to total internal reflection. These two curves are also multi-valued in part of their range, indicating that a photon may take one of two possible paths from event to PMT.

which specifies the relationship between the angle of incidence and angle of refraction at an interface between substances with different indices of refraction.

Dividing Equation (7.19) by (7.20) and Equation (7.21) by (7.22) gives, respectively,

$$\cot \theta_1 = \cot \alpha - \frac{r}{R} \csc \alpha \tag{7.26}$$

$$\cot \theta_2 = \frac{r}{a} \csc \left(\beta - \alpha\right) - \cot \left(\beta - \alpha\right). \tag{7.27}$$

We square each of Equations (7.26) and (7.27) and add one in order to obtain $\csc^2 \theta_{1,2}$ on the left-hand sides, using the Pythagorean identity $\cot^2 \theta_{1,2} + 1 \equiv \csc^2 \theta_{1,2}$. These may then be substituted into Equation (7.25) once it has been manipulated into the form



Figure 7.6: The time of flight τ_f of a photon from a test point inside the CTF vessel to a specific PMT. τ_f is shown as a function of β (the angle between PMT and test point) and a (the test point radial coordinate) in CTF 3; refer to Figure 7.4. β is given on the horizontal axis (in degrees) and τ_f on the vertical axis (in nanoseconds). From top to bottom on the left side of the figure (bottom to top on the right), the curves are shown for a = 0, 10, 30, 50, 70, 90 and 100 cm. Note the gap in the curves for a = 90 and 100 cm, centered at about $\beta = 100^{\circ}$, representing the "dark zones." The rectangle shows the area of the graph covered by the expanded view below.



Figure 7.7: Times of flight for the upper and lower branches of the functions $\alpha(\beta, a)$ for a = 90 cm (bottom curves) and 100 cm (top curves). β is in degrees, and τ_f in nanoseconds. The figure shows that the difference in time of flight between upper and lower branches is negligible (< 0.1 ns).

 $\csc^2 \theta_2 = (n_1/n_2)^2 \csc^2 \theta_1$. The result is

$$1 + \left[\frac{r}{a}\csc\left(\beta - \alpha\right) - \cot\left(\beta - \alpha\right)\right]^2 = \left(\frac{n_2}{n_1}\right)^2 \left[1 + \left(\cot\alpha - \frac{r}{R}\csc\alpha\right)^2\right].$$
 (7.28)

This equation may finally be solved for α . This cannot be done in closed form. In the practical case of the reconstruction software, the equation is manipulated so that one side is equal to zero, and then is fed to a custom numerical root-finding routine for specific values of a and β . The algorithm used combines bisection and linear interpolation with some additional tricks, for instance, tightening the assumption of function linearity as the interval containing the root becomes smaller. It typically finds a value for α to within 10^{-8} radians with 10–20 calls to the function being evaluated.

The form of α as a function of β for different values of a is shown in Figure 7.5. The time of flight as a function of β for different values of a is shown in Figure 7.6. For some values of a, the function $\alpha(\beta, a)$ is multi-valued. But this does not notably affect reconstruction; the differences in times of flight between the branches are negligible, as shown in Figure 7.7. It therefore is not very important which solution is obtained by the root finder. For other values of a, α is undefined. These "dark zones" due to total internal reflection, which are much more problematic, will be discussed in Section 7.6.3.

We next ask how to define the time of flight to a PMT for points outside the CTF vessel. One option that seems appealing at first glance is to avoid the question by forcing all events to be reconstructed within the vessel volume, for instance by defining the likelihood function of an event to be zero (and $-\log \mathcal{L} \to \infty$) when $|\mathbf{x}_0| > r$. This is physically realistic, as of course scintillation photons can only come from within the scintillator volume. Unfortunately, this tactic spoils the statistics of event position reconstruction developed in Section 5.5. It is intuitively clear that surface events should be reconstructed to have an average position at about $|\mathbf{x}_0| = r$, with roughly equal numbers of events reconstructed with positions inside and outside the vessel. Forcing all events to be reconstructed within the vessel introduces a statistical bias shifting the distribution of events in r towards r = 0.



Figure 7.8: Geometry for determining the time of flight from a test point (A) in the water buffer to a PMT (B) which is behind the CTF vessel as seen from the test point. The light traveling between them is refracted twice at C and D.

A seemingly better approach is to use the actual time of flight that would be seen if an event really did occur in the water buffer. For a test point between the vessel and PMT, or away from the lines of sight between the vessel and PMT; that is, one for which the condition

$$(0 \le \beta \le \alpha_{\rm c} \text{ and } |\mathbf{x}_0| \ge r) \text{ or } (\alpha_{\rm c} < \beta < \alpha_{\rm c} + \pi/2 \text{ and } |\mathbf{x}_0| \ge r \sec(\beta - \alpha_{\rm c}))$$
(7.29)

holds; the time of flight is simply $\tau_f = (n_1/c)|\mathbf{x}_i - \mathbf{x}_0|$.

For points behind the vessel as seen from the PMT, from which light must be refracted twice, consider the geometry shown in Figure 7.8. Here, the total time of flight is $\tau_f = n_1(\ell_1 + \ell_3)/c + n_2\ell_2/c$. From the figure, we may derive seven equations in the unknowns $\ell_1, \ell_2, \ell_3, \theta_1, \theta_2, \alpha_1$ and α_2 . The first four are roughly analogous to Equations (7.19) to (7.22), the fifth is Snell's Law, and the last two come from the properties of Δ COD:

$$\ell_{1} \cos \theta_{1} = R \cos \alpha_{1} - r \qquad n_{1} \sin \theta_{1} = n_{2} \sin \theta_{2}$$

$$\ell_{1} \sin \theta_{1} = R \sin \alpha_{1} \qquad \ell_{2} = 2r \cos \theta_{2}$$

$$\ell_{3} \cos \theta_{1} = a \cos \alpha_{2} - r \qquad \alpha_{1} + \alpha_{2} + 2\theta_{2} = \pm (\pi - \beta)$$

$$\ell_{3} \sin \theta_{1} = a \sin \alpha_{2}$$

The \pm sign in the last equation comes from the ambiguity in whether the test point A is on

the same side of line \overrightarrow{OB} as the points of refraction C and D (as shown in Figure 7.8), or the opposite side from them (imagine segment \overrightarrow{DA} being extended until point A is below the *x*-axis). As a result, the solution will be multi-valued for some values of β . Just as in the case of a test point inside the vessel (times of flight for the two branches are shown in Figure 7.7), this is not expected to affect the time of flight much.

Without going into detail, this system of equations may be reduced to a single equation in one variable,

$$\cot \alpha_1 - \frac{r}{R} \csc \alpha_1 = \cot \left[\alpha_2(\alpha_1)\right] - \frac{r}{a} \csc \left[\alpha_2(\alpha_1)\right]$$
(7.30)

where

$$\alpha_2(\alpha_1) = \pm (\pi - \beta) - \alpha_1 + 2\sin^{-1} \left\{ \frac{n_1}{n_2} \left[1 + \left(\cot \alpha_1 - \frac{r}{R} \csc \alpha_1 \right)^2 \right]^{-1/2} \right\}$$
(7.31)

and solved numerically for α_1 . The lengths of the path segments are then

$$\ell_1 = \sqrt{R^2 + r^2 - 2Rr\cos\alpha_1} \tag{7.32}$$

$$\ell_2 = 2r \cos\left[\frac{\alpha_1 + \alpha_2 \mp (\pi - \beta)}{2}\right]$$
(7.33)

$$\ell_3 = \sqrt{a^2 + r^2 - 2ar\cos\alpha_2}. \tag{7.34}$$

7.6.3 The problem of dark zones

For each PMT in the CTF, there exists a corresponding region inside and behind the CTF vessel from which light traveling in a straight line through each fluid volume cannot reach that PMT. We refer to these regions as "dark zones." They are illustrated in Figure 7.9, showing the time of flight from a test point to a PMT at (-R, 0, 0) as a function of the test point position. Figure 7.9a shows the time of flight for a system with a single index of refraction. Figure 7.9b shows the time of flight for the CTF 3 system. The situation is worse for CTF 2, since the index of refraction of PXE is even greater than that of pseudocumene. For each PMT triggered in an event, the event likelihood function within that PMT's dark



Figure 7.9: The time of flight function from a test point to a PMT located at (-R, 0, 0), at the left edge of each diagram. The time of flight is shown in the xy-plane as contours of constant time with colored bands between the contours. (Contours are spaced at intervals of 0.625 ns.) On the left (part a) is the time of flight in a single index of refraction system. Contours are simply circles centered on the PMT. In this particular diagram, n = 1.504, but the graph would look similar except for the width of the contours for any n. On the right (part b), the time of flight is shown for the CTF 3. The smaller circle represents the CTF vessel, and the large circle on the outer border of each graph is the location of the set of PMT photocathodes. The dark gray triangular regions in the graph on the right are the "dark zones." Unless it is scattered, a photon cannot travel from a point in the dark zone of a PMT to the center of that PMT's photocathode. Note also the discontinuity separating points behind the CTF vessel from those in the rest of the water buffer.

zone is identically zero. In an ideal world, these dark zones would have little effect on the reconstruction, because it would be impossible for an event to have occurred within the dark zone of any PMT that triggered.

In reality, though, a photon emitted within the dark zone of a PMT may reach that PMT in several ways. It may have been absorbed outside of the dark zone, and then re-emitted in a different direction to end up at the PMT. Because PMTs are finite in extent, it may have been refracted into the edge of that PMT and detected. Finally, the CTF vessel is not a perfect sphere, so a photon may be refracted at the nylon film in an unexpected direction.



Figure 7.10: The logarithms of likelihood functions in the xz-plane of a ²¹⁴Po event occurring very near the edge of the CTF 2 vessel. This event occurred within a radon source, nominally located at the very top of the vessel, in Run 795. The maximum values are in blue; very small values are shown in red and then white. At left (part a) is shown the logarithm of the likelihood function using a single value of n = 1.565, that of PXE, throughout the CTF tank. The circle is the CTF vessel. Contour intervals are 0.5 units. The black area is the region where $\log \mathcal{L}_{max} - \log \mathcal{L} < 0.5$. At right (part b) is shown the log likelihood function using the time of flight function derived implicitly in Section 7.6.2. The union of the dark zones of each triggered PMT, where the likelihood function is identically zero, is shown in dark gray. Here, the event position occurs within some of these dark zones! The likelihood function values outside the vessel are non-zero only in disconnected regions.

Any of these effects may lead to an event that occurred in the dark zone of one or more triggered PMTs. If we naïvely try to maximize the likelihood function (or equivalently, minimize $-\log \mathcal{L}$) to find the position of the event, we will be unable to do so, as shown in Figure 7.10. Either the minimization routine running over $-\log \mathcal{L}$ will fail to converge, or it will find a false minimum, most likely at a distance less than the minimum distance of the dark zone from the center of the CTF.

The dark zones are generally not a problem for events whose true position is within about 50–60 cm of the vessel center. For events outside that volume, the fraction of events on which the minimization fails to converge increases rapidly as a function of radial coordinate. This

behavior can be tested to some precision using data from the source runs in CTF 2 and CTF 3, in which a small (few cm diameter) radon source was lowered into place at various positions inside the CTF vessel. These source runs will be discussed in more detail in Sections 7.7 and 7.8.

Removal of photoelements

Conceptually, the simplest way to work around the problems caused by the dark zones is to ignore timing information from the PMTs that probably should not have triggered. Let the last point reached by the minimization routine before it failed to converge be \mathbf{x}^* . For each PMT, find the time of flight function from that point, $\tau_f^i(\mathbf{x}^*)$. If τ_f^i is undefined, then \mathbf{x}^* is in the dark zone of PMT *i*, so ignore PMT *i*. Then, try minimizing the likelihood function again, considering only the PMTs that remain. To be doubly sure that all potentially problematic PMTs are removed, we modify this procedure by multiplying \mathbf{x}^* by $0.95r/|\mathbf{x}^*|$ before finding the times of flight, putting the point deep into the region where dark zones occur.

In practice this method does not work well at all. For CTF 2 source run 805, which had a nominal position at z = -62 cm, the fraction of ²¹⁴Po events for which position reconstruction failed to converge was 8.5%. This technique of removing problematic PMTs only reduced the reconstruction failure rate to 5.0%. For source run 803, nominally at z = -83 cm, removing some PMTs reduced the reconstruction failure rate from 48.5% to 31.0%. That is, only about 40% of events for which reconstruction at first fails are fixed by the PMT removal technique. Furthermore, in each of these runs, the average number of PMTs that were discarded from events for which PMT removal was necessary was about 25. This number is much higher than might be expected; it may be a result of the high efficiency with which the PMT light cones guide light from within the vessel into each photocathode. It seems statistically unwise to discard more than 25% of one's data set for each event. For all these reasons, a different solution was sought.



Figure 7.11: Geometry needed for a change of coordinate systems. As before, A is a test point, B is a PMT, and O is the center of the CTF vessel. ℓ_1 and ℓ_2 are the path a photon would take from A to B. We wish to switch from the vessel-centered coordinate system (a, β) to the PMT-centered coordinate system (ρ, θ) .

Cubic spline interpolation

By this point, one may be thinking that it would be easier to simply ignore the dark zones by continuing the time of flight function through them somehow. One possible approach to take is to use a cubic spline interpolation between the times of flight for test points within the vessel and those in the water buffer. If the test point is anywhere except a dark zone or the region behind the dark zone as seen from the PMT under consideration, then the previously derived implicitly defined time of flight function is used. Otherwise the cubic spline interpolation is used.

For this task, the vessel-centered coordinate system (a, β) is not the most natural to use. It is simpler to convert to a PMT-centered polar coordinate system and consider the variables $\rho \equiv |\mathbf{x}_i - \mathbf{x}_0|$ (distance from the PMT to an event) and θ (angle of the event away from the axis through the vessel center and PMT). Two equations immediately result from applying



Figure 7.12: Cubic spline interpolation over dark zones. At left (part a) is the time of flight function in CTF 3 with a cubic spline interpolation over the dark zone for a PMT at the left edge of the figure. At right (part b) is the logarithm of the likelihood function in the xz-plane obtained by using the interpolated time of flight on the same ²¹⁴Po event shown in Figure 7.10. There are no more dark zones, but some discontinuities still appear in the upper left of the plot. The CTF vessel is represented by the smaller circle in part a and by the circle in part b.

the Law of Cosines on $\triangle ABO$ in Figure 7.11:

$$a(\rho,\theta) = \sqrt{R^2 + \rho^2 - 2R\rho\cos\theta}$$
(7.35)

$$\beta(\rho,\theta) = \cos^{-1}\left[\frac{R^2 + a^2(\rho,\theta) - \rho^2}{2Ra(\rho,\theta)}\right] = \cos^{-1}\left[\frac{R - \rho\cos\theta}{a(\rho,\theta)}\right].$$
(7.36)

The method chosen for the cubic spline interpolation of the time of flight function is as follows. On any curve of constant ρ (part of a circle centered at the PMT), τ_f is a function only of the angle θ . Let the gap over which we wish to interpolate be $\theta \in [\theta_-, \theta_+]$. On the θ_- side of the gap, τ_f may be calculated (numerically) as a function of (ρ, θ) instead of (a, β) using the change of coordinates, Equations (7.35) and (7.36). Its derivative, $(\partial \tau_f / \partial \theta)_{\theta_-}$, is determined numerically. On the θ_+ side of the gap (in all cases, $\theta_+ = \sin^{-1}(r/R) = \pi/2 - \alpha_c$), τ_f is constant and equal to $n_1 \rho/c$. Given these parameters, the cubic spline interpolation of τ_f on a curve of constant ρ in the interval $\theta \in [\theta_-, \theta_+]$ is defined by

$$\tau_f(\rho,\theta) = \frac{n_1\rho}{c} + \left(\frac{\partial\tau_f}{\partial\theta}\Big|_{\theta_-} - 3\Delta\tau_f\right) \left(\frac{\theta_+ - \theta}{\Delta\theta}\right)^2 - \left(\frac{\partial\tau_f}{\partial\theta}\Big|_{\theta_-} - 2\Delta\tau_f\right) \left(\frac{\theta_+ - \theta}{\Delta\theta}\right)^3, \quad (7.37)$$

where $\Delta \theta \equiv \theta_+ - \theta_-$ and $\Delta \tau_f \equiv n_1 \rho/c - \tau_f(\rho, \theta_-)$. The cubic spline interpolation matches both the function value and the first derivative with respect to θ at the two endpoints of the gap.

The time of flight function with added cubic spline interpolation is shown for CTF 3 in Figure 7.12a. At first one might expect that the cubic spline interpolation solves the problem of the spatial reconstruction failing to converge. Unfortunately, this may still happen as a result of the discontinuities shown in Figure 7.12b; they may arise from the multi-valued nature of the time of flight function. It is also possible that the derivative $\partial \tau_f / \partial \theta$ is not calculated numerically to sufficient accuracy. Roughly 0.2% of the ²¹⁴Po events in the CTF 2 source run 803 still do not converge when τ_f is defined by Equation (7.37) over the dark zones.

Analytic continuation

A third approach to consider, which does not suffer from the problem of numerical approximations where functions are spliced together, is to expand the time of flight function inside the vessel in a Taylor series. This also avoids any problems due to multiple branches in the time of flight function. In this approach, the time of flight function is taken as the first several terms of the series, and it is carried over through the dark zones and water buffer (ignoring the true time of flight function for the water buffer).

In the (ρ, θ) coordinate system, the time of flight function for a single index of refraction system is trivial: $\tau_f(\rho, \theta) = n\rho/c$. We now consider the time of flight function in the two index of refraction CTF geometry as a perturbation of this function. The function for the two index of refraction system is plotted as a function of ρ and θ in Figure 7.13a.



Figure 7.13: On the left in part (a) is shown the time of flight τ_f to a PMT from a point inside the CTF 3 vessel with coordinates (ρ, θ) in the PMT-centered coordinate system, as a function of θ for various values of ρ . θ is shown on the horizontal axis in degrees (its maximum value for a point inside the vessel is $\sin^{-1}(r/R) \approx 17.6^{\circ}$). From top to bottom, the curves represent $\rho = 4.20, 4.05, 3.80, 3.55, 3.30, 3.05, 2.80, 2.55, and 2.40 m$. The envelope of the curves is caused on the top and bottom by the spherical shape of the vessel, and at upper right by the boundary of the dark zone. On the right in part (b), the difference $-\Delta \tau_f \equiv \tau_f(\rho, \theta = 0) - \tau_f(\rho, \theta)$ is shown as a function of θ for the same values of ρ (curves for greater ρ are higher up). On both graphs, the longest curve is that for $\rho = 3.05$ m. The vertical axes are in units of nanoseconds.

It is clear that the time of flight function on the axis through the vessel center and PMT has the form (within the scintillator volume) of

$$\tau_f(\rho, \theta = 0) = \frac{n_1}{c} (R - r) + \frac{n_2}{c} \left[\rho - (R - r) \right].$$
(7.38)

The first term is the time it takes light to reach the PMT from the closest point on the vessel, and the second term is the elapsed time from event to vessel surface. Neglecting scattering, this is an exact result; a photon traveling along this axis is not refracted. Notice this is equivalent to the time of flight in a single index of refraction system with $n = n_2$ and a time offset (which will not affect the reconstruction) of $(n_1 - n_2)(R - r)/c$.

In Figure 7.13b is shown the time of flight as a function of θ for several fixed values of ρ in

CTF	C_{00}	C_{01}	C_{20}	C_{21}	C_{22}	C_{40}	C_{41}	C_{60}	C_{80}
2	15.45	5.220	-3.448	-0.418	$\sim +0.07$	-14.869	-7.2	-99.7	~ -950
3	15.25	5.017	-2.442	-0.219	$\sim +0.05$	-9.997	-3.8	-63.8	~ -600

Table 7.4: Coefficients of the Taylor expansion of $\tau_f(\rho, \theta)$ about (R, 0) as defined in Equation (7.39): C_{ij} is the coefficient of $\theta^i [(\rho - R)/r]^j$. Here, $C_{00} = n_1 R/c + (n_2 - n_1)r/c$, and $C_{01} = n_2 r/c$; all other C_{0j} are zero, as are all C_{ij} with *i* odd. All C_{ij} are given in units of nanoseconds. Notice that the absolute values of the C_{ij} $(i \ge 2)$ are greater for CTF 2 than for CTF 3. This was expected: since $n(\text{PXE}) > n(\text{PC}) > n(\text{H}_2\text{O})$, the CTF 2 is a greater perturbation from a single index of refraction system than the CTF 3 is.

the range 2.55–4.05 m, with Equation (7.38) subtracted out. This must be an even function of θ since Figure 7.11 can be flipped across the horizontal axis with no effect on the time of flight. Therefore we write

$$\tau_f(\rho,\theta) = \frac{n_1 - n_2}{c}(R - r) + \frac{n_2}{c}\rho + \sum_{i=1}^{\infty} \sum_{j=0}^{\infty} C_{2i,j} \left(\frac{\rho - R}{r}\right)^j \,\theta^{2i} \tag{7.39}$$

(θ is raised only to even powers).

The coefficients may be determined using numerical methods; the lower-order ones are tabulated in Table 7.4 for both CTF 2 and CTF 3. (The coefficients for CTF 2 and CTF 3 are different, since pseudocumene and PXE have different indices of refraction; an interesting extension of this work would be to determine the coefficients as functions of the index of refraction of the scintillator.) Using these values in the Taylor expansion, Equation (7.39), and neglecting all higher-order terms, the error in time of flight is at most 0.1 ns (equivalent to 3 cm, much less than the detector resolution) even at the edges of the scintillator volume.

The time of flight function defined by the Taylor series expansion (only including the terms with coefficients given in Table 7.4) is shown in Figure 7.14a. Note that the function value drops off quickly away from the CTF vessel; this is a desirable feature since most events should be reconstructed near it or inside it. Because of this drop-off, the likelihood functions become small outside of the vessel much more quickly than with other time of flight functions, as seen in Figure 7.14b. More importantly, unlike the exact time of flight function or the cubic spline interpolation, no discontinuities or undefined regions appear



Figure 7.14: The CTF 3 time of flight function (part a, at left) and a CTF 2 event likelihood function (part b, at right) defined by the truncated Taylor series of Equation (7.39) and Table 7.4. The CTF vessel is represented by the smaller circle in part a and by the circle in part b. The event shown in part b is the same as in Figures 7.10 and 7.12b. Note that the likelihood function becomes small outside of the vessel much more quickly than in previous figures.

in the likelihood functions resulting from the Taylor series of the time of flight. This is of course a consequence of the definition of the time of flight as a polynomial series.

One result of using this Taylor series expansion for the time of flight function in reconstructing the CTF 2 source runs was surprising. For a small fraction (a few thousandths) of ²¹⁴Po events in each run, the minimization routine did not converge. However, when these events were examined individually with the evdisp tool, they were found not to be reconstructed correctly with any other time of flight function either. The minimization routine reported convergence, but the position obtained, in all cases, is several meters outside the CTF vessel. The problem common to all of these events appears to be that a single TDC channel is triggered more than 15 ns before any other channel. This is likely a result of PMT dark noise. Once recognized, this problem was worked around in software, but it affects so few events that the gains are negligible.
7.7 Testing algorithms with CTF 2 source runs

The CTF 2 source runs consist of a series of fifteen data acquisition periods with run numbers in the range 791–810, acquired during September 2000. (A few additional runs also received an official number in this range but were of too short a duration to provide statistically useful data.) During each run, a small ²²²Rn source was held at a constant position inside the CTF vessel. These runs allow comparison of the performance of position and energy reconstruction software with the actual position and energy of events due to radioactive decays in the radon source. In particular, decays of the radon daughter ²¹⁴Po are easy to identify due to its short half-life of $164 \,\mu$ s; they appear as the double event or "coincidence" ²¹⁴Bi \rightarrow ²¹⁴Po \rightarrow ²¹⁰Pb. The α decay of ²¹⁴Po is monoenergetic, and because of α quenching, it appears to have an energy within the neutrino energy window of 250–800 keV. These properties make ²¹⁴Po events ideal for testing reconstruction software.

7.7.1 The source calibration hardware

Since the CTF 2 source runs are described in detail in reference [72], we will only remark upon factors relevant to tests of the reconstruction software. As has been mentioned already, two different radon sources were used. Each consisted of a small quartz vial, roughly cylindrical in shape, filled with radon-laced PXE scintillator. Quartz was used since its index of refraction is near that of the scintillator, and it is transparent to the near-UV wavelengths of scintillation light. Both sources were partially quenched by the presence of oxygen in the scintillator, causing events to produce less scintillation light than expected. The quenching was unintentional, but the oxygen contamination unfortunately could not be removed; stripping the PXE would also purge it of radon.

Each source was attached to the end of a set of neutrally buoyant stainless steel rods, each about 1 m in length. These rods were inserted into the north end pipe at the top of the CTF tank and lowered until the source was inside the CTF vessel. (A nitrogen blanket was



Figure 7.15: The CTF 2 source insertion system. A set of stainless steel rods was lowered into the CTF vessel from above. A thin nylon string was attached to a hinged rod at the end, allowing the source (hanging from the end of the rod) to be raised into an off-axis configuration by pulling on the upper end of the string. Figure adapted from reference [72].

used to prevent normal outside air from contaminating the CTF scintillator.) By raising or lowering the rods, as shown in Figure 7.15, the source could be positioned at the desired z-coordinate. Rods were added or removed at the top of the assembly as needed to allow for the low clearance of the ceiling of the clean room atop the CTF tank. The lowest rod was hinged, with a cord attached near the hinge, permitting it to be pulled up at a right angle to the rest of the rods so the source could be positioned away from the CTF vessel z-axis. In this configuration, the source was held 48 cm from the vessel axis, and could be positioned anywhere on a cylinder of this radius by rotating and raising or lowering the rod assembly.

Source 1, used in Runs 791–795, had a diameter of 1.7 cm and length of 9.5 cm. The scintillator in this source was quenched by oxygen such that the CTF detected 200 ± 16 photoelectrons from ²¹⁴Po α decays in the source, compared to the expected number of 300 ± 25 . Over the period this source was used, its rate of ²¹⁴Po events fell from 50 to

30 mBq as the radon decayed away. Source 2, used in Runs 797–810, had a diameter of 1 cm and length of 5.5 cm. Its scintillator was more strongly quenched, with the ²¹⁴Po events having a yield of 100 ± 12 photoelectrons. Its rate of ²¹⁴Po decays ranged from 250 to 100 mBq. Only Source 2 could be positioned off-axis.

The nominal positions of the sources are reported for each run in columns 3–7 of Table 7.5. These positions have some systematic uncertainty. In the z direction, the position of the top of the CTF vessel was uncertain to about 8 cm because it is possible that the vessel shifted in position during the CTF 2 water filling operation. In the xy-plane, the angle of rotation of the hinged arm holding the source was difficult to set accurately. Furthermore, when the arm is raised to the horizontal, the remainder of the metal rods are tilted slightly away from the vertical. The positions are therefore recorded in [72] as having possible systematic errors of

- $\Delta x = \Delta y = \pm 10 \text{ cm}$ for off-axis points
- $\Delta x = \Delta y = \pm 5 \text{ cm}$ for on-axis points
- $\Delta z = \pm 5 \,\mathrm{cm}.$

7.7.2 Accuracy of reconstruction software with ²¹⁴Po events

The ²¹⁴Po events in Runs 791–810 were identified using cuts on the data set. The number of candidate ²¹⁴Po events in each run is reported in column 2 of Table 7.5. A candidate event was required to meet the following criteria:

- it was detected by the Group 2 electronics;
- the photoelectron yield was in the range 140–260 photoelectrons (for Source 1) or 70–140 photoelectrons (for Source 2);

- the coincidence time after the corresponding Group 1 event was in the range $10-1000 \,\mu s$;
- the corresponding Group 1 event had a photoelectron yield in the range 100–1200 photoelectrons (the lower limit was set to avoid accidental coincidences with ¹⁴C events);
- neither event of the coincidence triggered the muon veto system.

For each run, histograms were made of the x, y, and z coordinates of the reconstructed positions of every candidate event. Each histogram was fit to a Gaussian curve. In Table 7.5, the results of these fits for the reconstruction software developed at Milan and Munich are reported. Both codes were optimized for the scintillator response function of pseudocumene, but this is not expected to cause large errors.

In the Milan software results, note that the reconstructed positions of runs 802–805, having a very negative z coordinate, are offset several cm in the positive x direction. This phenomenon was investigated for the present work; it was found to be a result of TDC channel 15 (attached to PMT 76) often triggering earlier than it should have. In Run 793, in which the source was nominally at the center of the vessel, this channel triggered first in 872 of over 3000 ²¹⁴Po events. (In second place was TDC channel 47, attached to PMTs 52 and 53, which triggered first in only 90 events). It is not clear whether this is a calibration table error or an electronics noise problem, nor why the problem did not appear in the Munich software results.

Further investigations with Run 793 showed that the probability for each single-PMT TDC channel to be the first triggered in an event, which should have a uniform value over PMTs of 1% (with a Poissonian error of $\pm 0.2\%$ given the number of events in the run), in fact ranges between 0 and 2.2%. A histogram of the number of times each single-PMT TDC channel was the first to trigger showed no obvious peak. Moreover, increased probability for a channel to trigger first in an event does not always correlate with an increased average

	214 D			1	• . •			3 4.1	1			1.6	1 1	
Run	214Po	IN	omir	ial po	ositio	n		Milan	code			Munic	h code	
	events	x	y	z	ρ	r	x	y	z	$ \Delta \mathbf{x} $	x	y	z	$ \Delta \mathbf{x} $
793	3113	0	0	0	0	0	3.7	0.3	-4.9	6.1	0.3	0.1	-1.4	1.4
791	2403	0	0	32	0	32	3.4	4.4	28.4	6.6	0.0	3.8	35.7	5.3
797	13377	0	48	0	48	48	4.1	52.2	-3.4	6.8	-0.1	54.0	1.9	4.4
798	3452	-48	0	0	48	48	-47.2	4.7	0.9	4.9	-52.3	5.7	4.2	8.3
799	3039	27 -	-40	0	48	48	26.7	-41.3	-2.1	2.5	22.4	-43.8	2.7	6.5
810	6423	27 -	-40	11	48	49	31.0	-38.4	10.6	4.3	26.3	-40.7	17.4	6.5
800	11833	40 -	-27	-30	48	57	46.7	-25.5	-37.1	9.9	40.7	-27.9	-32.3	2.6
801	5269	-40	27	-30	48	57	-41.6	22.8	-33.4	5.6	-47.1	23.4	-32.1	8.2
805	3337	0	0	-62	0	62	6.1	-1.1	-68.4	8.9	1.4	-1.7	-69.1	7.4
807	6682	-40	27	-56	48	74	-42.3	28.1	-63.1	7.5	-47.6	30.0	-65.1	12.2
809	2504	$-40 \cdot$	-27	-56	48	74	-41.5	-18.3	-62.2	10.8	-49.6	-19.3	-63.6	14.5
804	7544	0	0	-75	0	75	6.1	2.9	-81.9	9.7	0.0	3.0	-82.6	8.2
802	8939	0	0	-78	0	78	9.0	-1.1	-87.5	13.1	0.4	-1.7	-88.2	10.3
803	1133	0	0	-83	0	83	8.0	-1.5	-94.2	13.8	-1.0	-2.1	-95.7	12.9
795	1811	0	0	100	0	100	-	-	-	-	-	-	-	-

Table 7.5: CTF 2 source runs, in order of increasing distance r from the center of the CTF vessel. This table gives the total number of events meeting the ²¹⁴Po cuts, the nominal position of the source, and the mean positions obtained by reconstruction with the Milan software (Section 7.1.1) and the Munich software (Section 7.1.2) as reported in reference [72]. The effective index of refraction used in the Milan code was unfortunately not recorded. Reconstructed positions for Run 795, at the north pole of the vessel (z = 100 cm), were not provided. All positions are given in cm, relative to the nominal center of the CTF vessel. $|\Delta \mathbf{x}|$ is the distance between reconstructed position and nominal position. For both codes, it tends to increase as r increases. This effect appears mainly due to a bias in the reconstruction codes that pushes apparent source positions farther away from the origin by 10–15%.

number of photoelectrons detected on that channel. This indicates that the TDC calibration tables are not as accurate as one might desire. The present software developed at Princeton works around the problem of TDC channel 15 by discarding data from this channel in the position reconstruction of CTF 2 runs, but cannot solve the more general issue of ensuring calibration table accuracy.

Table 7.6 reports the results of these Gaussian fits of the reconstructed source positions for the reconstruction software developed at Princeton, using various methods for calculating the time of flight described in Section 7.6. Because the average photoelectron yield of ²¹⁴Po events was observed to vary significantly between runs 791, 793, and 795, the required energy range of 140–260 photoelectrons for candidate Group 2 events produced in Source 1 was relaxed to 100–300 photoelectrons in these analyses. This changed the number of candidate ²¹⁴Po events to 2468 for Run 791, 3136 for Run 793, and 2136 for Run 795. The photoelectron yield of Source 2 events showed no such variations.

The first observation to be made about the CTF 2 source run reconstruction using the Princeton software is that the three time of flight functions used in creating Table 7.6 produce very similar results. In comparing the various reconstruction codes, it is sufficient to consider only the truncated Taylor series approximation to the time of flight function.

Comparing the values of $|\Delta \mathbf{x}|$ in Tables 7.5 and 7.6, while interesting, is ultimately not very informative. The problem is that the nominal x and y coordinates of the source are not known independently; they were calculated using the 48-cm length of the hinged arm and the observed angle of the arm in the xy-plane. However, that angle could not be measured with much precision, and has a significant potential error. It is more telling to note the difference between the horizontal distance ρ of the reconstructed source from the z-axis and the actual length of the source arm. This difference, which we will denote $\Delta \rho$, is reported in Table 7.7. We also assume there is a constant z offset due to the uncertainty in the position of the CTF vessel. If we calculate this offset by calculating the difference Δz between reconstructed and nominal source z coordinate, and determining the average $\overline{\Delta z}$ independently for each of the

D	Nominal parities				C.	1	/DV	· T 1)	<u> </u>	1.	• ,	1					
Run	Non	nınal	posr	tion	Sin	gle n =	$= n(\mathbf{PX})$.E)	Cubic	spline	interpo	plation	Trune	cated 1	aylor s	series	
	x	y	z	r	x	y	z	$ \Delta \mathbf{x} $	x	y	z	$ \Delta \mathbf{x} $	x	y	z	$ \Delta \mathbf{x} $	
793	0	0	0	0	0.5	-0.6	-1.1	1.3	0.5	-0.6	1.1	1.3	0.5	-0.6	1.1	1.3	
791	0	0	32	32	-0.1	2.7	32.4	2.7	-0.1	2.7	32.6	2.8	-0.1	2.7	32.6	2.8	
797	0	48	0	48	-1.5	49.6	4.7	5.2	-1.6	49.7	4.8	5.3	-1.5	49.7	4.8	5.3	
798	-48	0	0	48	-49.1	4.0	4.8	6.3	-49.4	4.1	4.9	6.5	-49.5	4.1	4.9	6.6	
799	27	-40	0	48	21.9	-41.7	4.8	7.2	22.0	-41.7	5.0	7.2	22.0	-41.7	4.9	7.2	
810	27	-40	11	49	25.6	-38.0	16.4	5.9	25.7	-38.0	16.6	6.1	25.7	-38.0	16.6	6.1	
800	40	-27	-30	57	38.3	-26.3	-25.7	4.7	38.2	-26.4	-25.7	4.7	38.2	-26.5 ·	-25.7	4.7	
801	-40	27	-30	57	-44.3	21.1	-26.2	8.2	-44.4	21.1	-26.4	8.2	-44.5	21.2 ·	-26.4	8.2	
805	0	0	-62	62	2.0	-1.6	-59.2	3.8	2.0	-1.5	-60.0	3.2	2.0	-1.5 ·	-60.1	3.1	
807	-40	27	-56	74	-43.2	26.9	-54.5	3.5	-43.3	27.0	-55.3	3.4	-43.5	$27.2 \cdot$	-55.4	3.6	
809	-40	-27	-56	74	-43.4	-18.6	-54.3	9.2	-43.6	-18.6	-55.3	9.2	-43.6	-18.6 ·	-55.4	9.2	
804	0	0	-75	75	0.7	2.9	-71.5	4.6	0.5	3.0	-72.3	4.1	0.5	3.0 -	-72.5	3.9	
802	0	0	-78	78	0.6	-2.7	-74.0	4.9	0.5	-2.6	-75.1	3.9	0.6	-2.6 ·	-75.0	4.0	
803	0	0	-83	83	-0.9	-1.4	-79.8	3.6	-1.0	-1.5	-81.2	2.5	-0.8	-1.3 ·	-81.3	2.3	
795	0	0	100	100	-10.6	2.6	101.7	11.0	-9.6	1.7	101.4	9.8	-9.7	2.1	99.7	9.9	

Table 7.6: CTF 2 source runs, in order of increasing radial coordinate r. For each run, this table gives the mean position of the ²¹⁴Po events as reconstructed by the software developed at Princeton. The first set of coordinates comes from assuming the CTF is a single index of refraction system with n equal to that of PXE. The second set comes from the cubic spline interpolation of the time of flight function over dark zones, and the third comes from the truncated Taylor series for the time of flight with coefficients given in Table 7.4. See Section 7.6.3 for more details. All positions are given in cm, relative to the nominal center of the CTF vessel. $|\Delta \mathbf{x}|$ is the distance between reconstructed position and nominal position.

		Nom	inal		M	ilan co	de	Mu	nich co	ode	Princeton code			
Run		posit	ion		Δz	= -5.1	l cm	Δz	= -2.9	$0\mathrm{cm}$	$\overline{\Delta z} = +2.7 \mathrm{cm}$			
	x	y	z	ρ	$ \Delta \rho $	Δz	$ \Delta \mathbf{x} _{c}$	$ \Delta \rho $	Δz	$ \Delta \mathbf{x} _{\mathrm{c}}$	$ \Delta \rho $	Δz	$ \Delta \mathbf{x} _{\mathrm{c}}$	
793	0	0	0	0	3.7	-4.9	3.7	0.3	-1.4	4.3	0.8	1.1	1.8	
791	0	0	32	0	5.6	-3.6	5.8	3.8	3.7	7.6	2.7	0.6	3.4	
797	0	48	0	48	4.4	-3.4	4.7	6.0	1.9	6.1	1.7	4.8	2.7	
798	-48	0	0	48	0.6	0.9	6.0	4.6	4.2	8.5	1.7	4.9	2.8	
799	27 -	-40	0	48	1.2	-2.1	3.2	1.2	2.7	5.7	0.9	4.9	2.4	
810	27 -	-40	11	48	1.4	-0.4	4.9	0.5	6.4	9.3	2.1	5.6	3.6	
800	40 -	-27 -	-30	48	5.2	-7.1	5.6	1.3	-2.3	1.4	1.5	4.3	2.2	
801	-40	27 -	-30	48	0.6	-3.4	1.8	4.6	-2.1	4.7	1.3	3.6	1.6	
805	0	0 -	-62	0	6.2	-6.4	6.3	2.2	-7.1	4.7	2.5	1.9	2.6	
807	-40	27 -	-56	48	2.8	-7.1	3.4	8.2	-9.1	10.3	3.3	0.6	3.9	
809	-40 -	-27 -	-56	48	2.7	-6.2	2.9	5.2	-7.6	7.0	0.6	0.6	2.2	
804	0	0 -	-75	0	6.8	-6.9	7.0	3.0	-7.6	5.6	0.6	2.5	0.6	
802	0	0 -	-78	0	9.1	-9.5	10.1	1.7	-10.2	7.5	2.7	3.0	2.7	
803	0	0 -	-83	0	8.1	-11.2	10.1	2.3	-12.7	10.1	1.5	1.7	1.8	
795	0	0	100	0	-	-	-	-	-	-	9.9 -	-0.3	10.3	

Table 7.7: "Corrected" distances between nominal and reconstructed CTF 2 source positions for software developed at Milan, Munich, and Princeton. The Princeton code uses the truncated Taylor series for a time of flight function. In this table, $|\Delta \rho|$ is the difference between the horizontal distance of the reconstructed position ρ from the z axis, and the nominal ρ of the source (either 0 or 48 cm). Δz is the difference between the reconstructed z coordinate and the nominal z coordinate. See the text for the definition of $|\Delta \mathbf{x}|_c$.

three reconstruction codes, we may define a "corrected" distance between the reconstructed and true positions of a source with $|\Delta \mathbf{x}|_{c} = \sqrt{(\Delta \rho)^{2} + (\Delta z - \overline{\Delta z})^{2}}$.

When this transformation is performed, the results are given in Table 7.7. The corrected distances of reconstructed positions from nominal source positions now range from 1.8–10.1 cm for the Milan software, 1.4–10.3 cm for the Munich software, and (excluding Run 795) 0.6– 3.9 cm for the Princeton software. In other words, the results of the Princeton software are much more consistent with the assumption of a constant z offset, and no other bias, than are the other two codes. Part of the blame for this probably lies with PMT 76 not being specifically ignored in the Milan and Munich codes. Still, the spread in Δz is much greater for the other two codes than for the Princeton software: the standard deviation for Δz is 1.9 cm for the Princeton code, 3.3 cm for the Milan code, and 6.1 cm for the Munich code. This effect appears mainly due to a bias in the other two reconstruction codes that pushes apparent source positions farther away from the origin by 10–15%. If only the six source runs having the source closest to the origin are considered, the average Δz obtained with the Munich code is +2.9 cm. This value is reassuringly consistent with the z offset $\overline{\Delta z} = +2.7$ cm seen in the Princeton software results.

7.7.3 Precision of reconstruction software with ²¹⁴Po events

From Equation (5.41), we expect a position resolution at the center of the CTF vessel of

$$\delta a = \frac{c\sigma}{n} \sqrt{\frac{3}{\epsilon} \frac{\pi(1+\delta)}{\pi+2\delta}}$$
(7.40)

where $\delta \equiv (\epsilon - N)/N$, ϵ is the total photoelectron yield of an event, and N is the number of triggered TDC channels (not PMTs), calculated with Equation (5.31). The total number of TDC channels having at least one functional PMT was T = 63 or 64 for all CTF 2 source runs. As usual, n is the index of refraction (n(PXE) = 1.565), σ is a representative time half-width of the scintillator PDF, and c is the speed of light. σ was estimated by sampling the scintillator PDF used in the reconstruction at 1 ns intervals and fitting it to a pure Gaussian curve; the result was $\sigma = 3.3$ ns. For ²¹⁴Po events in Runs 791–795, with a photoelectron yield of about 200, the expected position resolution is $\delta a = 9.0$ cm. For Runs 797–810, with a photoelectron yield near 100, it is $\delta a = 12.1$ cm.

The actual resolutions obtained with the Milan, Munich, and Princeton reconstruction codes are shown in Table 7.8. For both sources, the resolutions obtained with the Princeton software are 2–4 cm poorer than predicted. There are several possible reasons. This may be due to late-arriving photons from internal reflection. (Though not directly affecting the TDC channel values, they do increase the ADC channel values that control the peak width and time of the p_n scintillator response order statistic functions.) The finite extent of the radon source contributes a bit. It is possible that the scintillator PDF used, with a characteristic response time of 3.3 ns, is too narrow. This PDF was obtained via Monte Carlo methods, while the CTF 1 PDF (which had a response time of 5.1 ns) was derived from real CTF data. Finally, another important factor is the apparent inaccuracy in the timing calibration tables discussed in the previous section.

Resolutions obtained with the Milan and Munich codes are poorer than those of the Princeton code by an additional 1–2 cm. As expected, all reconstruction codes yield slightly worse resolutions when the source is farther from the center of the CTF vessel. The effect is not, however, a very strong function of the source radial coordinate. It is interesting to note that the resolution in y is better than the resolution in x or z for most runs with all reconstruction codes. The reasons for this are not known.

The Princeton software results in Table 7.8 are reported using the truncated Taylor series for the time of flight function. When other approximations are used, the results are nearly identical for a source at the center of the vessel (shown in Figure 7.16). However, the cubic spline interpolation in particular develops serious problems at the edge of the vessel (Figure 7.17a), while the Taylor series (Figure 7.17b) and effective index of refraction approximations still produce results with a nearly Gaussian spread.

7.7.4 Data contamination by mis-reconstructed surface events

In Figures 7.16 and 7.17, several isolated events are present in each run that lie well outside the Gaussian curves representing the radon source. These originate in one of two ways. They may be *bona fide* ²¹⁴Po events in the CTF scintillator due to radon contamination, having no relationship to the source. However, they may also be events that occurred inside the radon source but whose positions have been badly mis-reconstructed, beyond the expected Gaussian error, by the reconstruction software. The latter case is a serious concern because it would lead to events appearing to contaminate the innermost cubic meter of scintillator, though they actually originated on the surface of the vessel.



Figure 7.16: Histograms of the reconstructed z positions of ²¹⁴Po events in Run 793, with the radon source nominally at the center of the CTF vessel. Note that the vertical axes of the graphs are logarithmic. The histogram at left (a) uses the cubic spline interpolation of the time of flight function, and at right (b) is a histogram using the truncated Taylor series. Gaussian fits to the histograms are shown in bold.



Figure 7.17: Histograms of the reconstructed z positions of ²¹⁴Po events in Run 795, with the radon source nominally at the north pole (z = +100 cm) of the CTF vessel. The histogram at left (a) uses the cubic spline interpolation, and at right (b) is a histogram using the truncated Taylor series. The spline interpolation has serious problems at the edge of the vessel, while the truncated Taylor series still produces almost Gaussian results. Use of a single index of refraction produces a histogram (not shown) with a more Gaussian tail on the outside of the vessel, but does not change the size of the non-Gaussian inner tail at z < 50 cm.

Run	Non	ninal	posi	tion	M	ilan co	de	Mu	nich c	ode	Princeton code			
	x	y	z	r	σ_x	σ_y	σ_z	σ_x	σ_y	σ_z	σ_x	σ_y	σ_z	
Sour	ce 1													
793	0	0	0	0	12.6	10.7	13.6	12.2	11.5	13.3	11.3	10.8	12.4	
791	0	0	32	32	12.8	10.3	13.0	12.5	11.3	13.4	11.3	10.4	12.0	
795	0	0	100	100	-	-	-	-	-	-	20.5	18.1	15.4	
Sour	ce 2													
797	0	48	0	48	16.9	13.9	17.1	16.1	14.9	16.7	14.5	13.5	15.1	
798	-48	0	0	48	15.3	13.4	16.3	14.9	14.5	17.5	14.1	13.5	15.8	
799	27 -	-40	0	48	17.1	13.4	17.3	16.6	14.6	17.3	15.2	13.1	15.7	
810	27 -	-40	11	49	16.8	13.5	17.4	16.3	14.9	17.2	14.5	13.5	15.8	
800	40 -	-27	-30	57	17.3	13.3	18.0	16.0	14.4	17.4	14.6	13.2	15.8	
801	-40	27	-30	57	16.0	13.2	16.0	15.6	14.6	17.1	14.3	13.3	15.3	
805	0	0	-62	62	16.6	14.1	16.9	15.9	15.8	17.3	14.6	14.4	15.9	
807	-40	27	-56	74	16.5	13.9	16.3	17.0	15.9	18.8	14.8	14.0	16.0	
809	-40 -	-27	-56	74	16.8	13.7	16.6	17.7	15.5	18.5	15.0	14.2	16.0	
804	0	0	-75	75	18.4	14.7	17.0	18.0	16.7	17.8	15.6	15.2	15.3	
802	0	0	-78	78	18.7	14.5	17.1	17.3	16.4	18.2	15.0	14.4	15.0	
803	0	0	-83	83	18.4	14.1	16.1	17.6	15.8	18.7	15.2	13.9	14.2	

Table 7.8: Resolutions obtained with different reconstruction software for ²¹⁴Po events in the CTF 2 source runs. All values are the σ parameter, reported in cm, of Gaussian fits to the histograms of reconstructed x, y and z coordinates. Results for Milan and Munich codes are taken from reference [72]. The Princeton code used the truncated Taylor series for a time of flight function. The runs are arranged first by source (the first three used Source 1 and the rest used Source 2), and second in order of increasing distance r from the center of the CTF vessel.

In Run 793, seven events are clearly identifiable as being outside the Gaussian tail of ²¹⁴Po events occurring within the radon source. (All seven are visible in Figure 7.16a.) If these events are due to background, they most likely come from the CTF 2 radioactive "column" described in Section 6.4.2. This is made more likely by the observation that these events are all within 50 cm of the z-axis. Since the Gaussian peak of the histogram extends from z = -50 cm to +50 cm, statistically it should be hiding 7 ± 3 more background events.

Runs 793 and 795 each lasted about 20 hours, and so should have about the same number of background radon events due to the "column." Hence we expect to see 7 ± 3 background events having r < 50 cm. In fact the number of such events in Run 795 is ten. (Not all of the events shown in Figure 7.17b as having |z| < 50 cm also have r < 50 cm.) As a result we can report the number of mis-reconstructed events having r < 50 cm to be 3 ± 4 . This is out of a total of 2136 events, of which presumably ~ 14 are background. We conclude that the fraction of surface events that will be mis-reconstructed into the innermost 50-cm radius of scintillator is $(0.15 \pm 0.20)\%$. This is equivalent to an upper limit of 0.38% at 90% confidence level.

7.7.5 Results with ²¹⁴Bi events

In selecting candidate ²¹⁴Po α events, one automatically also selects candidate ²¹⁴Bi β decays from Group 1 events. These events have a much greater energy range than ²¹⁴Po events, whose energy spectrum would be monoenergetic if not for variance in the emission of scintillation photons and photoelectron yields of PMTs. Though ²¹⁴Bi events are less representative of energies in the neutrino window, they do give an indication of the effectiveness of the reconstruction software at higher energies.

One problem inherent in testing the reconstruction with β decays is the production of related γ rays. The isotope ²¹⁴Bi decays into ²¹⁴Po with emission of an electron and antineutrino having several different total energies. The highest combined $e^- + \bar{\nu}_e$ energy is 3.26 MeV, resulting in a ²¹⁴Po atom with its nucleus in the ground state. However, the branching ratio for an $e \bar{\nu}$ pair at this energy is only 18%. The production of an electron and antineutrino at lower energies yields an excited ²¹⁴Po nucleus, which rapidly emits one or more γ rays to make up the remainder of the 3.26 MeV energy difference. Since γ rays have an absorption length in scintillator of several cm, they are not approximated well by the assumption of a point-like event. Fortunately, as shown in Figure 7.18, the ²¹⁴Bi events are largely γ ray-free up to about 1.3 MeV (\approx 410 photoelectrons).

The results of reconstruction using the Milan and Munich codes have not been reported for ²¹⁴Bi events. With the Princeton software, the average reconstructed source positions



Figure 7.18: The total energy spectrum of the electron and γ rays emitted by a decaying ²¹⁴Bi nucleus. At top (part a) is the exact spectrum shown convoluted with the energy resolution of the CTF 2 detector, a Gaussian function having $\sigma_E(E)$ given by 2.0 keV^{1/2} \sqrt{E} . (The proportionality constant was derived from the observed ²¹⁴Po α peak energy of 201.3± 16.1 photoelectrons in Run 793, combined with the conversion factor from photoelectrons to keV of about 3.14 keV/pe [72]. In an ideal detector, the proportionality constant would be given exactly by the square root of this conversion factor, 1.78 keV^{1/2}.) The total energy spectrum is shown, as well as the component due to the 3.26 MeV pure β decay (18% of decays) and the component due to the e^- plus the pair of γ lines at 1.73 and 1.77 MeV (36%). The latter component is a β decay energy spectrum shifted to the right by the energy of the γ rays. Up to about 1.3 MeV, most of the energy spectrum results from the pure β decay component. Below (part b) is shown the actual energy spectrum observed for ²¹⁴Bi events in Run 793. The ranges of the two graphs' *x*-axes are roughly equivalent. In the second figure, events with fewer than 100 photoelectrons have been excluded.



Figure 7.19: Position resolution of the CTF 2 at the center of the vessel as a function of photoelectron yield. The black curve is the prediction of Equation (7.40). Blue points represent the reconstructed resolution in the x coordinate of the ²¹⁴Bi events of Run 793, grouped in 100-photoelectron-wide bins. The green point at 200 photoelectrons shows the resolution in x of the ²¹⁴Po events of Run 793 (Source 1), while the green point at 100 photoelectrons shows the resolution in x of the ²¹⁴Po events in Run 798 (Source 2).

with these events for each run are not very different (within 1–2 cm) from the results for the ²¹⁴Po events, so they are not tabulated here. The typical position resolution (the σ of a Gaussian fit to histograms of each reconstructed coordinate) is 10.5–12 cm for Source 1, and 11.2–12.9 cm for Source 2. This is significantly less than the position resolution for ²¹⁴Po events because the average ²¹⁴Bi photoelectron yield is much larger, and Equation (7.40) goes roughly as $\epsilon^{-1/2}$ for large ϵ .

Figure 7.19 shows the predicted dependence of the resolution upon the photoelectron yield, Equation (7.40), as well as the observed resolution in the x coordinate for ²¹⁴Bi events in Run 793 grouped into bins 100 photoelectrons wide. Points for the ²¹⁴Po events of Source 1 and Source 2 are also shown. The analogous graphs for the y and z coordinates are not presented, but they are similar. Notice that the real data reach a minimum resolution of



Figure 7.20: Histograms of the reconstructed z positions of 214 Bi events in (a) Run 793 and (b) Run 795. Both histograms use the truncated Taylor series for the time of flight. The γ ray components of these events cause a significant non-Gaussian tail of events in both runs, even though the Gaussian curves themselves are narrower than for the corresponding 214 Po events. Compare with Figures 7.16b and 7.17b.

9.7 cm in the 300–400 photoelectron bin, and beyond that the resolution worsens again. This is because a significant fraction of the ²¹⁴Bi spectrum above this energy consists of γ ray emissions. Except for a near-constant offset that worsens the resolution by about 2 cm, for which some possible reasons were discussed in the previous section, the data below 400 photoelectrons follow the predicted curve fairly well.

In addition to worsening the position resolution, γ rays also cause histograms of the reconstructed positions of ²¹⁴Bi events to have non-Gaussian tails. Figure 7.20 shows histograms of the reconstructed z positions of ²¹⁴Bi events for Runs 793 and 795. Comparison with Figures 7.16b and 7.17b shows that some events from the source are now reconstructed farther away from the nominal source position. This is a potential problem since the CTF experiences a high rate of external γ rays (hundreds per day) which we do not want to confuse with true internal events produced in the scintillator. Curiously, however, the number of ²¹⁴Bi events from Run 795 that are reconstructed within 50 cm of the origin is still only ten, just as with the ²¹⁴Po events.



Figure 7.21: The condition of the CTF 3 vessel in February 2005, one month before the CTF 3 source calibration tests. The bottom half of the vessel departs significantly from a spherical shape. This is suspected to be a significant factor in unexpected results obtained from the source runs.

7.8 Testing algorithms with CTF 3 source runs

No source calibration runs were specifically planned at the time CTF 3 was brought into operation. Although it was always intended to perform source runs eventually, the accidental spill of August 2002 caused activities with pseudocumene to cease for 2.5 years. As a result, a position calibration of CTF 3 could not be performed until March 2005.

At this point, the CTF 3 was in less than optimal condition for a source calibration. The last operation before the pseudocumene spill was a silica gel column purification test in batch mode. Some pseudocumene was lost during the test, leaving the vessel underfull. With pseudocumene having a lower density than the surrounding water buffer, the result was a vessel shaped somewhat like an upside-down teardrop or hot-air balloon. The condition of the vessel in February 2005 is shown in Figure 7.21. Nevertheless, eight source calibration runs, numbered 2532–2539, were acquired during March 17–18, 2005.



Figure 7.22: Diagram of the attachment of the CTF 3 source to the hinged metal arm. The scintillator is contained in the quartz sphere at left, and the assembly—connected via a spring—is threaded into the end of the arm. Figure courtesy of S. Kidner.

7.8.1 The source calibration hardware

The source insertion system used for CTF 3 was mainly similar to that of CTF 2. Several metal rods could be inserted into the vessel from the top of the CTF water tank. The bottom rod was hinged, and the lower arm holding the source could be raised to a horizontal orientation by pulling upward on a plastic monofilament line attached below the hinge. The distance from hinge to source in this case was 54 cm.

Unlike the source system of CTF 2, the CTF 3 source was attached to the hinged arm via a spring, as shown in Figure 7.22, returning to the design of the original CTF 1 source insertion system. This feature permitted bringing the source in contact with the nylon vessel safely. With the hinged arm extended, the source could be slowly raised, at the same time rotating the arm back and forth about the z axis at a slight angle, and stopping when the increased force required for the rotation indicated that the source was in contact with the nylon film. This was in fact done in Run 2537. The nominal position of the source in that run was (x, y, z) = (52, -13, 98) cm. This position has a radial coordinate of 112 cm, indicating that the top of the vessel was much too flat to be a perfect sphere.

The source itself was a radon source with a 214 Po event rate of 0.6 Bq. The radon-infused pseudocumene/PPO scintillator was contained in a spherical quartz vessel with a radius of about 1.5 cm. As in CTF 2, some oxygen contamination was present, quenching the average photoelectron yield of 214 Po events in the source to 126 ± 15 . As a result, if all of the 75 surviving PMTs received independent timing data, we would predict a position resolution at the vessel center of 11.0 cm. The software cuts used to determine candidate events for the results of the reconstruction were identical to those used with the CTF 2 tests, except that only Group 2 events having a photoelectron yield in the range 70–200 were accepted.

7.8.2 The CTF 3 source run results

Tabulated in Table 7.9, the results of the CTF 3 source calibration runs were less than ideal. The candidate 214 Po events were reconstructed using the Milan code (with an effective index of refraction n = 1.75) and the Princeton code. The code developed at Munich is seemingly no longer available. The Princeton results are from the truncated Taylor series time of flight; except in Run 2537, other time of flight approximations give results that are the same within a few mm. As with CTF 2 source runs, histograms of the reconstructed coordinates were fit to Gaussian curves to obtain reconstructed source positions.

In this case, for the sake of completeness, the results of reconstructing ²¹⁴Bi decays (Group 1 coincidence events) with the Princeton code are also shown. Only small (few-cm) differences are seen from the Princeton code results for ²¹⁴Po events. There is a general trend for the ²¹⁴Bi reconstructed positions to be 2–3 cm greater in x and about 1 cm smaller in y than the ²¹⁴Po reconstructed positions. It is possible that the Group 1 and Group 2 electronics systems have different systematic biases that cause these differences.

Rı	un	214 Po		Nomi	nal posi	tion		Mila	n code	$(^{214}Po $	α 's)		\Pr	inceton	code	
		events	x	y	z	ρ	r	x	y	z	$ \Delta \mathbf{x} $		x	y	z	$ \Delta \mathbf{x} $
25	532	19718	0	0	82.0	0	82	-1.8	4.6	80.0	5.3	α :	-1.3	4.7	84.8	5.6
												β :	0.7	3.9	84.8	4.8
25	533	9647	0	0	49.3	0	49	1.2	4.7	47.6	5.1	α :	-1.1	4.6	49.4	4.7
												β :	0.9	4.1	49.6	4.2
25	534	4174	0	0	-53.5	0	54	2.2	4.8	-57.0	6.3	α :	-2.4	5.0	-62.1	10.2
												β :	0.8	4.3	-61.5	9.1
25	35	4235	0	0	-80.0	0	80	-2.1	-2.8	-84.7	5.9	α :	-1.9	3.1	-91.2	11.7
												β :	0.9	2.4	-90.3	10.6
25	536	8298	54	0	21.5	54	58	49.0	5.0	18.0	7.9	α :	53.6	4.8	18.0	6.0
												β :	55.1	4.9	18.0	6.1
25	537	15490	52	-13	97.8	54	112	48.3	-3.2	92.0	11.9	α :	51.3	-3.5	95.4	9.8
												β :	53.6	-4.6	91.9	10.3
25	38	4552	-54	0	21.5	54	58	-54.0	-1.0	16.0	5.6	α :	-58.2	-1.3	15.9	7.1
												β :	-54.4	-2.5	15.4	6.6
25	39	7302	0	0	0.0	0	0	-2.5	6.3	-1.7	7.0	α :	-2.8	6.5	-2.0	7.4
												β :	-0.1	5.4	-2.4	5.9

Table 7.9: CTF 3 source runs. This table gives the total number of coincidences meeting the ²¹⁴BiPo cuts, the nominal position of the source, and the mean positions obtained by reconstruction with the Milan software (Section 7.1.1) and the Princeton software. The effective index of refraction used in the Milan code was n = 1.75 [164]. Only ²¹⁴Po α events were analyzed with the Milan software. Both ²¹⁴Po α and ²¹⁴Bi β events (first and second lines for each run, respectively) were analyzed with the Princeton software. All positions are given in cm, relative to the nominal center of the CTF vessel. $|\Delta \mathbf{x}|$ is the distance between reconstructed position and nominal position.

The results of runs 2532, 2533, and 2536–2539, with the source in the upper hemisphere of the CTF vessel, are not too far from the nominal source positions, aside from an odd offset in the positive y direction seen with both reconstruction codes. Perhaps, as in the CTF 2 source runs, the offset is due to a misfiring TDC channel or an error in the TDC calibration tables. When the raw data were inspected, however, no electronic channel was obviously misbehaving, the TDC channel 15 (or its calibration table entry) having settled down since the days of CTF 2. It is also quite possible that this offset is real. Since the inner diameter of the north end pipe in CTF 3 is substantially larger than in CTF 2, it was difficult to insert the set of rods in an exactly vertical orientation. In runs 2536–2539, both reconstruction codes also show a small ~ 3 cm offset in the negative z direction, which could also easily be a real effect.

In runs 2534 and 2535, located on the z axis below the center of the CTF, both reconstruction codes (the Princeton code especially) produce disappointing mismatches with the nominal source positions. These errors are thought to result from the non-spherical deformities in the CTF vessel. With the vessel shaped like an upside-down teardrop, light produced by an event has to travel through less scintillator and more water to reach PMTs in the bottom half of the PMT superstructure than would be the case for a spherical vessel. Hence the lower PMTs receive scintillation light earlier than one would expect, making the event be reconstructed closer to them than it actually was. This error will be most pronounced when the lower PMTs are the first ones hit, *i. e.*, for events in the bottom half of the vessel. The effect can be easily seen in Figure 7.23, a plot of the reconstructed z positions found by both reconstruction codes as a function of the nominal z positions of the source.

We can estimate the size of the maximum vessel deformity Δr by noting that the reconstruction error Δz is 11.2 cm for Run 2535. To first order, the relation between these quantities is given by

$$\Delta r \approx \frac{n_{\rm PC}}{n_{\rm PC} - n_{\rm H_2O}} \,\Delta z. \tag{7.41}$$

The result is a 45-cm deformity. Such a large vessel deformation is not supported by



Figure 7.23: Scatter plot of the reconstructed z positions of ²¹⁴Po events in the CTF 3 source runs as functions of the nominal source positions in z. The Princeton code results are shown as the black stars, and the Milan code results as the green stars. The line $z_{\text{recon}} = z_{\text{nominal}}$ is also shown for comparison.

Figure 7.21, from which the size of the deformity may be estimated at ~ 30 cm. However, changes in the shape of the vessel affect the angle of refraction at the vessel surface, which also has some effect on photon times of flight. Additionally, the equation above makes the simplifying assumption that the PMTs are collinear with the true and reconstructed event positions (nearly on the z axis) rather than off at an angle as in reality.

7.8.3 Position resolution for CTF 3 source runs

The position resolutions obtained using the Princeton reconstruction software, using the truncated Taylor series for the time of flight function, are shown in Table 7.10 for both 214 Po and 214 Bi events. As with the CTF 2 source runs, the resolutions are a couple of cm poorer than predicted. Except in Runs 2535 and 2537, they are essentially independent of the source position, although (similar to the CTF 2 runs) the resolution is anisotropic, being consistently better in y than in x or z.

Run	1	Nomi	nal po	sition	1	214	Po eve	nts	²¹⁴ Bi events			
	x	y	z	ρ	r	σ_x	σ_y	σ_z	σ_x	σ_y	σ_z	
2532	0	0	82.0	0	82	13.1	11.7	14.3	12.0	11.0	12.8	
2533	0	0	49.3	0	49	13.0	11.6	14.2	11.9	11.0	12.3	
2534	0	0 ·	-53.5	0	54	13.2	11.7	14.2	12.1	11.3	12.7	
2535	0	0 -	-80.0	0	80	14.4	12.8	13.1	12.1	11.5	12.0	
2536	54	0	21.5	54	58	12.9	11.7	13.9	12.2	11.2	12.2	
2537	$52 \cdot$	-13	97.8	54	112	18.8	17.0	16.4	15.1	14.0	13.4	
2538	-54	0	21.5	54	58	13.2	11.8	14.0	12.3	11.2	11.9	
2539	0	0	0.0	0	0	13.0	11.7	14.2	11.8	11.2	12.6	

Table 7.10: Resolutions obtained with the Princeton reconstruction software for ²¹⁴Po and ²¹⁴Bi events in CTF 3 source runs. All values are the σ parameter, reported in cm, of Gaussian fits to the histograms of reconstructed x, y, and z coordinates. The time of flight function used was the truncated Taylor series (other time of flight approximations gave very similar results).

A study of possibly mis-reconstructed events analogous to that performed with the CTF 2 source runs benefits from a much higher source rate and a much lower ²¹⁴BiPo background rate in the CTF 3. In Run 2539, nominally located at the center of the vessel, only five ²¹⁴Po events out of 7302 are reconstructed outside the Gaussian curve surrounding the radon source (Figure 7.24a). Compare this with the seven out of 3136 ²¹⁴Po events reconstructed outside the Gaussian curve in Run 793 of the CTF 2 source runs. In the 20-hour long Run 793, most of these events were probably due to background from radon in the scintillator; in Run 2539, with only 2.1 hours of live time, it is likely that most of them really were mis-reconstructed events from the radon source.

Considering Run 2537 (Figure 7.25a), at the surface of the vessel, we find that only two ²¹⁴Po events are within 50 cm of the nominal vessel center, and only eight are within 60 cm of the center. This is consistent with the apparent rate of event mis-reconstruction in Run 2539. In the worst case, if all of these events were produced in the source and mis-reconstructed, the fraction of point-like surface events seen inside a 60 cm radius will be 5×10^{-4} , or 0.05%.

When we instead look at 214 Bi events (Figures 7.24b and 7.25b), again the non-Gaussian tails seen in the CTF 2 source runs are present. In Run 2537, there are now 29 214 Bi events



Figure 7.24: Histograms of the reconstructed z positions of events in Run 2539, with the radon source nominally at the center of the CTF 3 vessel. Note that the vertical axes of the graphs are logarithmic. The histogram at left (a) is of 214 Po events, and that at right (b) is of 214 Bi events. Both histograms used the truncated Taylor series for the time of flight. Gaussian fits to the histograms are shown in bold.



Figure 7.25: Histograms of the reconstructed z positions of events in Run 2537, with the radon source touching the CTF vessel near the top. The histogram at left (a) is of 214 Po events, and at right (b) is of 214 Bi events. In both runs, the 214 Bi event histograms have narrower Gaussian fits, but also have non-Gaussian tails due to the partial γ ray character of the events.

reconstructed to be within 50 cm of the origin, and 83 events within 60 cm. The fraction of γ -like surface events that are mis-reconstructed to lie within a 60-cm volume is therefore at least 0.5%; probably greater, since a reasonable fraction of ²¹⁴Bi events do not emit any γ rays. The Monte Carlo analysis described in Section 9.4.1 suggests a figure more on the order of 2%.

Chapter 8

Internal Contamination in the CTF

As already mentioned in Chapter 6, the main purposes of the Counting Test Facility are to provide a proof-of-concept for the Borexino experiment, to test materials and cleaning methods potentially to be used in Borexino, and to ensure that the radiopurity of the pseudocumene + PPO scintillator is sufficient for Borexino (or at least, at levels beyond the sensitivity of the CTF itself). There is one primary source of information for all three of these goals: the recorded photomultiplier tube hits, which are grouped by electronics and reconstruction algorithms into single events. Each such event is presumed to represent the decay of one radioactive atom in the CTF or the passage of one muon through it. Neutrinorelated events are also in principle observed, but they are impossible to distinguish from the much more common radioactive background noise. This "background noise," though most undesirable in Borexino, represents the signal we are interested in studying in the CTF.

The radioisotopes that are most common in CTF are expected to be the same as those in Borexino, since the scintillator to be used in both detectors is the same, and the methods of construction and purification are similar. Specifically, they fall into the categories of ¹⁴C intrinsic to the hydrocarbons of the scintillator; cosmogenic isotopes such as ¹¹C and ⁷Be produced by muons passing through the detector; common radioactive isotopes of noble gases (specifically ³⁹Ar, ⁸⁵Kr and ²²²Rn) that may be present in the detector's water tank or nitrogen blanket; the heavy element decay chains originating with ²³²Th and ²³⁸U; and other "miscellaneous" naturally occurring radioactive isotopes such as ⁴⁰K and ⁸⁷Rb. Tables 2.2, 2.3, and 2.4 in Section 2.2 provide more details.

Events in the CTF may also be divided into three categories based on their positions in the detector. The reconstructed distributions of these positions were discussed earlier, in Section 5.5. Internal events are those that result from impurities or particulates in the scintillator itself, from cosmogenic isotopes and neutrons produced by cosmic rays in the scintillator, and from noble gases diffusing into the scintillator from outside (this category also includes the decay products of ²²²Rn). Neutrino events would also be considered internal events if they were feasible to observe. Except perhaps for the noble gases, the distribution of these events would ideally be uniform throughout the CTF scintillator volume. Due to time and space constraints, only the heavy isotopes of the uranium and thorium decay chains will be discussed in this chapter in detail. Other, lighter radioisotopes in the CTF scintillator will be mentioned briefly at the end.

Surface events, to be discussed in the following chapter, are a result of impurities in the nylon film of the CTF vessel, as well as impurities and particulates originally in the scintillator that eventually adhere to the vessel film. The distribution of these events should in principle be restricted to the vessel surface, idealized as a sphere of radius 1 m. External events, also discussed in the next chapter, are those produced by radioactive decays of isotopes outside the scintillator, for instance in the water buffer or the PMTs. These events are visible only because γ rays produced by them may travel some centimeters, into the scintillator volume. The volume of scintillator in CTF is much smaller than in Borexino, so it has a much smaller volume to surface ratio. Consequently, surface and external events are much more prominent in CTF than they will be in Borexino. (This is not to say that they will not be a problem in Borexino, however!) In studying CTF data, we must ensure that our samples of internal events are not contaminated much by surface or external events.

For the purposes of this chapter and the next, we will focus mainly on the later history of the CTF 3, after the most recent of the purification tests in June 2002. The largely disappointing

results of the tests themselves have been re-analyzed and discussed *ad nauseam* [42, 44, 67, 92], and it seems unlikely that many novel things can be said about them. (An exception to this policy is made in the case of the ²¹⁰Po activity, regarding which some new results will be presented in Section 8.3.) It is the opinion of the author that the most interesting part of the CTF 3 history, or at least the part most amenable to analysis, is that following the pseudocumene spill of August 2002, after which the virtual shutdown of operations provided an extensive data acquisition period on undisturbed scintillator.

For all the analyses presented here, an arbitrary cutoff date of June 22, 2005 (Run 2563) was set. Since CTF data acquisition is ongoing, attempting to keep the present work up-to-date with the most recent data until the moment of its publication would be a futile attempt at chasing a moving target.

8.1 Particle identification techniques

Particle identification, as applicable to radioactive background, encompasses a variety of techniques by which an individual scintillation event may be identified with confidence as being produced by a specific type of decay. The most important of these techniques is the use of coincidence events. Other event parameters may also provide a great deal of information; for instance, α/β discrimination, particle energy, and event position. In Borexino, when an event is positively identifiable as a certain radioactive decay, it may be tagged and excluded from the data set, improving the neutrino signal-to-noise ratio.

Much of the time it may be impossible to say that a particular event is definitively due to the decay of a certain isotope. However, it may still be possible to determine the concentration of a specific isotope statistically, using the overall observed energy spectrum in the CTF, the α/β parameters, the spatial distribution of events, and the presence of positively identifiable isotopes in the same radioactive decay chain. In this case, the energy spectrum of the isotope may be statistically subtracted from the overall energy spectrum, again improving

the neutrino signal-to-noise ratio. Unfortunately, in this case, the statistical uncertainties in the neutrino data set will be greater.

A thorough discussion of particle identication methods in Borexino is available in reference [68]. This section will focus on particle identification in CTF 3.

8.1.1 Event position

The reconstructed position of an event does not directly help with particle identification. However, it does identify the most likely source of the radioactive contamination. One of the most interesting uses of the CTF is as a radiopurity testing facility for the materials intended to be used in Borexino. The component having the most stringent radiopurity requirement is the scintillator. In this chapter, we therefore focus upon contamination specific to the scintillator. In order to exclude external and surface events from the data sample currently of interest, radial cuts can be used—analysis is only performed on events whose reconstructed radial coordinate is less than some set value. (In certain special cases, other three-dimensional shapes may also be used to define cuts on the detector volume.)

An additional known problem with events near the surface of the vessel is that of light loss. Surface events are empirically observed to suffer a reduction in light yield of about 15% relative to those at the center of the detector; this effect will be described a bit more in the next chapter. The cause of this reduction may be optical effects related to the nylon film and the different indices of refraction at the scintillator-film-water interface, but it is not completely understood. The light loss is a fairly strong function of the radial coordinate, so it may be neglected (simplifying analysis considerably) by using a radial cut.

Two natural radial cuts are those which exclude events with reconstructed radial coordinates outside 50 cm and 65 cm, respectively. The first cut represents a volume of $\frac{1}{8}$ of the total amount of CTF scintillator, and the second cut represents a scintillator mass of 1000 kg. The first is useful in analysis of event samples that may be heavily contaminated from external

 γ rays, while the second may be used for samples that are not so sensitive (coincidence analyses, for instance). A look at Figure 5.4 shows that the resolution error for these volumes (due to assuming that the number of events reconstructed within these radii is equal to the true number of events occurring within them) is negligible for a reconstruction position resolution of 10–15 cm.

We must however be careful, due to the fact that CTF 3 events in the lower hemisphere seem to be reconstructed lower along the z axis than their true positions. To review the relevant results of Section 7.8, ²¹⁴Po events in the radon source in run 2534 (nominally on the z axis at z = -53.5 cm) were reconstructed at an average height of z = -62.1 cm, and in run 2535 (nominally at z = -80.0 cm) they were reconstructed at z = -91.2 cm. By linear interpolation, we expect a reconstructed position of z = -50 cm to correspond to a true position of z = -42.5 cm, and -65 cm to a true position of -56.1 cm. The physical volume in the CTF to which the 50-cm cut on reconstructed radial coordinates corresponds to a physical region squashed by a scale factor of 1.081. The 65-cm radial cut corresponds to a physical region squashed by a scale factor of 1.073. All analyses that obtain a total number of events (or a number of events per unit volume or unit mass) using a radial cut must be corrected by multiplying the result with one of these scale factors. These factors may also be calculated using the measured ²¹⁴Bi event positions; the results are each smaller by about 0.5%, a negligible difference.

8.1.2 Event energy

Many radioactive isotopes have distinctive energy spectra. This is especially true for isotopes which decay via emission of a monoenergetic α particle. But to estimate an event energy, one first needs a conversion factor between the total number of photoelectrons observed by the PMTs and the actual energy of the radioactive decay. This factor, the "light yield," is provided by observation of the shape of the ¹⁴C energy spectrum in CTF, which will be described later.

Decay	Species	Half-life	E_{α}	E_{γ}	Pred.	E_{quen}	$_{\rm ched}$ [keV]	$\sigma_E \; [\text{keV}]$
chain		$ au_{1/2}$	[MeV]	[MeV]	Q(E)	Pred.	Obs.	Pred.
$^{235}\mathrm{U}$	223 Ra	11.4 d	5.717	0.159	12.87	603	-	65
	219 Rn	$3.96\mathrm{s}$	6.819	-	11.44	596	-	64
	215 Po	$1.78\mathrm{ms}$	7.386	-	10.70	690	-	69
	211 Bi	$2.15\mathrm{m}$	6.623	-	11.69	567	-	63
232 Th	220 Rn	$55.6\mathrm{s}$	6.288	-	12.12	519	-	60
	216 Po	$145\mathrm{ms}$	6.779	-	11.49	590	-	64
	^{212}Bi	$60.6\mathrm{m}$	6.051	0.040	12.43	527	-	61
	212 Po	$299\mathrm{ns}$	8.784	-	8.88	989	-	83
^{238}U	222 Rn	$3.82\mathrm{d}$	5.490	-	13.16	417	410 ± 6	54
	218 Po	$3.05\mathrm{m}$	6.002	-	12.50	480	483 ± 6	58
	214 Po	$164\mu{ m s}$	7.687	-	10.31	746	751 ± 7	72
	210 Po	$138.4\mathrm{d}$	5.305	-	13.40	396	395 ± 10	53

Table 8.1: Energy quenching and Gaussian peak widths predicted in the CTF 3 and Borexino scintillator for the most important α -decay isotopes in the three heavy-element decay chains. For ²²²Rn and its daughters, the quenched energies actually observed in the CTF 1 are also listed [128]. For species whose main decay mode includes a γ ray, the energy of the γ is presumed not to be significantly quenched. The α decay listed for ²¹²Bi has a 36% branching ratio; its main decay channel is β^- emission. Notice that all of these species but ²¹²Po have a quenched energy in the neutrino window, 250–800 keV.

As discussed in Section 3.1, the amount of light emitted by the CTF scintillator following an α decay is "quenched" with respect to an electron or γ ray interaction of equal energy. The amount of α particle quenching is a function of the specific scintillator used and any impurities in the scintillator (such as oxygen). the original α particle energy. The observed energy of an α particle is reduced by the quenching factor Q(E), which is itself a function of the true energy. An empirical formula for the quenching factor in the Borexino and CTF scintillators is given by [148, 165]

$$Q(E) = 20.3 - (1.3 \,\mathrm{MeV^{-1}})E = 20.3 - \frac{E}{0.77 \,\mathrm{MeV}}.$$
 (8.1)

The observed energy is then $E_{obs} = E/Q(E)$. α quenching factors and observed energies for the most important α -emitting isotopes in CTF 3 are tabulated in Table 8.1.

Since the CTF has a finite energy resolution, the observed spectrum of an α emitter is not actually monoenergetic. The energy of an event is determined, as described in Section 6.3.2,

by the number of photoelectrons that the PMTs of the detector see. For a monoenergetic event, this value will be a discrete random variable with Poisson statistics. When the number of photoelectrons is reasonably large, its distribution in an ideal CTF-like detector would approximate a Gaussian curve whose width is proportional to the square root of the mean observed energy, $\sigma_E = k\sqrt{E_{obs}}$. (Equivalently, the energy resolution σ_E/E would be proportional to $1/\sqrt{E_{obs}}$.) The proportionality constant k would be equal to the square root of the conversion factor from the number of photoelectrons observed to the event energy. For CTF 3, this conversion factor is approximately 3.90 keV/photoelectron in runs with 75 active PMTs (those numbered 2392–2552). We would therefore predict that, if CTF were an ideal detector, $k = 1.97 \text{ keV}^{1/2}$.

Of course, CTF is not ideal; for instance, its ADC electronics do not have perfect accuracy. Nevertheless, the α energy peaks are still empirically observed to be modeled well by Gaussian functions, as will be seen many times in this chapter. The main effect is a broadening of the observed peak widths in the energy spectrum. If we suppose that σ_E is still proportional to $\sqrt{E_{\text{obs}}}$, an estimate for the real value of k may be obtained by using the CTF 3 source run with the source at the center of the vessel, Run 2539. Selecting only ²¹⁴Po events using the method of coincidences, we find that a histogram of the event photoelectron yield spectrum (Figure 8.4b) is a Gaussian curve with mean value 126 photoelectrons and $\sigma = 15.0$ photoelectrons. Converting these values into energies with the conversion factor 3.90 keV/photoelectron and solving for k yields $k = 2.64 \text{ keV}^{1/2}$. The predicted values of σ_E using this value of k are also tabulated in Table 8.1.

In the remainder of this dissertation, the quenched energies $E_{\rm obs}$ of α decays will generally be treated as if they were the real energies of the decays, and the subscript "_{obs}" will be dropped.

8.1.3 α/β discrimination

It is very useful to be able to tell decays via α emission apart from β decays, since all α emitters have a quenched energy within the neutrino window, but events due to neutrino scattering have β -like characteristics. Even leaving aside the question of neutrinos, α/β discrimination is an important part of particle identification in the CTF as well. Three main approaches have been taken to identify radioactive decays in the CTF as α or β decays. The simplest is the tail-to-total ratio. The other two methods, the Gatti filter and the DPSA skew, both use the values acquired by the DPSA board [81].

One possible measure of the ability of a parameter to discriminate between α and β events is the factor of merit, defined as

$$D \equiv \frac{\Delta S}{\sqrt{\sigma_{\alpha}^2 + \sigma_{\beta}^2}},\tag{8.2}$$

where ΔS is the absolute value of the difference between the mean values of the α and β distributions, and $\sigma_{\alpha,\beta}$ are the standard deviations of the two distributions. When two distributions are cleanly separated, ΔS is much greater than the width of either, and D is large. Conversely, if the distributions overlap significantly, at least one of them has a width comparable to the distance between them, so D is small (≤ 3).

One may also think about the problem in terms of the classification efficiency. Suppose that, for some α/β discrimination parameter χ , the mean value of the α distribution is χ_{α} and that of the β distribution is χ_{β} (that is, $\Delta S = |\chi_{\alpha} - \chi_{\beta}|$). Without loss of generality, assume $\chi_{\alpha} < \chi_{\beta}$. Now consider a fixed value χ_0 which will be used to classify an event as either α -like or β -like depending on whether the event's value for χ is less than or greater than χ_0 , respectively. The classification efficiency $\epsilon_{\alpha}(\chi_0)$ for α 's is the fraction of α events that are correctly classified at a given χ_0 , and likewise for $\epsilon_{\beta}(\chi_0)$.

As ϵ_{α} and ϵ_{β} are both monotonic functions of χ_0 , we may think of one variable as a function of the other, $\epsilon_{\beta}(\epsilon_{\alpha})$. Naturally, they have some negative correlation; if there is any overlap between the distributions, as ϵ_{β} increases, ϵ_{α} must decrease and vice versa. The better the α/β separation is, the greater will be the value of ϵ_{α} at which $\epsilon_{\beta} = \epsilon_{\alpha}$. If one requires a very pure sample of β -like events, for instance neutrino events in Borexino, one can require a higher value of ϵ_{α} (in order to reject more α events) at the cost of having a lower value for ϵ_{β} (a larger number of β -like events are mistakenly rejected).

If we assume that the distributions of χ for α and β events are Gaussian with means $\chi_{\alpha}, \chi_{\beta}$ and standard deviations $\sigma_{\alpha}, \sigma_{\beta}$, then the value of χ_0 at which $\epsilon_{\alpha} = \epsilon_{\beta}$ is given by

$$\chi_0^{\rm eq} = \frac{\chi_\alpha \sigma_\beta + \chi_\beta \sigma_\alpha}{\sigma_\alpha + \sigma_\beta}.$$
(8.3)

At $\chi_0 = \chi_0^{\text{eq}}$, equal fractions of α and β events will be classified correctly. This fraction, the value of ϵ_{α} (equivalently, ϵ_{β}) itself, is given by

$$\epsilon_{\alpha}(\chi_{0}^{\text{eq}}) = \epsilon_{\beta}(\chi_{0}^{\text{eq}}) = \frac{1}{2} \left[1 + \operatorname{erf}\left(\frac{\Delta S}{\sqrt{2}\left(\sigma_{\alpha} + \sigma_{\beta}\right)}\right) \right].$$
(8.4)

This measure of the efficiency of discrimination has a lower bound of $\frac{1}{2}$ (for the case of no separation, where the distributions have the same mean value), and an upper bound of 1 (when the separation is perfect). We will refer to it henceforth simply as the discrimination efficiency, $\epsilon_{\alpha\beta}$.

Finally, we may define a contamination fraction C. For a given value of χ_0 , what is the fraction of events classified as being β -like that are really α 's? Suppose that the event sample contains $N_{\alpha} \alpha$ events and $N_{\beta} \beta$ events. The number of β events having $\chi > \chi_0$ is $N_{\beta}\epsilon_{\beta}(\chi_0)$, and the number of α events with $\chi > \chi_0$ is $N_{\alpha}[1 - \epsilon_{\alpha}(\chi_0)]$. The contamination fraction is then

$$C(\chi_0) = \frac{N_\alpha [1 - \epsilon_\alpha(\chi_0)]}{N_\beta \epsilon_\beta(\chi_0) + N_\alpha [1 - \epsilon_\alpha(\chi_0)]}.$$
(8.5)

In the simplified case where $\epsilon_{\alpha} = \epsilon_{\beta}$ and $N_{\alpha} = N_{\beta}$, we have $C(\chi_0^{\text{eq}}) = 1 - \epsilon_{\alpha\beta}$.

We now go into more detail in describing the three α/β discrimination parameters commonly used in the CTF: the tail-to-total ratio, Gatti parameter, and DPSA board skewness parameter.

The tail-to-total ratio R32

Recall from Section 3.1.2 that the scintillation light that results from emission of an α particle has a larger slow-decay component. Let the collected charge during a time period dt be q(t) dt. Then the tail-to-total ratio R32, defined as the ratio of the pedestal-subtracted electronics channels ADC 3032 and ADC 3000, is simply the value

$$\frac{\int_{32\,\mathrm{n}s}^{532\,\mathrm{n}s} q(t)\,\mathrm{d}t}{\int_{0\,\mathrm{n}s}^{500\,\mathrm{n}s} q(t)\,\mathrm{d}t}.$$

Because of its larger slow-decay component, an α decay will tend to have a greater value of R32 than a β decay. (The parameters R16 and R48 are defined analogously but in practice are not as useful. There is an optimum starting time for collecting the "tail" signal that maximizes the α/β discrimination efficiency [83].) In the laboratory, the α/β separation obtained between ²¹⁴Bi and ²¹⁴Po has a factor of merit D = 4.60 [81] and a discrimination efficiency $\epsilon_{\alpha\beta} = 99.956\%$. The two distributions are seen in Figure 8.1a to be cleanly separated.

Unfortunately, α/β separation using this parameter works much better in the laboratory than in the actual CTF. Consider the events in CTF 3 source run 2539 that have an apparent event energy between 400 and 800 keV. This sample includes both Group 1 and Group 2 events in the ²²²Rn decay chain. However, the extra quenching that resulted from oxygen contamination of the source caused α decays other than that of ²¹⁴Po to appear at energies lower than 400 keV. In practice, therefore, the sample includes only ²¹⁴Po α decays and some of the ²¹⁴Pb and ²¹⁴Bi β decays. For this sample of events, the factor of merit for the α/β discrimination using the R32 parameter is D = 1.63. The discrimination efficiency is $\epsilon_{\alpha\beta} = 88.7\%$. Significant overlap between the two distributions is apparent in Figure 8.1b. As an additional point of comparison, the discrimination efficiency for ²¹⁴BiPo coincidences in the Borexino source runs¹ using the tail-to-total ratio is between 84% and 95%, although

¹As the Borexino detector was not yet filled with scintillator at the time of these source runs, the source (consisting of radon gas dissolved in scintillator, enclosed in a small spherical quartz vial) was suspended in the empty (air-filled) volume of the detector.



Figure 8.1: α/β discrimination with the R32 parameter. The top graph (part a) is the distribution of R32 in the laboratory, and clearly shows the separation between α and β events. Taken with slight modifications from reference [81]. The bottom graph (part b) is the distribution of R32 seen in CTF 3 Run 2539, for events with apparent energies between 400 and 800 keV.
not enough information is provided to determine it precisely [68]. Some reasons for the poorer performance of α/β discrimination in the CTF and Borexino detectors as compared to the small-scale lab experiments include light scattering, lower light collection efficiency, and the time jitter of multiple PMTs. All of these effects tend to smear out the pulse shapes of events, pushing the α and β pulses closer together with any α/β discrimination parameter.

At lower energies, the α/β separation is poorer. For the events in CTF 3 runs 2300–2349, the main α -producing species present is ²¹⁰Po (which will be discussed further in Section 8.3). It has a predicted quenched energy of 396 keV. Using the R32 parameter to perform α/β separation on a sample of events with apparent energies in the range 250–500 keV yields a factor of merit of only D = 0.95. The discrimination efficiency is only $\epsilon_{\alpha\beta} = 75.3\%$. (Some of this loss of efficiency compared to the Rn source runs is also due to the limited volume of the radon source compared to the much larger radial cut over which ²¹⁰Po events were sampled.) In this region of the energy spectrum, the sum of the two distributions does not even exhibit a local minimum or shoulder; only a single peak is visible. Unless an event has a value of R32 less than 0.25 or greater than 0.4, the R32 parameter provides little information in this energy range.

The Gatti parameter

The Gatti parameter is the characteristic value yielded for each event by the Gatti optimum filter, a mathematical technique that parametrizes how similar a given function is to one of two reference functions [153]. In the case of the CTF, the given function is the set of values recorded in the DPSA buffer of an event. Recall that the value stored in the n^{th} element of the DPSA buffer is proportional to the total photoelectron charge yielded by an event between time zero and time $t_n = 8.3 \text{ ns} \times n$. The DPSA buffer holds 150 values, but as the width of the charge collection gate is 500 ns, only about 45 of these are meaningful. The remainder may be used as measurements of the pedestal values for the DPSA buffer. We define f_n to be the value stored in the n^{th} bin of the DPSA buffer, minus the average pedestal value, normalized such that the sum $\sum_n f_n$ is one. The pedestal value is determined from the original contents of the first 14 bins. The sum is taken only over the bin numbers n that hold a meaningful value: $n = 15, \ldots, 60$. When n is outside this range we define $f_n = 0$.

The reference functions in question are the averages of the DPSA buffer values for events known to be α and β decays. That is, the value α_n is the average (pedestal-subtracted, normalized) value in the n^{th} bin of the DPSA buffer for known α events, and likewise for β_n . Obtaining a sample of β events is easy; one takes advantage of the ²¹⁴BiPo coincidences and uses only the first event in each coincidence. Since the DPSA board processes only Group 1 events, though, ²¹⁴Po cannot be used for the sample of α events. (A second DPSA board that processes Group 2 events exists, but ensuring that the two boards had compatible calibrations would be difficult.) In this case, one uses the α decays produced by ²²²Rn and ²¹⁸Po in runs with large amounts of radon.

With these quantities defined, the Gatti parameter $G_{\alpha\beta}$ is then

$$G_{\alpha\beta} = \sum_{n} \frac{\alpha_i - \beta_i}{\alpha_i + \beta_i} f_n.$$
(8.6)

It can be shown that the Gatti parameter has the maximum discrimination capability of any parameter defined as a linear combination of the values f_n . Indeed, in the laboratory the factor of merit for the Gatti parameter is D = 5.08, compared to 4.60 for the tail-to-total ratio used on the same set of data [81]. The discrimination efficiency in the lab is $\epsilon_{\alpha\beta} =$ 99.987% (compared to 99.956%). In the Borexino detector source runs, the discrimination efficiency for ²¹⁴BiPo coincidences using the Gatti parameter with the Borexino DAQ system is roughly 95.5% [68].

In the CTF, however, interpreting the Gatti parameter is somewhat difficult. The reason is that the DPSA electronics seem to exhibit varying time offsets over long time periods. As the Gatti parameter depends delicately upon the exact time binning of the DPSA board, it is hard to consistently analyze sets of data that span a period of several months. For this reason, the Gatti parameter was not used in this work.

The DPSA skewness

Fortunately, another α/β discrimination method is available that uses the DPSA board but is not as sensitive to shifts in the board's time offset. O. Smirnov has suggested the use of the skewness of the DPSA buffer as a discrimination parameter. Skewness is a measure of the asymmetry of a peaked function—a function with a larger tail for values less than the mean has negative skewness, otherwise it has positive skewness. The skewness is independent of any shift in the position of the function peak.

To define the skewness, the normalized values f_n of the DPSA buffer for an event are considered as a probability distribution. (As before, f_n is considered to be zero when nis outside the useful range of about 15–60). The expectation value of the bin number for this distribution is given by $\bar{n} = \sum_n n f_n$. The second and third central moments of the distribution are obtained as follows:

$$\mu_2 = \sum_n (n - \bar{n})^2 f_n$$

$$\mu_3 = \sum_n (n - \bar{n})^3 f_n$$

Finally, the skewness of the distribution is defined as

$$\gamma_1 = \frac{\mu_3}{\mu_2^{3/2}}.$$
(8.7)

When a scatter plot is made of CTF data with event energy on one axis and the DPSA skewness on the other, it can be seen that there exists a slight negative correlation, apparently linear, between the two parameters for β events (Figure 8.2). In order to correct



Figure 8.2: Scatter plots of the DPSA skewness (top), γ_1 , and the DPSA skewness with energy correlation removed (bottom), γ_1^* , against the apparent event energy for Group 1 events in the CTF 3 source run 2539. The β events form a straight line across the tops of the plots. The lower-energy α decays of ²²²Rn and ²¹⁸Po are at center left. (All α events in this run appear to be shifted to lower energies; α quenching is increased due to oxygen contamination in the source vial.) The α decays of ²¹⁴Po are not shown since Group 2 events are analyzed by a different DPSA board.



Figure 8.3: α/β discrimination with the parameter γ_1^* , the skewness of the DPSA buffer with energy correlation removed. The graph shows the distribution of γ_1^* for CTF 3 runs 2300–2349, for events with apparent energies in the range 250–500 keV. The performance of the α/β discrimination is not particularly good; here D = 1.12 and $\epsilon_{\alpha\beta} = 80.3\%$; but it is better than that of the R32 parameter in this energy range (D = 0.95, $\epsilon_{\alpha\beta} = 75.3\%$).

for it, we define a linear combination of the skewness and energy which has the correlation empirically removed:

$$\gamma_1^* \equiv \gamma_1 + \frac{E}{26.3 \,\mathrm{MeV}}.\tag{8.8}$$

The distribution of γ_1^* for β events in the neutrino energy window can then be modeled (again, determined in a purely empirical way) by a Gaussian with average value $\gamma_1^* = 0.026$ and $\sigma_{\gamma} = 0.011$.

For events with apparent energies in the range 250–500 keV in CTF 3 runs 2300–2349, the performance of the α/β discrimination is not particularly good: D = 1.12 and $\epsilon_{\alpha\beta} = 80.3\%$ (Figure 8.3). This is, however, still better than the performance of the R32 parameter with the same set of events (D = 0.95, $\epsilon_{\alpha\beta} = 75.3\%$). Although the γ_1^* parameter is probably more effective at higher energies, this is difficult to test. The higher-energy ²¹⁴Po and ²¹²Po α decays always appear as Group 2 events, which are not analyzed by the same DPSA board.

It is a pleasant surprise that the effectiveness of α/β discrimination with the γ_1^* parameter does not seem to be affected by the position of an event. In Run 2539, the radon source was located at the center of the vessel. In Run 2537, the source was near the top of the vessel, in contact with the nylon film. Yet the distributions of the γ_1^* parameter for α and β events in Runs 2537 and 2539 are nearly identical, with the exception that the events in Run 2537 appear to be reduced by about 15% in energy—a well-known phenomenon affecting surface events.

8.1.4 Coincidence events

The use of coincidences for particle identification has already been discussed somewhat in Chapter 7 in the context of the CTF source runs. The idea is simple: radioactive isotope A decays into isotope B, which is itself radioactive and has a very short lifetime. The necessary condition is that the mean life of isotope B must be much shorter than the average time between unrelated events of similar energies in the detector. A histogram of the time delays in these coincidences should look like a decaying exponential curve whose characteristic decay time is the mean life of isotope B. Coincidences in the CTF are particularly noticeable because the second decay will be observed by the Group 2 electronics system if it happens within 8 ms of the first event. The principal types of coincidence of any relevance that are short enough to have the second event detected by the Group 2 electronics in the CTF are summarized in Table 8.2.

If there exist several unrelated such pairs of isotopes AB, A'B', etc., then either the ratio between the half-lives of any of the isotopes B, B', etc., must be large (at least several times e) so that they fall into a clearly separable hierarchy of decay times, or else the different types of coincidences must be easily distinguishable for other reasons. The former is the case with, for instance, the isotopes of polonium in each of the three heavy-element decay series. The latter is the case for the coincidence ${}^{85}\text{Kr} \rightarrow {}^{85\text{m}}\text{Rb} \rightarrow {}^{85}\text{Rb}$, which has a first event energy less than 200 keV, much less than the high-energy β decays of the isotopes

Coincidence	Progenitor	τ of 2nd event	Branching ratio
²¹² Bi $\xrightarrow{\beta}$ ²¹² Po $\xrightarrow{\alpha}$ ²⁰⁸ Pb	232 Th	$431\mathrm{ns}$	64%
$^{85}\mathrm{Kr} \xrightarrow{\beta} ^{85\mathrm{m}}\mathrm{Rb} \xrightarrow{\gamma} ^{85}\mathrm{Rb}$	$^{85}\mathrm{Kr}$	$1.46\mu{ m s}$	0.43%
²¹⁴ Bi $\xrightarrow{\beta}$ ²¹⁴ Po $\xrightarrow{\alpha}$ ²¹⁰ Pb	238 U / 222 Rn	$237\mu{ m s}$	100%
219 Rn $\xrightarrow{\alpha} ^{215}$ Po $\xrightarrow{\alpha} ^{211}$ Pb	$^{235}\mathrm{U}$	$2.57\mathrm{ms}$	100%
220 Rn $\xrightarrow{\alpha}$ 216 Po $\xrightarrow{\alpha}$ 212 Pb	232 Th	$210\mathrm{ms}$	100%

Table 8.2: The coincidences potentially relevant to the CTF and Borexino, arranged by increasing mean life τ . The decay of the first isotope is immediately followed by the decay of the second isotope. The branching ratio given is the probability that each coincidence will occur in the decay chain of the listed progenitor.

of bismuth. Without this property, this coincidence, with the metastable 85m Rb atom having a mean life of 1.46 μ s, would be difficult to distinguish from a 212 BiPo coincidence. Indeed, when a radial cut is not made, selections of supposed 85 Kr coincidences tend to be highly contaminated by 212 BiPo coincidences occurring on the vessel surface, for which the observed energy of the first event is highly suppressed [44].

By making additional cuts on the energy and relative spatial position of the coincidence events, one can reduce the odds that a coincidence is not merely the result of unrelated atoms decaying within an accidentally short time frame. If one of the isotopes in the coincidence emits an α particle or (in the case of ^{85m}Rb \rightarrow ⁸⁵Rb) a monoenergetic γ ray, an energy cut can exclude accidental coincidences in the wrong energy range. A histogram of the distance $s \equiv |\mathbf{x}_1 - \mathbf{x}_2|$ between reconstructed spatial positions $\mathbf{x}_{1,2}$ of the two events should be proportional to $s^2 e^{-\frac{s^2}{2\sigma^2}}$, where $\sigma \equiv \sqrt{\sigma_1^2 + \sigma_2^2}$ and $\sigma_{1,2}$ are the resolutions of the position reconstruction for each event.² The probability that a coincidence will be observed with a distance between reconstructed event positions of greater than 60 cm is small, on the order of a few percent.

²The reason for this is that each of \mathbf{x}_1 and \mathbf{x}_2 has Gaussian components with respective widths $\sigma_{1,2}$, assuming isotropic uncertainties. Hence the difference vector $\mathbf{s} \equiv \mathbf{x}_1 - \mathbf{x}_2$ has components with Gaussian distribution and width $\sigma = \sqrt{\sigma_1^2 + \sigma_2^2}$. By assumption, the mean values of corresponding components of \mathbf{x}_1 and \mathbf{x}_2 are equal, so the mean of each component of \mathbf{s} is zero. The volume element in radial coordinates is $4\pi s^2 ds$, yielding a distribution for s of $4\pi s^2/(2\pi\sigma^2)^{3/2} e^{-\frac{s^2}{2\sigma^2}}$. We see that the s^2 is therefore a purely geometric factor.



Figure 8.4: Histograms of the photoelectron yield for ²¹⁴BiPo coincidences in Run 2539, one of the CTF 3 source runs. The ²¹⁴Bi energy (part a) is shown at left. The $\beta + \gamma$ spectrum of ²¹⁴Bi is complex and is not suitable for making energy cuts; here events with fewer than 100 photoelectrons have been excluded in order to reduce the probability of accepting accidental coincidences. On the right (part b) is shown the observed ²¹⁴Po energy spectrum. As a monoenergetic α emitter, ²¹⁴Po has a Gaussian spectrum when observed in a detector with finite energy resolution. (This spectrum is quenched by the presence of oxygen to a lower energy than would normally be observed in the detector.)



Figure 8.5: Histograms of (part a) the coincidence time, at left, and (part b) the apparent event separation, at right, for ²¹⁴BiPo coincidences in Run 2539. Both graphs are logarithmic. The fit shown at left is to a decaying exponential plus a constant term, and the fit shown at right is to the function $As^2 e^{-\frac{s^2}{2\sigma^2}}$. Fit parameters are given in the main text.

As an illustration of the method of coincidences, Figures 8.4 and 8.5 show histograms of the event energies, coincidence time, and apparent position separation for ²¹⁴BiPo coincidences in one of the CTF 3 source runs, Run 2539. These events were selected using the same cuts described in Section 7.8. The coincidence time histogram was fit to the function $Ae^{-t/\tau} + B$ with A, B, and τ varying freely; the results were

- $A = 323 \pm 6;$
- $B = -4.2 \pm 0.9;$
- $\tau = 251 \pm 6 \,\mu \text{s}$ (compare to the known value of $237 \,\mu \text{s}$).

The negative value obtained for B is a bit surprising (the existence of background should yield a positive result); it may simply be due to a statistical deficit of events in the last few bins of the histogram. It does not affect the general results that the observed mean life matches the expected value, and that the ratio B/A is small.

The position separation histogram was fit to the function $As^2e^{-\frac{s^2}{2\sigma^2}}$ with A and σ varying freely. The result is that $\sigma = 17.7 \pm 0.1$ cm. Recalling the values of $\sigma_{1,2}$ for Run 2539 (Table 7.10), $\sigma_1 \approx 11.9$ cm and $\sigma_2 \approx 13.0$ cm. (Since the uncertainties are not exactly isotropic, these values were calculated as the arithmetic mean of the σ values for each axis.) Thus, we would have predicted a σ of ~ 17.6 cm, a nice agreement. The histogram diverges slightly from the fit function when s > 50 cm or so. This results from the γ ray component of the ²¹⁴Bi events, which yields non-Gaussian tails in the position reconstruction; refer to Section 7.7.5. In any case, the number of coincidences with s > 60 cm is 286, which is only 3.9% of the total (7302). That is, a cut requiring s < 60 cm on the apparent spatial separations of coincidences will have an efficiency of 96.1% for inclusion of ²¹⁴BiPo events. This is actually a conservative estimate. Events in the source runs are quenched by the presence of oxygen, which reduces the number of observed photoelectrons and increases the uncertainties in position.

8.2 Heavy element decay chain coincidences

It is possible to obtain a measurement of the presence in CTF 3 of isotopes in each of the three naturally occurring heavy-element decay chains. The ²³⁸U, ²³⁵U, and ²³²Th decay chains each include one isotope with a very short half-life (see Table 8.2), so the presence of each chain can be observed using the method of coincidences. It should be stressed, though, that these measurements do not necessarily imply a specific value for the concentration of the parent isotope of each chain. When some isotopes have long lifetimes, it is not safe to assume that secular equilibrium (the condition in which every isotope in a chain produces the same rate of decay events per unit volume) prevails, especially since purification operations on the scintillator will invariably favor some elements over others.

8.2.1 Measurement of ²²²Rn/²³⁸U using the ²¹⁴BiPo coincidence

The predominant coincidence seen in the CTF is the ²¹⁴BiPo that occurs in the ²³⁸U decay chain. However, the meaning of the number of these coincidences that are observed is open to interpretation. The reason is that the decay chain (Figure 2.4a) includes ²²²Rn, a noble gas with a long half-life (3.82 days) and a high mobility. As a result, radon atoms present in the water of the CTF external tank ($\mathcal{A} \approx 30 \text{ mBq/m}^3$ [63]) may diffuse through the nylon shroud and vessel film to decay inside the CTF vessel. In addition, ²²⁶Ra atoms embedded in the vessel film produce radon upon their decay (emanation), which similarly may travel into the scintillator. These effects were discussed earlier in Section 4.4; they guarantee that secular equilibrium higher in the decay chain than ²²²Rn cannot be assumed. There is no coincidence to take advantage of higher in the chain, so it is difficult to tell how much ²³⁸U is present; we may only safely infer the presence of ²²⁶Ra somewhere in or near the scintillator. On the other hand, the half-life of each isotope between ²²²Rn and ²¹⁴Po is less than half an hour; hence, secular equilibrium between radon and the ²¹⁴BiPo coincidences is guaranteed.

One-ton data sample

In order to measure the rate of ²¹⁴BiPo coincidences, the following cuts are made in software on the data. (Because of the energy quenching caused by oxygen in the CTF 3 radon source, the cuts below are not identical to those made for the CTF 3 source runs in Section 8.1.4.) Each cut is listed together with the efficiency ϵ with which the cut will accept true ²¹⁴BiPo coincidences. The total combined efficiency of these cuts is 75.5%:

- The coincidence must consist of a Group 1 event and its corresponding Group 2 event. ($\epsilon \approx 100\%$)
- Neither event may be flagged by the muon veto system. ($\epsilon = 99\%$ [146])
- The energy of the candidate ²¹⁴Bi (first) event must be greater than 200 keV (to avoid accidental coincidences triggered by the high rate of ¹⁴C events) and less than 3.5 MeV. By integration of this cut over the ²¹⁴Bi energy spectrum convolved with a Gaussian resolution function having σ_E(E) = (2.64 keV^{1/2}) √E, ε may be estimated at 99.4%. (A graph of the ²¹⁴Bi theoretical spectrum, albeit convoluted with the CTF 2 energy resolution function instead, is shown in Figure 7.18a.)
- The observed energy of the candidate ²¹⁴Po (second) event must be in the range 300–1200 keV. ($\epsilon \approx 100\%$)
- The α/β discrimination parameter γ_1^* must be greater than -0.01 for the first event $(\epsilon \approx 100\%)$. The utility of this cut can be seen in Figure 8.3.
- With ²¹⁴Po having a mean life of 237 μ s, only coincidences with a time delay between the two events in the range 20–500 μ s are accepted ($\epsilon = 79.8\%$). The lower cut is made in order to avoid accepting events from shorter coincidences such as the ²¹²BiPo in the thorium chain (discussed in the next section). The higher cut excludes as many accidental coincidences as possible without cutting too deeply into the set of ²¹⁴BiPo coincidences.

Runs	Dates	Live time	Tagged	Rate	238 U conc.
		[days]	coincidences	[ev/day/ton]	$[10^{-16} \text{ g/g}]$
2300 - 2346	03/04/12- $03/10/20$	93.86	32	0.48 ± 0.09	4.5 ± 0.8
2350 - 2399	03/10/31 - 04/04/02	103.07	18	0.25 ± 0.06	2.3 ± 0.5
2400 - 2447	04/04/03- $04/09/08$	110.36	31	0.40 ± 0.07	3.7 ± 0.7
2450 - 2499	04/09/08-04/12/09	70.84	34	0.68 ± 0.12	6.4 ± 1.1
2500 - 2531	04/12/10-05/03/13	55.42	34	0.87 ± 0.15	8.1 ± 1.4
2541 - 2563	05/03/18– $05/06/22$	40.47	48	1.69 ± 0.24	15.7 ± 2.3
2300 - 2563	03/04/12 - 05/06/22	474.03	197	0.59 ± 0.04	5.5 ± 0.4

Table 8.3: The ²¹⁴BiPo coincidence rate in several sets of CTF 3 runs. Dates are given in yy/mm/dd format. The number of coincidences listed is the actual number observed; the rest of the columns are corrected for the cut efficiencies and the 65-cm radial cut scale factor. The last column lists the ²³⁸U contamination that would be implied by the assumption of secular equilibrium. Since in fact radon is able to enter the detector by diffusion through nylon, this is not a real value of the ²³⁸U contamination, only an absolute upper limit.

• The spatial distance $|\mathbf{x}_1 - \mathbf{x}_2|$ between the reconstructed positions of the two events must be less than 60 cm, again in order to exclude accidental coincidences. From the results of Section 8.1.4, $\epsilon = 96.1\%$.

One more cut is made for the data analysis: the radial coordinate of the average position of the two events, $r_{\text{avg}} = \frac{1}{2} |\mathbf{x}_1 + \mathbf{x}_2|$, is required to be less than 65 cm. An average position was used as input for the cut, rather than the position of only one of the two events, in order to improve the coincidence position estimates.³ This corresponds to one ton of scintillator. Together with the scale factor calculated for a 65-cm radial cut in Section 8.1.1, the overall efficiency of the analysis is 70.3%. The number of ²¹⁴BiPo coincidences observed must be divided by this value to obtain an estimate of the true number. As the number of surface α events mis-reconstructed to lie within a 60-cm radial cut was previously seen (in Section 7.8) to be less than 0.05% of the total surface rate, we expect negligible contamination in the data sample from surface events, particularly because of the numerous other cuts.

³Statistically speaking, the best position estimate would be a linear combination of the two positions, weighted according to the expected σ values of the reconstruction for each type of event and energy. In practice, weighting each event with a factor of $\frac{1}{2}$ instead appears to make essentially no difference.



Figure 8.6: Energy spectrum of the 197 candidate ²¹⁴BiPo events. At left (part a), the first event energy; at right (part b), the second event energy. The fit shown is to a Gaussian with mean value 770 keV and $\sigma = 67$ keV.



Figure 8.7: At left (part a), the coincidence times of the 197 candidate ²¹⁴BiPo events. The fit shown is to a decaying exponential plus constant term, which here is negative. The e-folding time of the fit is fixed to the ²¹⁴Po mean life. The leftmost bin of the histogram is empty due to the exclusion of candidate events with $t < 20 \,\mu$ s. At right (part b) is shown a histogram of the distance s between reconstructed positions of the event pairs. The histogram goes to zero at s = 0 due to a geometric factor of s^2 in front of the expected Gaussian distribution. Refer to the discussion in Section 8.1.4.

In Table 8.3, we present the results of the analysis for several sets of CTF 3 runs, grouped into periods of 50–100 days live time each. A total of 197 candidate coincidences were found over 474 days of live time. It can be seen that the event rate (corrected for the cut efficiencies and the 65-cm radial cut scale factor) reaches a minimum in the period of Runs 2350–2399 of 0.25 ± 0.06 events per day per ton. This minimum may be a statistical fluke (it is consistent within 3σ with the surrounding two periods). However, the increase in event rate seen starting with Run 2450 is clearly not a statistical accident. It appears to correspond with construction work in Hall C of LNGS that commenced on September 1, 2004. The even greater rate observed for Runs 2541–2563 is perhaps due to a bit of radon contamination entering the vessel during the CTF 3 source runs, Runs 2532–2540.

The lowest event rate observed puts an upper limit on 238 U contamination of less than 2.3×10^{-16} grams per gram of scintillator. This rate is consistent with previously reported 238 U concentration equivalents of $(3.5 \pm 0.7) \times 10^{-16}$ g/g [44] and $(2.8 \pm 0.2) \times 10^{-16}$ g/g [42]. Recall for comparison that the value reported for CTF 1 was $(3.5 \pm 1.3) \times 10^{-16}$ g/g [138]. These are most likely all upper limits. It is hard to tell what portion of this rate is due to uranium in the scintillator in secular equilibrium with its decay products, to radon diffusing at a constant rate through the nylon vessel from water in the external tank, and to radon produced at a constant rate by the decay of radium atoms embedded in the nylon film, end caps and pipes. And even if the radon production were negligible, strictly speaking these results would still only directly give the concentration of 226 Ra, not that of 238 U itself.

The peak of the energy spectrum of the 197 ²¹⁴Po events found in Runs 2300–2563 can be used as a check on the accuracy with which the α quenching function is known. This peak is shown in Figure 8.6b. Fitting it to a Gaussian curve yields a mean value of 770 ± 5 keV, with $\sigma = 67 \pm 5$ keV. This is a slightly higher energy than expected (compare to 751 keV, a 2.5% difference), and the peak is slightly narrower than the prediction (72 keV). The reason for these discrepancies may result from the radial cut at 65 cm; the quenching formula (8.1) was derived from data with no radial cut, some of which therefore exhibited the light loss seen near the vessel surface. The spatial distribution of the ²¹⁴Po events was observed, but the number of events was too small to say anything meaningful about the presence or absence of non-uniform features within the central 65-cm radius volume of scintillator.

Expected level of data contamination

This discussion would be incomplete without investigating possible background in the ²¹⁴BiPo sample. To estimate it, the numbers of singles events⁴ that pass the ²¹⁴Bi and ²¹⁴Po cuts, separately, were determined. The singles events in the sample were required to have a reconstructed radial position of less than 72 cm (this is the maximum possible radial position of an event in a ²¹⁴BiPo coincidence that passes the two cuts $|\mathbf{x}_1 - \mathbf{x}_2| < 60$ cm and $\frac{1}{2} |\mathbf{x}_1 + \mathbf{x}_2| < 65$ cm). The number of possible pairs of them that meet both spatial cuts is 30.9% of the total number of possible pairs (a purely geometric factor based on the assumption of a uniform spatial distribution of singles events).

From this sample of singles events, the number that passed all the ²¹⁴Bi-specific cuts in Runs 2300–2563 was 3.60×10^4 , and the number that passed all ²¹⁴Po-specific cuts was 2.08×10^4 (a total of 1.72×10^4 passed both sets of cuts). The live time during which a singles event passing the ²¹⁴Po cuts would be mistakenly accepted as a ²¹⁴Po event occurring between 20 and 500 μ s after a ²¹⁴Bi-like singles event—was 17.3 s. The total live time of these runs was 474.03 days. Therefore, considering also the geometric factor, the expected number of accidental coincidences in the total data sample is 2.7×10^{-3} . In other words, the probability that the data sample includes an accidental coincidence is one in 370.

The histogram of coincidence delay times for all 197 candidate ²¹⁴BiPo coincidences in Runs 2300–2563 is shown in Figure 8.7a. A fit of this histogram to a decaying exponential plus constant term, $Ae^{-t/\tau} + B$, with τ fixed to the mean life of ²¹⁴Po, yields a constant term of $B = -0.034 \pm 0.041 \,\mu \text{s}^{-1}$. At the 2σ level, this corresponds to fewer than 23 background events (12% contamination) in the sample. Because of the calculations above, it is believed

⁴Recall that singles events are those Group 1 events that have no corresponding Group 2 event.

that the true background is much lower still; nevertheless, an error of this magnitude is small compared to the overall statistical uncertainties in Table 8.3 due to the small sample size.

As an additional check, the number of coincidence event pairs which passed all cuts except the time cut was observed. Instead of the standard 214 BiPo time cut, the cut used for this check was that of a coincidence delay time between 3 and 8 ms, where essentially no 214 BiPo coincidences should fall. The total number of coincidences falling into this "random cut" time window was only two. The resulting estimated data contamination in the 214 BiPo time window, after multiplying by the ratio of time window sizes (0.48 ms/5 ms), was 0.19 events, again highly consistent with zero.

8.2.2 Measurement of ²³²Th using the ²¹²BiPo coincidence

The decay scheme of 232 Th is shown in Figure 2.4b. The detection of 232 Th using the 212 BiPo coincidence is in principle similar to the use of the 214 BiPo coincidence just described. The mean life of 212 Po (431 ns) is much shorter than that of 214 Po, so contamination of the 212 BiPo sample by accidental coincidences is even less of a worry. However, we will see below that the 214 BiPo coincidences themselves may contaminate the 212 BiPo data to some extent.

For this measurement we use some of the same cuts on the data sample as before. In particular, the coincidence cut (the requirement of a Group 1 + Group 2 coincidence), the muon veto cut, and the spatial separation cut are kept. We assume that the spatial separation cut now has an efficiency of $\approx 98\%$ rather than 96.1% as before, since the ²¹²Bi β decay does not produce γ rays. These cuts have a combined efficiency of 97%. The radial cut, requiring the average position of the two coincidence events to have a radial coordinate less than 65 cm, is also kept. The cuts that differ from the case of the ²¹⁴BiPo coincidences are as follows:

- The α/β discrimination cut on the first event is dropped. Because of the very short delay between the first and second event in a ²¹²BiPo coincidence, the ADC 3032 and DPSA board electronic channels may still be gathering data when light from the second event arrives, so the resulting parameters R32 and γ_1^* may not have meaningful values.
- The energy of the first candidate event is required to be in the range 200–2500 keV. (The Q value of the ²¹²Bi β decay is 2.246 MeV.) The efficiency of this cut, calculated by integration over the ²¹²Bi energy spectrum convolved with a Gaussian resolution function having $\sigma_E(E) = (2.64 \text{ keV}^{1/2}) \sqrt{E}$, is $\epsilon = 94.1\%$.
- The energy of the second candidate event is required to be in the range 500–1500 keV (the expected quenched energy of the ²¹²Po α decay is 989 keV). ($\epsilon \approx 100\%$)
- The delay time between the two events is required to be in the range 100 ns-2 μs (ε = 78.4%). Below a 100-ns coincidence time, the efficiency of the Group 2 trigger becomes significantly less than 100%; see, for instance, Figure 9.10.

The total combined efficiency of all these cuts is 71.6%. In addition, we must take into account the vertical scaling factor for the 65-cm radial cut of 1.073 (written as an efficiency, 93.2%) and the fact that the β decay of ²¹²Bi has only a 64% branching ratio (36% of the time it decays by α emission to ²⁰⁸Tl). As a result, the final number of candidate ²¹²BiPo coincidences observed must be divided by 42.6% to obtain an estimate of the true number of ²¹²Bi decays.

We now consider the data contamination expected from stray ²¹⁴BiPo coincidences with these cuts. Statistically, the number of ²¹⁴BiPo coincidences with delay times in the $0.1-2 \mu s$ window should be 1.0% of the 197 with delay times in the 20–500 μs window. Furthermore, tightening the high-end energy cut on the first event of the coincidence from 3.5 MeV to 2.5 MeV reduces the number of ²¹⁴Bi events accepted by a factor of 68.7%. The result is an expected contamination of 1.4 events. The probability that the sample contains an accidental coincidence of two unrelated events, due to the extremely short time window, is negligible.

The result of running the analysis on Runs 2300–2563, 474.03 days of live time, was that a total of nine candidate ²¹²BiPo coincidences were observed. The background-subtracted number of ²¹²BiPo coincidences is then 7.6. Considering the various cut efficiencies, scaling factor, and branching ratio, this corresponds to a ²¹²Bi decay rate of 0.037 ± 0.015 events/day/ton. As in the case of ²³⁸U, it is not entirely safe to assume secular equilibrium. There are several long-lived isotopes in the thorium chain before the ²¹²BiPo coincidence, and strictly speaking, the method yields only the activity of the last of them. (The radon isotope in the chain, "thoron," does not break secular equilibrium is assumed, the concentration of ²³²Th in the scintillator is then $(1.1 \pm 0.4) \times 10^{-16}$ g/g. This value is consistent with the upper limits on thorium contamination reported for CTF 1 [138] and CTF 2 [73].

Of the nine candidate events, the mean value for the ²¹²Po candidate energy was 851 keV, with a standard deviation of 113 keV. This value is significantly different (14%) from the expected quenched energy of 989 keV. Part of this difference is presumably due to the accidental inclusion of one or two ²¹⁴BiPo coincidences, having lower-energy α particles, in the data. However, the statistics are too poor to draw any strong conclusions.

It should be noted that the thorium decay chain also has a second coincidence with a longer mean delay—the decay of ²²⁰Rn followed by that of ²¹⁶Po, which has a 145-ms half-life. This coincidence has several advantages over the ²¹²BiPo coincidence: both events are α decays, so may be more easily tagged; unlike the ²¹²BiPo, this double- α coincidence has essentially a 100% branching ratio; and the isotope ²²⁰Rn itself has a half-life of less than a minute, suggesting the possibility of looking for ²²⁴Ra \rightarrow ²²⁰Rn \rightarrow ²¹⁶Po *triple-\alpha* coincidences. (See Figure 2.4b for the decay scheme.) Indeed, this idea is already being explored in the CTF 3 data [166]. It is even conceivable that the β -decay of the 10.6-hr

half-life isotope ²¹²Pb, which happens between this triple- α coincidence and the ²¹²BiPo coincidence, may be tagged through likelihood-based techniques similar to those described in Section 8.4. If so, a significant contributor of background in the neutrino energy window may be neutralized in Borexino.

8.2.3 Measurement of 235 U using the 219 Rn $\rightarrow {}^{215}$ Po coincidence

Although ²³⁵U has a much lower natural isotopic abundance than its less radioactive sister ²³⁸U, it also has a much greater intrinsic activity due to its shorter half-life. Therefore it is conceivable that the rate of production of ²³⁵U daughter nuclei is high enough to be observable in the CTF. The possibility arises of observing events in the ²³⁵U decay chain via the coincidence of the α decay of ²¹⁹Rn followed by the rapid α decay of ²¹⁵Po ($\tau = 2.57$ ms). Indeed, since the mean life of ²¹⁹Rn is itself only 5.7 s, we could even look for a triple- α coincidence similar to that described immediately above for ²³²Th. The decay scheme of ²³⁵U is shown in Figure 8.8, with more information available in Table 8.4.

For this analysis, we make the following cuts on coincidence events:

- The two events must be a Group 1 + Group 2 pair.
- Neither of the events may have a muon veto flag ($\epsilon \approx 100\%$).
- The time difference between the events must be in the range 1–8 ms ($\epsilon = 63.3\%$). This cut has a low efficiency because of the necessity to exclude ²¹⁴BiPo coincidences at the low end of the range.
- The energies of the two events must each be in the range 300–1000 keV ($\epsilon \approx 100\%$).
- The distance between reconstructed positions of the Group 1/Group 2 coincidence must be less than 60 cm ($\epsilon \approx 98\%$).
- The radial coordinate of the average of the reconstructed positions must be less than 65 cm.



Figure 8.8: Pictorial representation of the Uranium-235 decay chain. Energies shown are Q values for β emitters, and α kinetic energy for α emitters. Times shown are half-lives. Isotopes shaded blue are β emitters with a spectrum endpoint above the 250 keV lower limit of the neutrino energy window. Isotopes shaded yellow are α emitters. (In addition, ²²⁷Ac has a low-probability α -decay branch.) Secular equilibrium is likely to prevail only for ²²⁷Ac and its daughters. Due to the low isotopic abundance of ²³⁵U, none of these isotopes should present a problem for Borexino.

Species	Decay	Q value	$E_{\alpha} \text{ or } E_{\beta+\nu}$	E_{γ}	Branching	Half-life
	mode	[MeV]	[MeV]	[MeV]	ratio $[\%]$	$ au_{1/2}$
$^{235}\mathrm{U}$	α	4.679	4.397	0.205	57	$704\mathrm{Myr}$
231 Th	β^{-}	0.389	0.287	0.102	46	$25.5\mathrm{h}$
231 Pa	α	5.148	5.012	0.046	25.4	$32.8\mathrm{kyr}$
$^{227}\mathrm{Ac}$	β^{-}	0.044	0.044	-	54	$21.8\mathrm{yr}$
227 Th	α	6.147	6.038	-	24.5	$18.7\mathrm{d}$
223 Ra	α	5.979	5.717	0.159	53.7	$11.4\mathrm{d}$
219 Rn	α	6.946	6.819	-	81	$3.96\mathrm{s}$
215 Po	α	7.527	7.386	-	99.9	$1.78\mathrm{ms}$
$^{211}\mathrm{Pb}$	β^{-}	1.373	1.373	-	92.4	$36.1\mathrm{m}$
^{211}Bi	α	6.751	6.623	-	84	$2.15\mathrm{m}$
$^{207}\mathrm{Tl}$	β^{-}	1.422	1.422	-	99.8	$4.77\mathrm{m}$
$^{207}\mathrm{Pb}$	stable					

Table 8.4: Summary table of the Uranium-235 decay chain.

There are two potential sources of data contamination in the sample: ²¹⁴BiPo coincidences with an unusually long delay time, and accidental coincidences of unrelated events. The expected number of ²¹⁴BiPo coincidences with a delay time in the range 1–8 ms is 1.8% of the number having a 20–500 μ s delay. Furthermore, the energy cut of 300–1000 keV on the first event would accept only 4.9% of the ²¹⁴Bi events accepted by the ²¹⁴BiPo first energy cut of 200–3500 keV. Therefore the expected contamination of the data sample by ²¹⁴BiPo coincidences is 197 × 0.018 × 0.049 = 0.17 events.

The number of singles events in Runs 2300–2563 within a 72 cm radius of the origin that meet all of the cuts for each individual event in the coincidence is 1.9×10^4 . The total live time during which a second singles event would create an accidental coincidence is therefore 135 s. Given the total live time of 474.03 days and the geometric factor of 30.9%, the expected number of accidental coincidences in the data sample is 0.02.

That said, only a single coincidence, in Run 2309, passed all of the cuts listed above in all of Runs 2300–2563. The coincidence in question has $E_1 = 849 \text{ keV}$, $E_2 = 380 \text{ keV}$, and $|\mathbf{x}_1 - \mathbf{x}_2| = 53 \text{ cm}$. Upon further investigation, however, it turned out to be a false positive. As noted previously, the reconstruction software (used to generate the data on which this higher-level analysis was run) has the capability to skip over uninteresting low-energy and muon events. In this case, dozens of events between the two in the supposed coincidence, events 209,893 and 209,951, had not been reconstructed in order to save time, including the true Group 1 event to which the second event in the "pair" corresponded. These two events are separated by a time of 53 seconds, and are thus completely unrelated.

With no candidate coincidences, the 90% CL upper limit is the well-known value $-\log (10\%) \approx 2.3$ events. Dividing by the combined efficiency of the cuts (62.0%) and multiplying by the 65-cm radial cut scaling factor yields a 90% CL upper limit on the ²³⁵U daughter coincidence rate of 0.008 events/day/ton, or less than one event per day in the Borexino Fiducial Volume. This is equivalent to a ²³⁵U contamination in the scintillator of $< 1.2 \times 10^{-18}$ g/g.

In nature, the observed isotopic abundances of ²³⁸U and ²³⁵U, respectively, are 99.28% and 0.71% by mass. If these isotopes occur in the CTF 3 scintillator in the same proportions as in nature, the expected mass of ²³⁸U would then be 140 times that of ²³⁵U. In this case, we may estimate (assuming secular equilibrium in the ²³⁵U decay chain) that the amount of ²³⁸U in the scintillator is $< 1.7 \times 10^{-16} \text{ g/g} (90\% \text{ CL})$. Although not inconsistent with the ²³⁸U-equivalent value of $(2.3 \pm 0.5) \times 10^{-16} \text{ g/g}$ found above during the period of least ²¹⁴BiPo activity, this result does suggest that the true uranium contamination in the scintillator is significantly lower. In this case, much of the observed ²¹⁴BiPo coincidence rate would be due to radon atoms diffusing through the CTF vessel or emanating from it (or other nylon parts).

8.3 The decay products of ²¹⁰Pb

The greatest danger of radioactive contamination due to the heavy-element decay chains is presented by the isotope ²¹⁰Pb. This isotope is a decay product of ²¹⁴Po in the radon (or ²³⁸U) decay chain. Since the half-life of ²¹⁰Pb is 22 years, it will typically be out of equilibrium with radon and radon daughters in the scintillator. Atoms of ²¹⁰Pb deposited during the construction of the Borexino nylon vessels, as described in Section 4.4.4, may decay years later. Before being inserted into Borexino, the scintillator may also pick up lead atoms while stored in a steel container, where their production via radon emanation is an ongoing process, or while passing through metal pipes on whose surfaces lead atoms have been adsorbed.

The isotope ²¹⁰Pb is not itself a problem. The Q value of its β decay is 63 keV, so decay events are essentially undetectable, masked by the much greater rate of ¹⁴C decays. However, with a half-life of 5 days, the product ²¹⁰Bi itself β -decays with an energy of 1.16 MeV. A large portion of the β energy spectrum—62% before accounting for finite energy resolution—falls within the neutrino energy window of 250–800 keV. Additionally, the next isotope in the decay chain is ²¹⁰Po, which α -decays with a predicted quenched energy of 396 keV, also in the neutrino window. Both of these therefore represent a real threat to ⁷Be neutrino observation by Borexino. That of ²¹⁰Po is mitigated somewhat by the capability of α/β discrimination, but not eliminated completely.

The half-life of ²¹⁰Po is 138 days (it decays into stable ²⁰⁶Pb), far too long to perform any sort of event correlation. With an essentially invisible ²¹⁰Pb decay, and a ²¹⁰Bi decay that has no characteristics to permit distinguishing it from other neutrino window backgrounds, the ²¹⁰Po α decay is, however, the only handle available for study of this trio of isotopes in the CTF. The presence of ²¹⁰Bi and ²¹⁰Po, out of equilibrium with the ²¹⁴BiPo coincidences, is not only a theoretical worry. Examination of scatter plots of an α/β discrimination parameter against observed event energies shows a large concentration of α -like events at energies in the range 300–400 keV, near the predicted quenched energy of the ²¹⁰Po decay (396 keV). The known rate of α events due to the ²³²Th and ²³⁵U decay chains and the isotopes higher up in the ²³⁸U decay chain is far too low to explain this concentration.

8.3.1 Selection of ²¹⁰Po events

The means of study is via the DPSA skewness method of α/β discrimination. Singles events are selected to lie in the observed energy range 240–550 keV. Only those with a reconstructed radial position less than 65 cm are accepted. The histogram of the γ_1^* parameter of this samples of events is fit to a sum of two Gaussian curves. The presence of ²¹⁰Po may then be estimated in several ways. The simplest method (Method I) is to assume that the area under the Gaussian curve of α -like events is equivalent to the number of ²¹⁰Po events observed (plus a relatively small contribution from α events higher up in the ²³⁸U decay chain). The idea behind Method I is shown in Figure 8.3.

A second method (Method II) is to fit the energy spectrum of the selected events to the sum of a Gaussian curve (representing the α peak) and a decaying exponential (which the



Figure 8.9: Energy spectrum in Runs 2300–2349 for singles events whose reconstructed radial coordinate is less than 65 cm. At left (part a), the energy spectrum is that of all events. At right (part b), the energy spectrum is shown only for events having an α/β discrimination parameter $\gamma_1^* < 0$. In this case the number of α events relative to the background is much higher. Both histograms have been fit to the sum of a Gaussian curve and an exponentially dying curve.

author has found to model the β and γ event background fairly well in the neutrino energy window). The number of ²¹⁰Po events is then taken to be the area under the α peak. This method has the advantage of giving a visual indication of the energy spectrum of the α events seen, which may be compared to the expected Gaussian having a mean at 396 keV and a σ_E of 53 keV. Method II is illustrated in Figure 8.9a.

Method III is a hybrid of the first two methods. From the set of events selected by energy and radial coordinate, only those having the γ_1^* parameter with a value less than zero are considered. Again, the number of ²¹⁰Po events is taken to be the area under the α peak. This "raw" number is then renormalized to the entire range of the γ_1^* parameter by using the parameters of the distribution of γ_1^* for α events already observed via Method I. In the terminology of Section 8.1.3, the renormalization is accomplished by dividing the raw value by $\epsilon_{\alpha}(\gamma_1^* = 0)$, the estimated classification efficiency for α events at $\gamma_1^* = 0$. The uncertainties in results of this method are, as with Method I, correlated with uncertainties in the distribution of α events with respect to the γ_1^* parameter. An energy spectrum fit using Method III is shown in Figure 8.9b.

Estimated number of ²¹⁰ Po events					²¹⁰ Po energy [keV]			
Runs	Method I	Method II	Method III		Method II		Method III	
			raw	renorm'ed	\bar{E}	σ_E	\bar{E}	σ_E
2300 - 2346	1330 ± 135	830 ± 70	735 ± 40	1085 ± 165	392	45	388	48
2350 - 2399	665 ± 140	95 ± 40	230 ± 30	370 ± 125	401	21	380	45
2400 - 2447	875 ± 150	215 ± 50	305 ± 40	525 ± 160	364	31	372	48
2450 - 2499	1085 ± 100	675 ± 70	585 ± 35	875 ± 135	378	44	377	45
2500 - 2531	1070 ± 100	620 ± 60	670 ± 40	900 ± 110	379	40	376	46
2541 - 2563	1115 ± 85	665 ± 75	570 ± 40	890 ± 120	373	47	375	47

Table 8.5: The estimated number of ²¹⁰Po events in the central 65-cm-radius volume of the CTF 3 in six different periods. The number of events is estimated with three methods: the simple use of α/β separation parameters (I), an energy spectrum fit to all events (II), and an energy spectrum fit for only α -like events having $\gamma_1^* < 0$ (III). The "raw" column for Method III shows the raw number of events derived from the energy spectrum fit, while the "renorm'ed" column gives this number divided by $\epsilon_{\alpha}(\gamma_1^* = 0)$, the estimated fraction of the α distribution contained in the region $\gamma_1^* < 0$. In this table the value of $\epsilon_{\alpha}(\gamma_1^* = 0)$ has been calculated separately for each period. Finally, the parameters of the fit to the α peak in the energy spectrum are listed for Methods II and III.

To study the behaviors of these three methods of analysis, the estimated number of ²¹⁰Po events, determined using each of these methods, was calculated for six different time periods. The results are presented in Table 8.5. There is a noticeable and statistically significant difference between the estimated numbers of ²¹⁰Po events determined with the three methods. In every time period, the first method gives a higher number than the third, which in turn is higher than the second.

This discrepancy can be explained if the α and β event distributions over the γ_1^* parameter are not exactly Gaussian curves. Suppose instead that the α event distribution is asymmetric, with a longer tail on the negative side than the positive side, as shown by the dashed lines in Figure 8.10. The estimated classification efficiency for α events is then much greater than previous estimates for $\epsilon_{\alpha}(\gamma_1^*=0)$, which were in the neighborhood of 50%. This implies that, by using the assumption of Gaussian α/β discrimination curves, Methods I and III produce a higher value for the number of α -like decays than the true value. The reason for the ordering of the values produced by the three methods is then



Figure 8.10: A different way of looking at Figure 8.3 that can explain the discrepancy between the three methods used to detect ²¹⁰Po. The solid lines are the result of a fit of two Gaussian curves to the histogram. The dashed lines represent another possible pair of distributions for the α and β events. Since the total area under the Gaussian curve is much greater than under the true distribution, Method I gives the largest values for ²¹⁰Po. Method III, by fitting to an α peak in the energy spectrum, actually counts the true number of α events with $\gamma_1^* < 0$, but then divides by the fraction of the Gaussian curve's area that is to the left of $\gamma_1^* = 0$. This yields a figure smaller than Method I, but it is still too large. Method II performs a fit to the energy spectrum with no cut on the α/β discrimination parameter, so its results should be closest to the true value when the numbers of α and β events are similar. However, during Runs 2350–2447, the number of α decays is small, so the most accurate result then probably comes from applying Method III without dividing by the supposed efficiency of the γ_1^* cut.

clear. This hypothesis is further supported by the observation that the number of ²¹⁰Po events estimated using Method III *without* then dividing by the supposed efficiency of the γ_1^* cut (the "raw" column in Table 8.5) is usually similar to or slightly less than the number obtained with Method II. The exceptions happen in Runs 2350–2447, in which the number of α events is so small compared to the number of β events that the energy spectrum fit of Method II does not work well.

By considering the four other periods in the table, we obtain a weighted average of $91.2 \pm 5.2\%$ for the ratio between the results of Method II and the raw data obtained with

Method III. This is an estimate for the true α classification efficiency $\epsilon_{\alpha}(\gamma_1^* = 0)$ in the range of the ²¹⁰Po energy spectrum.

The quenched energy of the ²¹⁰Po α decays may be estimated from the results of Method III to be about 380 keV, with a σ of 47 keV. This energy is lower than the prediction of Table 8.1 by about 4%. A line drawn through the calculated quenching factors of ²¹⁰Po and ²¹⁴Po (770 keV) observed in CTF 3 yields

$$Q(E) = 22.8 - (1.67 \,\mathrm{MeV}^{-1})E; \tag{8.9}$$

compare with Equation (8.1). Error on this estimate is difficult to determine since it results from a line through only two data points.

8.3.2 Effects of purification tests on the ²¹⁰Po activity

In this section, the effects of the various 2002 CTF 3 purification tests on the ²¹⁰Po activity are discussed. Most prior CTF analyses have not studied the effects of purification methods upon ²¹⁰Po, instead focusing upon isotopes that can be studied more easily via the coincidence method. Given the dominance of ²¹⁰Po in the neutrino window radioactive background for most of the CTF 3 history, it is important to understand how to remove it from the scintillator. Five sets of runs were selected for this study. It may be useful to refer to Table 6.9 to recall their historical contexts.

- Runs 2069–2073 (Dec. 30, 2001 Jan. 9, 2002; 8.70 days livetime) were acquired just before the second batch of pseudocumene was introduced into the CTF vessel. At this point no purification operations had taken place yet, although the original batch of scintillator had been allowed to remain in the vessel undisturbed for about 32 days in order to allow radon to decay away.
- 2) Runs 2117–2122 (Mar. 9–20, 2002; 6.53 days livetime) were acquired 25 days after the first silica gel column test and immediately before the first water extraction test.

Each of these tests was run in "loop mode"; scintillator was continuously circulated through the CTF vessel and the purification system for several days during the test.

 Runs 2144–2149 (Apr. 13–24, 2002; 8.35 days livetime) were acquired 17 days after the first water extraction test.

Shortly afterward, the third batch of new pseudocumene was added to the CTF, followed by the second water extraction test, the fourth batch of new pseudocumene, and the second silica gel column test. Ideally, more runs would have been selected for study between each of these activities. Unfortunately, these tests were performed in such quick succession that there was not time for radon activity to die off between them, making it unfeasible to perform 210 Po analyses of the interim periods.

- Runs 2208–2214 (Jul. 28 Aug. 12, 2002; 10.33 days livetime) were acquired some time (38 days) after the end of the final purification test.
- 5) Runs 2247–2253 (Nov. 2–20, 2002; 13.43 days livetime) were acquired beginning 135 days after the end of the last purification test. These runs were used as a control to check that the ²¹⁰Po activity was decaying away as expected.

The ²¹⁰Po activity was determined for each of the five sets of runs using the Method II described in Section 8.3.1. Recall that in this method, the energy spectrum of all events, with no α/β discrimination cut, is fit to a Gaussian (the ²¹⁰Po energy peak) plus an exponentially decaying background. This method should give good results as long as ²¹⁰Po was the dominant background during the time periods under study. Table 8.6 summarizes the results of the study and demonstrates that this assumption holds.

In Table 8.6, the number of 214 BiPo coincidences and the number of other (non- 210 Po) singles events in the energy range 240–800 keV are shown as well. In every case, 210 Po is the dominant background in the neutrino energy window. The numbers of 214 BiPo coincidences are comparatively negligible, so we need not worry much about contamination of the 210 Po data by 238 U or 222 Rn daughters. Nevertheless, the numbers of 210 Po events seen have first

Data	Relative	Live time	Num	ber of even	Activity	Backgd.	
set	dates [days]	[days]	²¹⁰ Po	Backgd.	$^{214}\mathrm{BiPo}$	[ev/da]	y/ton]
1	0 - 9.9	8.70	2722 ± 60	508	13	331 ± 8	62
2	69.1 - 79.8	6.53	855 ± 42	435	5	138 ± 7	72
3	104.2 - 114.8	8.35	187 ± 20	158	5	22 ± 3	20
4	209.9 - 224.9	10.33	1062 ± 40	224	4	109 ± 4	24
5	307.2 - 325.2	13.43	459 ± 33	428	2	36 ± 3	34

Table 8.6: The ²¹⁰Po activity in each of five periods; see the text for details. The number of ²¹⁴BiPo coincidences and the number of "background" events (singles events with energies in the range 240–800 keV not part of the ²¹⁰Po peak) are also tabulated. All data have been subjected to a radial cut at 65 cm. For reasons of space, as they are not central to the topic, error bars on the ²¹⁴BiPo and background activities have been omitted. The tabulated activities incorporate the 65-cm radial cut scale factor of 1.073, but raw event numbers do not.

had the quantity $N_{\rm BiPo}/75.5\%$ subtracted twice (once for each of the ²²²Rn and ²¹⁸Po α decays that might be confused with ²¹⁰Po events) before being converted into activities.

It is worth noting that, if ²¹⁰Po were in equilibrium with ²¹⁰Pb in the scintillator, we would expect to see a rate of ²¹⁰Bi β decays in the neutrino energy window of about 62% the rate of ²¹⁰Po. The "background" activity rate of β -like events reported in the first row of Table 8.6 precludes an equilibrium concentration of ²¹⁰Pb in the original scintillator. Most polonium in CTF 3 has therefore entered the scintillator as itself, perhaps by desorption from the metal surfaces of the scintillator storage tanks.

The ²¹⁰Po activity $\mathcal{A}(t)$ in the scintillator is a moving target. Not only is it changed by purifications and contaminations, but left to itself, it decays away with a mean lifetime of $\tau = 200$ days if little ²¹⁰Pb is present. Therefore, a better way to look at the data is to normalize them to a baseline value. For simplicity, the date of each set of runs is taken to be the average of the beginning and end dates. If this date is symbolized as t_i for the i^{th} set of runs (i = 1, ..., 5), then the normalized ²¹⁰Po activity relative to the first set of runs may be defined as

$$\mathcal{A}_i^{\text{norm}} = \mathcal{A}_i \, e^{(t_i - t_1)/\tau}.\tag{8.10}$$



Figure 8.11: Activities in the five data sets during the purification tests, shown in bar graph format. For each data set, the left-hand bar is the real activity of ²¹⁰Po and the right-hand bar is the real "background" activity in the neutrino energy window, taken from the last two columns of Table 8.6. The middle, labeled, bar is the ²¹⁰Po activity normalized via Equation (8.10) to remove the effect of ²¹⁰Po decay by dividing out the expected exponential function of time elapsed since the first data set.

These normalized activities are plotted in Figure 8.11. The percent change from each normalized activity to the next is a result only of occurrences between the two sets of data, without any need to take the ²¹⁰Po radioactive decay into account.

As shown in Figure 8.11, a graph of the normalized 210 Po activities, the normalized activity decreased from 331 events/day/ton in period 1 to 195 events/day/ton in period 2. It appears that the combination of the second batch of scintillator and the first silica gel test reduced the 210 Po activity by only about 40%, a purification factor of 1.7. The purification factor predicted from theoretical considerations was about 4 [67]. Reasons for the failure to achieve this expected purification level are not well-understood.

The first water extraction test, in contrast, was much more effective. The normalized ²¹⁰Po activity was decreased from 195 events/day/ton to 37 events/day/ton between periods 2 and 3, a reduction of 81% (purification factor of five). Nevertheless, this purification factor was still much smaller than the theoretically expected result. Two hypotheses that have been suggested include the formation of hydrodynamic circulation stagnant loops in the scintillator that prevented pseudocumene from leaving the scintillator to pass through the water extraction column; and the possibility of ²¹⁰Po ions forming non-polar organic complexes with pseudocumene that shielded them from dissolving into a polar water phase [67]. It also seems likely that particulates enriched in ²¹⁰Po (for which evidence will be presented below) would not be strongly affected by water extraction.

It is clear that a major problem occurred at some point afterwards, between periods 3 and 4. Work done during this time interval included the last two additions of pseudocumene, the second water extraction test, and the second silica gel column test. At least one of these procedures caused the ²¹⁰Po activity to jump by an effective factor of 8.5. It is most unfortunate that these procedures were done in such quick succession that ²¹⁰Po data could not be analyzed from between them. Again, the source of this contamination is not clearly understood, though several possibilities have been considered [67].

Finally, between periods 4 and 5, the normalized polonium activity decreased from an activity of 315 events/day/ton to 171 events/day/ton, a reduction of 45%. This decline is rather surprising, as no major work was done on the detector between those times.

8.3.3 Evidence for ²¹⁰Po-enriched particulates in the scintillator

One hypothesis which perhaps may explain the apparent disappearance of ²¹⁰Po immediately following the last purification test is that of particulates in the scintillator. Consider the possibility that ²¹⁰Po is present inside the CTF vessel in two populations. One population consists of atoms or ions in solution in pseudocumene, with a homogeneous distribution



Figure 8.12: Sketch of the behavior over time of two hypothetical populations of ²¹⁰Po atoms inside the CTF, relative to the 65-cm radial cut (dashed circles). The CTF vessel is shown in vertical cross-section. One population, represented by the gray shading, is in solution in the scintillator. This population decreases only via normal radioactive decay, represented by the shading disappearing as time passes. The second population consists of ²¹⁰Po inside of or adsorbed onto particulates. A representative sampling of these particulates is shown by the small red circles. As time passes, these particles fall to the bottom of the CTF vessel at some terminal velocity v_z dependent upon their densities and radii. (The dotted curves show the maximum z positions of particulates at each given time.) ²¹⁰Po atoms attached to them also decay, illustrated by fading of the red coloring. Between times t_0 and t_1 , events observed within the 65-cm radial cut represent the activity of both populations. Between t_1 and t_3 (at time t_2 , for example), only part of the population in particulates is seen inside the radial cut. After time t_3 , events inside the cut represent solely the activity of the population in solution.



Figure 8.13: Logarithmic plot of the expected ²¹⁰Po activity observed within a 65-cm radial cut as a function of time, assuming two populations of ²¹⁰Po atoms: one in solution, and a second attached to particles that are falling out of the scintillator at a constant uniform velocity. Suppose the initial activities of the two populations are \mathcal{A}_1 and \mathcal{A}_2 . Between times t_0 and t_1 , the apparent activity observed within the radial cut, calculated as the number of events seen per unit time divided by the cut volume, will be $\mathcal{A}_{\text{rad}} = (\mathcal{A}_1 + \mathcal{A}_2) e^{-\frac{t-t_0}{\tau}}$. After time t_3 , it will be only $\mathcal{A}_{\text{rad}} = \mathcal{A}_1 e^{-\frac{t-t_0}{\tau}}$. Between t_1 and t_3 , it will decrease from the first curve to the second curve, causing the observed mean decay time during that period to be shorter than the ²¹⁰Po mean life. Assuming a uniform, constant terminal velocity of particles v_z yields the equation $t_3 - t_1 = \Delta z/v_z$, where Δz is the diameter of the cut volume, 1.3 m.

throughout the volume of scintillator. The second population consists of ²¹⁰Po attached to particles suspended in the scintillator fluid. Although these particles may also initially be uniformly distributed through the CTF vessel, they will slowly fall to the bottom at a constant velocity v_z under the influence of gravity.

Immediately after a thorough mixing (for instance, the second silica gel column test—the final purification test), both populations will be distributed homogeneously in the CTF vessel. Suppose the initial activities of the populations in solution and in particulates at time t_0 are \mathcal{A}_1 and \mathcal{A}_2 , respectively. For a while, the observed ²¹⁰Po activity within a 65-cm radial cut will decay according to the simple exponential law $\mathcal{A}_{rad}(t) = (\mathcal{A}_1 + \mathcal{A}_2) e^{-\frac{t-t_0}{\tau}}$. This is shown schematically in the top two diagrams of Figure 8.12.

At some point, however, after a time we label t_1 , the suspended particles will fall far enough that the top of the radial cut volume is clear of them. (At the later time t_2 , the situation is depicted in the bottom left diagram of Figure 8.12.) At still later times, no suspended particles remain inside the volume of the radial cut. This situation begins at time t_3 (the bottom right diagram of Figure 8.12). After time t_3 , only the ²¹⁰Po in solution remains inside the radial cut, so the observed activity evolves according to $\mathcal{A}_{rad}(t) = \mathcal{A}_1 e^{-\frac{t-t_0}{\tau}}$. Between times t_1 and t_3 , the actual observed activity falls from the upper curve $(\mathcal{A}_1 + \mathcal{A}_2) e^{-\frac{t-t_0}{\tau}}$ to the lower curve $\mathcal{A}_1 e^{-\frac{t-t_0}{\tau}}$, as shown in Figure 8.13. This transition between the two exponential decay curves must appear as a time period during which the mean decay lifetime is faster than the true mean life of ²¹⁰Po.

If for simplicity we assume that diffusion is negligible and that the particles are of uniform size, then given a graph of activity that looks like Figure 8.13, we may estimate the size of the particles. A uniform size implies a uniform terminal velocity. The terminal velocity of a spherical particle with radius r in a fluid of viscosity η and density ρ is the velocity v_z at which the combination of drag and buoyant forces counteract the downward pull of gravity. The force of drag on such a particle falling without turbulence is given by $F_{\uparrow} = 6\pi r \eta v_z$. On the other hand, the combination of gravitational and buoyant forces is given by $F_{\downarrow} = \frac{4}{3}\pi r^3(\rho_p - \rho)g$, with ρ_p being the density of the particle. Equating the two forces,

$$v_z = \frac{2}{9} \frac{r^2 g(\rho_p - \rho)}{\eta},$$
(8.11)

or equivalently,

$$r = 3\sqrt{\frac{v_z\eta}{2g(\rho_p - \rho)}}.$$

The velocity v_z in this equation is simply estimated as $v_z = \Delta z/(t_3 - t_1)$, where Δz is the height of the radial cut volume, 1.3 m.

When the actual CTF 3 ²¹⁰Po data from the first 1.25 years after the purification test are studied, as shown in Figure 8.14, they appear suggestively similar to the theoretical expectations of Figure 8.13. (The ²¹⁰Po activities shown in the figure are estimated by dividing the raw values seen with from Method III (described earlier in Section 8.3.1) by the estimated α classification efficiency $\epsilon_{\alpha}(\gamma_1^*=0) = 91.2 \pm 5.2\%$. No correction has been made for the 65-cm cut scale factor or for possible ²²²Rn chain contamination.) One may go so far as to draw exponential decay curves bracketing the data points to tentatively identify the initial activities: $\mathcal{A}_1 + \mathcal{A}_2 = 125$ events/day/ton, $\mathcal{A}_1 = 43$ events/day/ton. That is, $\mathcal{A}_2/\mathcal{A}_1 = 1.9$; particulates seem to account for almost ²/₃ of the initial ²¹⁰Po activity after the final silica gel test!

These decay curves seem to meet the real data at 35 days and 190 days after the beginning of Figure 8.14, or $t_1 = 39$ days and $t_3 = 194$ days after the end of the silica gel test, respectively. If this effect is due to falling particulates, we expect the ratio of 39 days to $t_3 - t_1 = 155$ days to match the ratio of the 35-cm distance from the top of the CTF vessel to the top of the 65-cm radial cut, to the diameter $\Delta z = 130$ cm of the radial cut. The former ratio is about 0.25, and the latter ratio is about 0.27—a suggestive result.

Let us plug in some numbers. The tentative values above yield $v_z = 9.7 \times 10^{-8} \text{ m/s}$. The density of pseudocumene is 0.875 g/cm^3 , and its viscosity is $0.960 \text{ cP} = 9.6 \times 10^{-4} \text{ Pa s}$ [167].



Figure 8.14: The decay of ²¹⁰Po within the 65-cm (one-ton) radial cut in Runs 2190–2346: June 24, 2002 to October 20, 2003. The first data shown here occur only four days after the last silica gel column purification test. Data are grouped into blocks of ten runs and shown with a logarithmic scale on the vertical axis. They have been bracketed with two exponential decay curves, each with the mean life of ²¹⁰Po; the coefficient of the upper curve is 125 events/day/ton, and that of the lower curve is 43 events/day/ton. The positions of the two vertical lines at 35 days and 190 days were selected to mark the (approximate) times at which the real data depart the upper curve and arrive at the lower curve, respectively. Compare with Figure 8.13.
Runs	Dates	Live time	²¹⁰ Po	238 U contam.		Rate	
		[days]	events	Min	Max	[ev/day]	/ton]
2300 - 2346	03/04/12- $03/10/20$	93.86	808 ± 62	85	254	$7.3~\pm$	1.2
2350 - 2399	03/10/31 - 04/04/02	103.07	250 ± 36	48	143	1.6 \pm	0.6
2400 - 2447	04/04/03- $04/09/08$	110.36	336 ± 46	82	246	$1.7~\pm$	0.9
2450 - 2499	04/09/08-04/12/09	70.84	641 ± 54	90	270	$7.0~\pm$	1.6
2500 - 2531	04/12/10-05/03/13	55.42	737 ± 59	90	270	$10.8~\pm$	2.1
2541 - 2563	05/03/18- $05/06/22$	40.47	628 ± 56	127	381	$9.9~\pm$	3.7

Table 8.7: The estimated rate of ²¹⁰Po events in the CTF 3 scintillator. The raw counts shown are estimated by dividing the raw values from Method III (shown in Table 8.5) by the estimated α classification efficiency 91.2 ± 5.2%. The estimated contamination of the data by ²³⁸U-chain α events is given for two extreme cases: the assumption that all ²¹⁴BiPo coincidences result only from radon diffusion/emanation ("min"), and the assumption of complete secular equilibrium ("max"). The ²¹⁰Po decay rate shown is taken from the raw count rate corrected for the 65-cm radial cut scale factor. The error shown includes a systematic error calculated from the two extreme cases of ²³⁸U-chain event contamination.

Assume a particulate density of $\rho_p = 3 \text{ g/cm}^3$. Then, for spherical particles, Equation (8.11) implies an estimated particle radius of 140 nm.

It is important to warn that there are many potential flaws in this argument. Figure 8.14 exhibits later deviations of the ²¹⁰Po activity from the expected exponential decay, even after all dust should have settled down. In reality one expects a distribution of particles of various sizes and shapes, not simply a collection of uniformly-sized spherules. This distribution, as well as the effects of diffusion, will broaden the distance between the times t_1 and t_3 shown in Figure 8.13. In fact, if sufficiently small particles are present, the observed activity will never settle down to the lower exponential curve, as the particles would remain suspended in the scintillator practically forever. Still, the resemblance of Figures 8.14 and 8.13 is impressive, implying that despite these assorted simplifications, the value of r = 140 nm may actually have some meaning.

8.3.4 Recent behavior of the ²¹⁰Po activity

We now discuss behavior of the ²¹⁰Po activity during the same time frame which was already analyzed for the heavy element decay chain coincidences, Runs 2300–2563. Table 8.7 shows the number of ²¹⁰Po events in each of six periods, estimated by dividing the raw values from Method III (from Table 8.5) by the estimated α classification efficiency $\epsilon_{\alpha}(\gamma_1^* = 0) =$ $91.2 \pm 5.2\%$ derived earlier. Method II was not used for the data due to the relatively small fraction of ²¹⁰Po events in the data for this time range. For each period, also shown are bounds on the likely data contamination by α decays higher up in the ²³⁸U decay chain. The minimum bound assumes that all ²¹⁴BiPo coincidences in the scintillator are a result of radon diffusion or emanation. Figures in this column are derived by taking the number of ²¹⁴BiPo coincidences seen in each period (tabulated in Table 8.3), dividing by the ²¹⁴BiPo efficiency of 75.5\%, and multiplying by two for the two α decays of ²²²Rn and ²¹⁸Po. (The decay of ²¹⁴Po occurs in a coincidence and is excluded from the sample of potential ²¹⁰Po decays with near 100% efficiency). The maximum bound assumes complete secular equilibrium in the ²³⁸U chain until ²¹⁰Pb. It is estimated as three times the minimum bound (assumes that six α decays occur higher up in the ²³⁸U decay chain for each ²¹⁴BiPo).

The ²¹⁰Po decay rate shown in the table is calculated by subtracting the average of the minimum and maximum ²³⁸U contamination from the number of candidate events, multiplying by the 65-cm radial cut scaling factor of 1.073, and dividing by the live time. The error shown for the rate equals the statistical error and a systematic error added in quadrature; the systematic error is estimated to be half of the difference between the maximum and minimum ²³⁸U contamination, multiplied by the 65-cm radial cut scaling factor, and divided by the live time.

The ²¹⁰Po rate reached its lowest point during Runs 2350–2447. Before this, as shown in Figure 8.15, ²¹⁰Po was still decaying away. (A fit to an exponential decay curve plus constant during the period of October 2002 to June 2004 yields an estimated mean life of 183 days, reasonably close to the true value of 200 days.) Observation of a decay curve



Figure 8.15: Decay of ²¹⁰Po event rate in Runs 2240–2418 (October 18, 2002 through June 7, 2004). The *y*-axis measures the rate in units of events/day/ton (note the logarithmic scale). The points shown in red are actual data analyzed using Method III above. They have not been corrected for the 65-cm radial cut scale factor, nor for the presence of radon events. Each data point represents a period of 20 runs. Horizontal bars show the time extent of the data set, while vertical bars show statistical errors at 1σ . These data were fit to a decaying exponential plus constant term in two independent least- χ^2 fits, each fit assigning equal weight to every data point. The green stars represent a fit to the data using the fixed value $\tau = 199.65$ days for the ²¹⁰Po mean life, while the blue stars show the best fit with the decay constant allowed to vary freely. In the first case, the constant term was 0.25 ± 0.01 events/day/ton. In the second case, the best-fit mean life was 182.7 ± 0.5 days, while the constant term was 0.94 ± 0.03 events/day/ton. In reality, this low a level of ²¹⁰Po was never attained; the rate began to rise again in September 2004.



Figure 8.16: The rise in the ²¹⁰Po event rate in Runs 2400–2559 (April 3, 2004 through June 13, 2005). Compare with Figure 8.15. The *y*-axis measures the rate in units of events/day/ton (note the logarithmic scale). The points shown in red are actual data analyzed using Method III above. They have not been corrected for the 65-cm radial cut scale factor, nor for the presence of radon events. Each data point represents a period of 20 runs. Horizontal bars show the time extent of the data set, while vertical bars show statistical errors at 1σ .

implies that ²¹⁰Po during this period was out of equilibrium with its progenitor ²¹⁰Pb, and therefore must have been introduced into the scintillator directly at some point. This phenomenon is also studied in reference [42].

The minimum event rate reached during Runs 2350–2447 (October 31, 2003 through September 8, 2004) was about 1.6 ± 0.6 events/day/ton of scintillator. Though this figure may seem low, it is equivalent to 160 events/day of ²¹⁰Po in the Borexino Fiducial Volume. With an α/β discrimination efficiency in Borexino on the order of 95%, 8 events/day will appear to be β -like. Furthermore, if most of these events result from the decay of ²¹⁰Pb within

the scintillator, another ~100 events/day of the β -decaying ²¹⁰Bi will occur in the neutrino energy window within the Fiducial Volume. This figure is worse than tolerable values by almost two orders of magnitude. It underscores the immense importance of keeping the Borexino scintillator, once it has been distilled, out of contact with any surfaces that may have been exposed to ²¹⁰Pb.

After Run 2450, the observed ²¹⁰Po rate increased again, back to a level of 10 events/day/ton. The increase is shown graphically in Figure 8.16. The growing ²¹⁰Po rate appears related to the increase in the number of ²¹⁴BiPo coincidences seen in the same period. However, it is important to note that the overall increase in the ²¹⁰Po rate (by about 10 events/day/ton) is far too large to be a logical consequence of the ²¹⁴BiPo rate increase (of about one event/day/ton). The two increases are correlated, but do not have a causative relationship.

8.3.5 Spatial distribution of ²¹⁰Po events

The spatial distribution of ²¹⁰Po events was studied for two periods: Runs 2200–2350, the initial period of radioactive decay towards equilibrium; and Runs 2450–2563, the more recent period during which the ²¹⁰Po rate has been increasing. Events were selected from the set of singles events with no muon flag by requiring $\gamma_1^* < 0$ and 240 keV < E < 550 keV. In the period of Runs 2200–2350, roughly 92% of such events within the central ton of scintillator are actually ²¹⁰Po. In the latter period of Runs 2450–2563, about 83% of the selected events within the one-ton radial cut are ²¹⁰Po.

During the initial period of radioactive decay towards equilibrium, ²¹⁰Po was mainly found at the surface of the vessel. The distribution of events within the central volume of scintillator was essentially uniform (Figure 8.17). In contrast, during the period in which the ²¹⁰Po rate started to grow again, the events are mainly found in a vertical stalactite-like column descending from the top of the vessel (Figure 8.18). It is clear that most of the increase in rate is due to the presence of this column. The column is also easily visible in figures



Figure 8.17: Spatial distribution of ²¹⁰Po-like events in Runs 2200–2350, the period during which the isotope was decaying away. The left-hand picture is a vertical cross-section through the CTF vessel *yz*-plane, while the right-hand picture is a horizontal cross-section through the *xy*-plane. Bins are 10 cm on a side, and the cross-section "slabs" are 20 cm thick. Color codes indicate the number of events reconstructed to lie within each bin. The nominal position of the CTF vessel is indicated by the 1-m radius solid circles. Dashed circles indicate the 65-cm radial cut. Events were selected to have $\gamma_1^* < 0$ and 240 keV< E < 550 keV. ²¹⁰Po events within the central volume of scintillator were distributed in a roughly uniform manner.

similar to Figure 8.18 constructed only for Runs 2450–2531; it is not an artifact caused by introduction of the CTF source calibration apparatus in Runs 2532–2540. A similar column was also observed earlier, in the positions of ²¹⁴BiPo coincidences during operation of the CTF 2; refer to Section 6.4.2.

Two hypotheses may explain the presence of the column of events. One hypothesis is that convection is occurring in the CTF scintillator fluid. In this hypothesis, beginning in summer or autumn 2004, the fluid began to circulate by traveling upward near the nylon vessel surface, reaching the top, and then descending in a column along the z-axis. (Conceivably it could also be circulating in the opposite direction.) Atoms of 210 Po are picked up from the nylon vessel by the passing currents. The downdraft at the axis of the



Figure 8.18: Spatial distribution of 210 Po-like events in Runs 2450–2563, a period in which the rate of 210 Po was growing again. These plots are exactly analogous to those of the previous figure, and 210 Po-like events were selected in the same way. Note the vertical column descending from the top of the vessel, a feature that was not present during the earlier period of high 210 Po levels! A large fraction of the event rates listed in Table 8.7 result from the presence of this column.

vessel then carries them along the z-axis through the central region of the CTF. The absence of the thorium-related 212 BiPo coincidences in the center of the vessel may be explained by noting that they occur on the nylon film only at the bottom, not anywhere else (this will be shown in the next chapter). At the bottom of the vessel where the downdraft current spreads out, the fluid velocities should be low so the probability of picking up particles from the surface there is small.

For convection to be possible, a temperature inversion must be present—that is, the temperature near the top of the CTF vessel must be lower than the temperature at the bottom. The fluid at the top becomes colder and denser, therefore sinking, while fluid at the bottom is warmed and becomes less dense, therefore rising. This condition can in principle be tested. Two temperature sensors are located inside the CTF water tank. One is located at a height of 1.5 m above the bottom of the water tank, and the other is at a height of

4.5 m, at the same height as the center of the CTF vessel. For the history of the CTF, the upper sensor has consistently measured a higher temperature than the lower sensor. Unfortunately, data from the period of interest have been lost as a side effect of an upgrade of the control system software used to acquire them.

If the convection hypothesis was correct, it would demonstrate that temperature inversion and convective cells in a large spherical detector are possible, and indicate that positive measures (for instance, active temperature control) must be taken to prevent them in Borexino. However, the still-existing CTF temperature records provide some evidence against the hypothesis. These records are also in line with the normal (*i. e.*, dT/dz > 0) temperature gradient seen in other large spherical detectors such as SNO and KamLAND , so it would be rather surprising for a temperature inversion to develop in the CTF.

The second hypothesis is that the column may be due to ²¹⁰Po falling out of the fluid in the north end pipe of the CTF. If a small leak from the atmosphere was present at the top of the pipe, then a gradient in concentration of ²²²Rn and its daughters would slowly diffuse downward through the pipe. Most of the radon daughters would decay before reaching the CTF vessel, but most of the ²¹⁰Pb would survive. This hypothesis therefore explains the relative scarcity of ²¹⁴BiPo coincidences compared to the increase in ²¹⁰Po events. Accordingly, this leak must be much smaller than a leak which would produce the column of ²¹⁴BiPo coincidences seen in the CTF 2. Still, due to the long half-life of ²¹⁰Pb, the ²¹⁰Po events seen in the column represent only a small fraction of the total amount of radon entering through the top of the pipe.

There is a small problem with this hypothesis. A radon leak should yield individual atoms of ²¹⁰Pb, not heavier particulates. But unless the ²¹⁰Pb is in the form of (or attached to) heavy particulates rather than individual atoms or ions, it would tend to diffuse into the volume of scintillator fluid isotropically from the bottom end of the pipe. A hemisphere-shaped region of high ²¹⁰Po contamination would be produced, rather than the distinct column that is observed.

One possible answer to the problem is that there is *not* in fact a significant radon leak, and the contamination in the column is in the form of particulates that are falling out of the end pipe. What would cause this? During late summer 2004, and continuing to the present (September 2005), there has been an unusual amount of work being done in Hall C. Some of it was related to work on the laboratory's water drainage system and sealing of the Hall C floors to make them spillproof. Other work, some of it right next to or on top of the CTF water tank, involved preparations for a test of scintillator distillation in CTF, which began in February 2006. Vibrations caused by these activities could very well be continually jarring dust particles loose from inside the CTF north pipe, letting them fall into the scintillator. If the dust was enriched in ²¹⁰Po, perhaps caused by previous circulation of ²¹⁰Po contamination through the end pipes in one of the purification tests or pseudocumene batch additions, this could explain the recent increase in ²¹⁰Po events being greater than the increase in ²¹⁴BiPo coincidences.

8.4 Tagging ²²²Rn and its daughters

It is noteworthy that the sum of the half-lives of all the radon daughters through ²¹⁴Po is less than an hour. ²²²Rn decays by α emission into ²¹⁸Po, which (with a half-life of 3.05 min) emits another α particle to become ²¹⁴Pb. This isotope then β -decays (Q = 1.02 MeV) with a half-life of 26.8 min into ²¹⁴Bi. From here, after the ²¹⁴Bi half-life of 19.7 min, the ²¹⁴BiPo coincidence follows. This sequence of events raises the possibility of tagging every decay in the radon decay chain from ²²²Rn through ²¹⁴Po. (The next isotope is ²¹⁰Pb, which was discussed already. With a half-life of 22 years, it breaks the chain.) If this can be done, then many events may be identified as belonging to the series of radioactive decays following the introduction of a radon atom into the scintillator, and individually excluded from the search for neutrinos. The CTF is not an ideal detector for performing such an analysis. Its diameter in units of its position resolution is relatively small (~15; the same ratio for the Fiducial Volume of Borexino, with a position resolution around 15–25 cm [68], is ~30). The α -decaying isotope ²¹⁰Po, which could produce false positives, is present in annoyingly large quantities. Even so, the CTF may provide a feasibility test for tagging radioactive decay chains over medium-length time periods on the order of an hour or less. Indeed, such a test has already been done for the triple- α series ²²⁴Ra \rightarrow ²²⁰Rn \rightarrow ²¹⁶Po (the half-life of ²²⁰Rn is 55.6 s and that of ²¹⁶Po is 145 ms) in the thorium decay chain. Results were mixed [166], perhaps due to the presence of ²¹⁰Po α events.

8.4.1 Selection of candidate events

For each of the 197 ²¹⁴BiPo coincidences selected in Section 8.2.1, the goal was to tag the decays of the ²²²Rn, ²¹⁸Po, and ²¹⁴Pb progenitor atoms. The most difficult of these is ²¹⁴Pb since α/β discrimination is of no use in separating this isotope out from generic CTF 3 background. Nor are we guaranteed that the β decay, whose energy spectrum has a non-negligible component extending to E = 0, will always be seen. In a detector with lower background rates, tagging the ²¹⁴Pb decay may be substantially more successful.

Even for the two α decays, background from ²¹⁰Po is a potential problem. If we consider all events within a 4-hr time window before each ²¹⁴BiPo coincidence, the probability of including the ²²²Rn progenitor in the window is close to 100%. We discard the 20 ²¹⁴BiPo coincidences that occur less than 4 hr after the beginning of a run. Then the average number of background ²¹⁰Po events included within a 65-cm radial cut, per BiPo coincidence, is 1.6. (The estimate was made by breaking down the data into 50-run long periods and using the rates determined for ²¹⁰Po and ²¹⁴BiPo events within each period. If the data are instead considered as a single set, the number of ²¹⁰Po events per BiPo is erroneously low, only 1.2.) Therefore some method must be used to select for the ²²²Rn/²¹⁸Po pairs and against random ²¹⁰Po events; otherwise the signal-to-noise ratio is much too low. Before implementing any sort of selection procedure, the pool of events in each 4-hr window was restricted to those having an apparent energy in the range 200 keV < E < 1200 keV and a reconstructed position within 60 cm of the average reconstructed position of the ²¹⁴BiPo coincidence. The efficiency of these cuts with respect to acceptance of the desired Rn chain events is probably on the order of 95–98%. No explicit radial cut was made on this pool of events.

For the selection procedure, a maximum likelihood method was chosen. For every possible triplet of events occurring in the 4-hr window before a coincidence, a likelihood function \mathcal{L} is evaluated; the triplet with the highest value for the likelihood function is presumed to be the 222 Rn/ 218 Po/ 214 Pb event triplet. The following event characteristics are incorporated into the value of $-\log \mathcal{L}$. Unless otherwise stated, each characteristic is distributed in a Gaussian manner and therefore contributes a term of the form $\frac{1}{2} [(\chi - \bar{\chi})/\sigma_{\chi}]^2$, with χ being the observed value for the event or event pair; $\bar{\chi}$ being the expected average; and σ_{χ} representing the width of the Gaussian.

The α/β discrimination parameter γ_1^* of the three events contributes one independent term for each event. In this case $\bar{\gamma_1^*} = -0.008$ and $\sigma_\gamma = 0.022$ for α decays, while $\bar{\gamma_1^*} = 0.023$ and $\sigma_\gamma = 0.011$ for β decays; refer to Figure 8.3, for instance. These values were estimated by inspecting the α and β distributions of γ_1^* for events over several different periods of runs. As an exception to the usual maximum likelihood methods, however, if $\gamma_1^*/\bar{\gamma_1^*} > 1$ for a particular event, the corresponding likelihood term is set to zero. This is a concession to the reality that an event having $|\gamma_1^*| \gg 0$ is far more likely to be one type of decay than the other, even though in principle prior densities should not enter into a likelihood function. In any event this modification does not affect the results very much.

The energy of the first candidate event E_1 should be similar to the quenched energy expected for ²²²Rn: $\bar{E}_1 = 410 \text{ keV}$, $\sigma_{E_1} = 54 \text{ keV}$. Likewise, the energy of the second candidate event E_2 should be similar to the quenched energy expected for ²¹⁸Po: $\bar{E}_2 = 483 \text{ keV}$, $\sigma_{E_2} = 58 \text{ keV}$. Finally, the energy of the third candidate event E_3 should be consistent with the spectrum of ²¹⁴Pb. In this last case the likelihood term is a constant plus $-\log [\rho_d(E)/\rho_{\text{max}}]$. Here $\rho_d(E)$ is the spectrum (obtained using the theoretical β spectrum shape, taking γ rays into account and incorporating corrections for the finite detector resolution), and ρ_{max} is the maximum function value it attains.

The time lapse Δt_1 between the first and second candidate events is expected to follow an exponential distribution, and thus contributes a term of the form $\Delta t_1/\tau_1 + \log \tau_1$, where τ_1 is the mean life of ²¹⁸Po, 4.40 min. The analogous terms for Δt_2 (between the second and third candidate events) and Δt_3 (between the third candidate event and the ²¹⁴BiPo coincidence) are defined similarly; the respective mean lives of ²¹⁴Pb and ²¹⁴Bi are $\tau_2 = 38.7$ min and $\tau_3 = 28.4$ min.

The spatial positions of the three candidate events and the two events in the ²¹⁴BiPo coincidence should be closely clustered together. We assume that the atom which progressively decays from radon through ²¹⁴Po moves by a negligible amount during that time period. Let the average position of these five events be $\bar{\mathbf{x}}$. Then, each of the five events contributes a term of the form $\frac{1}{2}[(\mathbf{x}_i - \bar{\mathbf{x}})/\sigma_r(E_i)]^2$. The function $\sigma_r(E_i)$ is the energy dependence of the spatial position resolution. Its value (in cm) is given by the formula $\frac{288}{\sqrt{E_i}}$ with E_i in keV.⁵ If the ²¹⁴Bi event has an energy greater than 1.5 MeV, an additional term (0.1 cm keV^{-1/2}) × $\sqrt{E_3} - 1500$ keV is added to $\sigma_r(E_3)$ to model the γ -ray produced behavior seen in Figure 7.19.

To summarize, the minus log of the likelihood function has an α/β term, an energy term, and a timing term for each of the three isotopes being searched for, as well as a total of five spatial terms, each quadratic in all three coordinates.

⁵The value 288 is derived from the observed position resolution in the CTF 3 source run 2539, which had a mean photoelectron yield for ²¹⁴Po events of 126 and a position resolution of ~ 13 cm, and from the photoelectron to energy conversion factor of 3.90 keV/photoelectron.

8.4.2 Monte Carlo simulation of the analysis

A simple Monte Carlo simulation has been used to study the probability that the event triplet with the maximum likelihood function value really is the set of events corresponding to a given ²¹⁴BiPo coincidence.

For each simulated ²¹⁴BiPo coincidence, a set of three progenitor events is simulated. The reconstructed energy spectra of the α decays (of ²²²Rn, ²¹⁸Po and ²¹⁴Po) are simulated with Gaussian random variables having the mean and σ values reported in Table 8.1. The energy spectra of ²¹⁴Pb and ²¹⁴Bi are simulated using sums of ideal β spectra shifted to higher energies by the corresponding γ ray (if any), weighted by the appropriate branching ratios, then convoluted with the energy resolution of the CTF 3.

The α/β discrimination parameter is simulated as a Gaussian random variable, with mean value and σ dependent upon whether the isotope is an α or β emitter. As above, the mean values used are -0.008 ($\sigma = 0.023$) for α decays and 0.022 ($\sigma = 0.011$) for β decays.

The reconstructed event positions are generated assuming that the true position is at the origin. For each of the independent coordinates x, y, z, the simulated position of an event with energy E is given by a Gaussian random variable with mean value zero and $\sigma = \sigma_r(E)$, calculated as described in the previous section.

Decay times between adjacent isotopes in the decay chain are generated by random variables with exponential distribution and mean values given by mean lifetimes of the product isotope.

The simulation incorporates a tunable number of events from two sources of background: ²¹⁰Po α -decay events, and generic β and γ events with an energy spectrum represented by a decaying exponential. Based on observation of the continuous β/γ spectrum in the CTF, the characteristic *e*-folding energy of the exponential is set to 235 keV. The number of each type of background event is generated with Poisson statistics, and the events are distributed uniformly and at random throughout both time and space. This is not entirely realistic, since most ²¹⁰Po and β -like events in the real CTF occur near the nylon vessel film. For ²¹⁴BiPo coincidences occuring within a strict radial cut of 65 cm from the center, however, this should not make much difference. It also neglects the presence of the polonium-enriched column in the most recent CTF 3 data, which should make the simulated estimates a bit optimistic for this period.

When the specific activity of background events is less than about 30 events/day/ton (total), the fraction of simulated progenitors that are tagged correctly by the procedure described in the previous section is about 90% or higher. Predictably, if the background is tuned to consist mostly of ²¹⁰Po α decays, the tagged ²¹⁴Pb events are more likely to be the true (simulated) ²¹⁴Pb decay than is the case with the two α -emitting radon chain isotopes. (For a 30-events/day/ton background of ²¹⁰Po, 91.3% of ²²²Rn events are correctly identified, 92.2% of ²¹⁸Po, and 92.7% of ²¹⁴Pb. 85.8% of triplets have all three events tagged correctly.)

On the other hand, if the background is mostly β -like in nature, then tagging ²¹⁴Pb is the most difficult. (If 30 events/day/ton of generic β -like events are generated in the energy range 200–1200 keV, then 96.0% of ²²²Rn events are identified correctly, 96.1% of ²¹⁸Po, and only 89.0% of ²¹⁴Pb. 86.8% of triplets have all three events tagged correctly.) When the two backgrounds are roughly equal, the figures are still a bit worse for ²¹⁴Pb than for the two desired α decays. This is undoubtedly a result of the fact that, unlike α decays, the energy spectrum of ²¹⁴Pb is not very sharply peaked. Hence the energy of an event does not provide much information about whether or not it is likely to be a ²¹⁴Pb decay.

For the periods of approximately 50 runs in length listed in previous tables, estimates of the data contamination in the tagged events are given in Table 8.8. The ²¹⁰Po rate for each period comes from Table 8.7. The rate of background events in the continuous β/γ spectrum comes from a fit of the energy spectrum to the ²¹⁰Po peak plus exponentially decaying background in the 240–800 keV range. The area under the background curve is then extrapolated to the range 200–1200 keV (assuming a characteristic decay energy

Runs	BiPo	210 Po	β	% mis-identified				Number mis-identified			
	events	rate	rate	222 Rn 2	¹⁸ Po ²	214 Pb	$\alpha/lpha$	222 Rn 2	¹⁸ Po ²	214 Pb	$\alpha/lpha$
2300 - 2346	28	7.3	32.9	5.4	5.0	12.6	6.4	1.5	1.4	3.5	1.8
2350 - 2399	17	1.6	28.9	4.2	4.0	10.9	4.9	0.7	0.7	1.9	0.8
2400 - 2447	29	1.7	33.7	4.4	4.2	12.0	5.2	1.3	1.2	3.5	1.5
2450 - 2499	31	7.0	33.3	5.6	5.1	12.7	6.4	1.7	1.6	3.9	2.0
2500 - 2531	30	10.8	46.2	6.9	6.3	15.7	8.1	2.1	1.9	4.7	2.4
2541 - 2563	42	9.9	41.2	6.5	6.1	14.7	7.6	2.7	2.6	6.2	3.2
Total	177			5.6	5.3	13.4	6.6	10.0	9.4	23.7	11.7

Table 8.8: Expected data contamination for the tagged events corresponding to the ²¹⁴BiPo coincidences in Runs 2300–2563, estimated by Monte Carlo simulation. The numbers in the "BiPo events" column are slightly less than those in Table 8.3 because they exclude coincidences that occur less than 4 hr after the beginning of a run. The rates of the ²¹⁰Po and continuous β background (in the range 200–1200 keV) are given in events/day/ton, including the correction for the 65-cm radial cut scale factor. The "% mis-identified" values are simulated with the Monte Carlo event generator described in the text. Values in the α/α columns indicate the fraction or number of ²²²Rn/²¹⁸Po-tagged event pairs in which at least one event is mis-identified.

interval of 235 keV) and adjusted by the radial cut scale factor. The result is a background β rate which is fairly constant at ~30 events/day/ton, except for a large increase in the most recent periods. Most likely the recent increase corresponds to the same phenomenon causing the increase in the ²¹⁰Po rate.

The results from the Monte Carlo are that for the entire data sample, the expected contamination of the ²²²Rn and ²¹⁸Po data sets is about 10 events per set (out of 177). Considered as pairs, at least one of the two α -decaying events will be mis-identified in about 12 cases. The ²¹⁴Pb fares much worse, with an expected contamination of about 24 events. This is largely due to the high rate of β/γ events in the continuous background. For all three species, in fact, a mis-tagged event is most often a generic β/γ event.

Monte Carlo results for Borexino

In the 100-ton Borexino Fiducial Volume, the rate of background events in the neutrino window is hoped to be in the single digits per day: at maximum, 0.1 events/day/ton. To be very pessimistic, assume a rate of 1 event/day/ton for ²¹⁰Po, 1 event/day/ton for ²²²Rn and each of its daughters, and 0.5 events/day/ton in the neutrino window for random β events. Recall that the rate of neutrino events themselves is predicted to be only 0.35 events/day/ton. These figures imply a total β event rate (including neutrinos) in the 200–1200 keV range of 1.15 events/day/ton. We further assume that a 95% efficiency in α/β separation is possible: $\epsilon_{\alpha}(0) = \epsilon_{\beta}(0) = 0.95$, with Gaussian distributions of the α/β discrimination parameter.

Under these unfavorable circumstances for Borexino, the identification rate for each of the three taggable species, simulated by Monte Carlo, is about 96.7%. This reduces the rate of unidentified ²²²Rn, ²¹⁸Po and ²¹⁴Pb events to less than $\frac{1}{3}$ of the neutrino event rate—small compared to the much more problematic background of generic β decays. In the remaining 3.3% of cases, by far the most common situation is that the pool of candidate events has fewer than three members (because the observed energy of the true ²¹⁴Pb event is below 200 keV, or one of the progenitors has a reconstructed position more than 60 cm from the average reconstructed position of the BiPo, or for some other reason). That is, nothing can be tagged, but neither are any events mis-tagged. In only 0.75% of cases is one of the three tagged events actually a mis-identified random β event, meaning that fewer than 1% of neutrino events will be mis-identified as a ²¹⁴BiPo progenitor. The possibility of excluding ²²²Rn and the following four isotopes in the decay chain from Borexino data, individually rather than statistically, is thus very favorable, even when other backgrounds are high.

8.4.3 Characteristics of the tagged events

Here we examine the results of running the tagging analysis on the real CTF data: 177 ²¹⁴BiPo coincidences that each occurred more than 4 hr after the beginning of a run.

A histogram of the energies of the candidate ²²²Rn events fits nicely to a Gaussian curve (Figure 8.19a). The mean value of the Gaussian is $393 \pm 4 \text{ keV}$, with $\sigma = 50 \pm 4 \text{ keV}$. The expected value (from Table 8.1) is 410 keV with $\sigma = 54 \text{ keV}$. The mean value is thus 4.1% lower than expected, in agreement with the 4% shift of the nearby ²¹⁰Po peak noted in Section 8.3.

The candidate ²¹⁸Po events also have an energy spectrum closely resembling a Gaussian (Figure 8.19b). In this case, a fit yields a mean energy of $473 \pm 5 \text{ keV}$ with $\sigma = 66 \pm 5 \text{ keV}$. As the expected ²¹⁸Po energy is 483 keV (with $\sigma = 58 \text{ keV}$), the observed energy is again a bit lower than expected, about 2.1%.

One may ask whether the nice Gaussian shapes of the candidate ²²²Rn and ²¹⁸Po energy spectra are merely artifacts of the selection method, since likelihoods are assigned to individual events in proportion to Gaussian functions peaked at the expected energies of ²²²Rn and ²¹⁸Po. It is possible to check for such an artifact by excluding all energy terms from the likelihood function during the analysis. Doing so degrades the accuracy of the likelihood fit, but the remaining terms—spatial distribution, α/β discrimination, and time lapses—still select in favor of the true ²²²Rn and daughter events. If this modification causes the energy spectra of the candidate events to assume a significantly different shape, then we would conclude that the Gaussian energy spectra are largely a result of the selection method. However, if the energy spectra remain mostly Gaussian in nature, it indicates that the maximum likelihood method does generally select the correct events. In the worst case, during Runs 2500–2531 (the period of highest ²¹⁰Po and β/γ background), the Monte Carlo simulation predicts a data contamination of 13–14% for each α -decaying species when the energy terms of the likelihood function are neglected.



Figure 8.19: The energies of the 177 candidate 222 Rn events (part a, left) and candidate 218 Po events (part b, right) selected using the maximum likelihood method. Both energy spectra are neat Gaussian curves having peaks within 3% of the expected values.



Figure 8.20: The energies of the 177 candidate 222 Rn events (part a, left) and candidate 218 Po events (part b, right) selected using the maximum likelihood method, without incorporating any energy-dependent terms. The energy spectra are still near-Gaussian, demonstrating that the selected events mostly are the true 222 Rn and 218 Po events with little accidental inclusion of background.



Figure 8.21: Energy spectrum of ²¹⁴Pb. At left (part a) is an ideal energy spectrum convoluted with the energy resolution function of the CTF, in this case simulated by the Monte Carlo generator. At right (part b) is the actual spectrum of 177 candidate ²¹⁴Pb events. Of these events, roughly 24 are believed to have been mis-identified.

The resulting energy spectra of the candidate ²²²Rn and ²¹⁸Po events after this modification to the analysis are shown in Figure 8.20. The ²²²Rn candidate event spectrum is still sharply peaked; a Gaussian fit to it yields a mean value of 391 keV, with $\sigma = 56$ keV. Some lowenergy noise appears toward the low end of the energy window. The ²¹⁸Po candidates fare similarly. A fit to a Gaussian gives a mean value of 457 kev ($\sigma = 63$ keV). Without the lowenergy spurious events, the peak might even still be in the range of 470 keV. We conclude that for the most part, the Gaussian energy spectra of the selected events really are present, and most selected events are the true ²²²Rn/²¹⁸Po pairs.

The expected energy spectrum of a pure sample of 10^6 simulated ²¹⁴Pb events is shown in Figure 8.21a, while the actual spectrum is shown in Figure 8.21b. This latter sample is expected to contain about 24 events that are not actually ²¹⁴Pb out of 177 events (13%).

Histograms for the original method of analysis (incorporating the energy terms in the likelihood function) were also generated for the elapsed times between the candidate ²²²Rn and ²¹⁸Po events, and between the candidate ²¹⁸Po event and the ²¹⁴BiPo coincidence. As already mentioned, the first of these histograms should look like an exponential decay with a mean life of $\tau_1 = 4.40$ min. The second distribution, being the distribution of the sum of



Figure 8.22: The time delays between the candidate 222 Rn and candidate 218 Po events (part a, left) and between the candidate 218 Po events and 214 BiPo coincidences (part b, right) selected using the maximum likelihood method. The fit at left is to a decaying exponential plus constant term. The fit at right is to the distribution of the sum of two exponential random variables, plus a constant term. Mean lives of the isotopes in question (218 Po: 4.4 min; 214 Pb: 39 min; 214 Bi: 28 min) were kept fixed for the fits.



Figure 8.23: The $\gamma_1^* \alpha/\beta$ discrimination parameter for the 177 candidate ²²²Rn events (part a, left) and candidate ²¹⁸Po events (part b, right) selected using the maximum likelihood method. Both histograms appear consistent with the expected distribution of γ_1^* for α decays; compare with Figure 8.3. The roughly Gaussian, symmetric shapes of the histograms may be artifacts of the assumed likelihood function, and does not necessarily contradict the hypothesis shown in Figure 8.10.

two exponential random variables ($\tau_2 = 38.7 \text{ min}, \tau_3 = 28.4 \text{ min}$), should take the form

$$f(t) \sim \frac{e^{-t/\tau_2} - e^{-t/\tau_3}}{\tau_2 - \tau_3}.$$
(8.12)

This function increases linearly from zero, peaks at about 33 min, and then decreases roughly exponentially. This histogram of the combined decay time was considered in preference to histograms of the individual decay times of ²¹⁴Pb and ²¹⁴Bi, for the reason that the sample of ²¹⁴Pb events is expected to be twice as contaminated as the sample of ²¹⁸Po events.

Each histogram was fit to the appropriate distribution plus a constant term. An attempt was made to allow the isotope mean lifetimes to float in the fit, but the resulting error bars were wide enough to be useless (e. g., $\tau_1 = 5.9 \pm 2.1$ min). Hence the mean lives were kept fixed to the true values, and only the multiplicative constants on the distributions, as well as the constant terms, were allowed to vary.

For the candidate 222 Rn and 218 Po events, the resulting number of pairs in the exponential distribution was 149 ± 48 , while the number of pairs fit in the constant background was 12 ± 36 . The time difference distribution is shown in Figure 8.22a. In the case of the elapsed times between candidate 218 Po events and 214 BiPo coincidences (shown in Figure 8.22b), the number of pairs in the expected distribution was 137 ± 20 , with an estimated 21 ± 17 constant background events. That is, both sets of histograms have a constant background term that is consistent (within 2σ) with zero. The errors in these values are too large to observe the expected data contamination of about 10 events.

Finally, histograms were generated for the $\gamma_1^* \alpha/\beta$ discrimination parameter of the candidate event pairs. These are shown in Figure 8.23. Both are consistent with the expected distribution of the parameter for α decays.

8.4.4 Looking for scintillator convection currents with tagged events

From the tagged events, we may estimate the scintillator velocity in the z direction near the z-axis in an attempt to observe convection. Each pair of 222 Rn/ 218 Po events is assumed to occur in one place, and after roughly an hour, the 214 BiPo coincidences occur at a (perhaps) slightly different location. The average velocity in z is given by the Δz between the 222 Rn/ 218 Po pair average position and the 214 BiPo coincidence average position, divided by the time difference. That is, labeling the events { 222 Rn, 218 Po, 214 Pb, 214 Bi, 214 Po } as {1, 2, 3, 4, 5}, we estimate

$$v_z = \frac{\frac{1}{2}(z_1 + z_2) - \frac{1}{2}(z_4 + z_5)}{\frac{1}{2}(t_1 + t_2) - \frac{1}{2}(t_4 + t_5)} = \frac{z_1 + z_2 - z_4 - z_5}{t_1 + t_2 - t_4 - t_5}.$$
(8.13)

The CTF 3 runs were analyzed in two periods: Runs 2300–2379, and Runs 2450–2563. The first period is that during which the distribution of ²¹⁰Po is essentially uniform. (Runs prior to about 2300 cannot be used; before that, the prevalence of ²¹⁰Po was too high for tagging of radon and its daughter isotopes to be feasible.) In the second period the column of events has formed. The ²¹⁴BiPo coincidences are selected in the same manner as in Section 8.2, and as before, coincidences that happen less than 4 hr after the beginning of a run are excluded.

In neither period is a fluid velocity observed that is inconsistent with zero. Histograms of the z velocity, calculated with the tagged progenitor events of each ²¹⁴BiPo, are shown in Figure 8.24. During the first period, only 33 candidate coincidences occurred, and the histogram does not even appear to have much of a peak. Fitting a Gaussian curve to it anyway yields a mean value of $v_z = -1.4 \pm 9.2 \text{ cm/hr}$, with a σ for the fit of $30 \pm 12 \text{ cm/hr}$. We may convert this to an upper limit on $|v_z|$, less than 15.3 cm/hr at 90% CL.

In the second period, there were 103 candidate coincidences. In this case the calculated z velocities were distributed with an obvious peak. (The cause of the sharp spike seen in Figure 8.24b is unknown.) A Gaussian fit resulted in a mean value of $v_z = -0.4 \pm 2.3 \text{ cm/hr}$, with a σ of 16 \pm 3 cm/hr. This value of v_z may be converted to the upper limit $|v_z| < 3.8 \text{ cm/hr}$ at 90% CL.



Figure 8.24: Distribution of velocities of the 222 Rn atom and its daughters in the z direction, calculated using event tagging and Equation (8.13). In part a (left) is shown the distribution for 33 214 BiPo coincidences in Runs 2300–2379. On the right (part b) is shown the distribution for 103 coincidences in Runs 2450–2563. A Gaussian fit was attempted to each distribution, but the results were less than satisfactory.

One may ask whether restricting the candidate events only to those within a certain distance from the z-axis yields more conclusive results, as perhaps a convection current would be more pronounced right on the axis. The answer is no. Histograms of v_z for events with a cylindrical radial coordinate ρ restricted to be within 20 cm of the z-axis show no additional signs of a non-zero fluid flow. Furthermore, scatter plots of the calculated values of v_z versus the ρ coordinate of the ²¹⁴BiPo coincidence average position do not show any correlation between the two variables, in either period. The data do not allow us either to confirm or to rule out convection currents in the CTF scintillator fluid.

The same Monte Carlo simulation as before was used in order to find out how strong a convection current would have to be in order to be detectable with the method of radon daughter tagging described here. For this purpose, an offset was added to the reconstructed z position of each event simulated in the ²²²Rn decay chain. The offset value was given by the time of the event (relative to the original ²²²Rn decay) multiplied by the assumed current velocity. Then, 100 ²²²Rn events and their daughters, together with an amount of background noise equivalent to that seen in Runs 2500–2531, were generated at a time. For each radon chain, the "most likely" radon daughters were tagged with the maximum

likelihood method. In each set of 100 radon chains plus noise, a histogram was constructed of the estimated atomic z-velocity, calculated with the tagged radon daughters as above, and a Gaussian fit was performed on the histogram. The mean and the estimated error in the mean were reported for each set.

From the simulation, it was concluded that the current velocity required for most samples of 100 radon chains to yield a mean estimated velocity farther from zero than the estimated error is on the order of $|v_z| \approx 8 \text{ cm/hr}$. As this is much larger than the previously estimated upper limit of 3.8 cm/hr, we require much better statistics in order to measure the actual value of the CTF scintillator velocity, if non-zero.

On the other hand, a lower limit for the velocity in the column may be estimated from the diffusion coefficient. A particle having diffusion coefficient D in the scintillator fluid will travel an RMS distance in the xy-plane of $\sqrt{\frac{2}{3}Dt}$ from its origin in a time t. (The factor of $\frac{2}{3}$ is a result of considering diffusion in only two dimensions of a three-dimensional problem.) If this distance was more than about 50 cm during the period it takes for fluid to travel the 2 m from the top to bottom of the CTF vessel, the column of contaminants shown in Figure 8.18 would spread out and become indistinct. That is, the lower limit is given by roughly

$$v_z > \frac{D}{19\,\mathrm{cm}}.\tag{8.14}$$

(This is a very conservative limit; one might claim that the maximum column spread in the xy-plane has a radius of only 20 cm, which would set the denominator of Equation (8.14) to 3 cm.) For individual radon atoms in pseudocumene, $D \approx 2 \times 10^{-5} \text{ cm}^2/\text{s}$, and the diffusion coefficient is known not to depend strongly upon the identity of the atom in question. We therefore have a lower limit on the velocity of $v_z > 0.09 \text{ cm/day}$ if the contamination in the column is in the form of individual atoms. (Less conservatively, this lower limit would be raised to 0.6 cm/day.) So, assuming the presence of convection, the column velocity is presumably in the range $0.09 \text{ cm/day} < v_z < 8 \text{ cm/hr}$.

If contamination in the column is in the form of particulate, this lower limit becomes smaller (less stringent). The Stokes-Einstein equation states that, for a spherical particle,

$$D = \frac{k_B T}{\pi \eta r} \tag{8.15}$$

(η being the viscosity of the liquid and r being the particle radius). As a result, the lower limit on column velocity is inversely proportional to the mean radius of the particles.

Even if we assume no convection (an assumption which, judging by the temperature record, seems likely to be accurate), particulates may be large enough to fall to the bottom of the CTF vessel under the influence of gravity. (The observed ²¹⁴BiPo coincidences may or may not follow the distribution of ²¹⁰Po, depending upon whether or not they are also mainly attached to particulates.) Using the terminal velocity formula that was derived earlier, Equation (8.11), and combining it with Equations (8.14) and (8.15), we therefore have

$$\frac{2}{9} \frac{r^2 g(\rho_p - \rho)}{\eta} > (19 \,\mathrm{cm})^{-1} \frac{k_B T}{\pi \eta r},$$

or equivalently,

$$r > \left[\frac{9/2}{19\,\mathrm{cm}} \frac{k_B T}{\pi g(\rho_p - \rho)}\right]^{1/3}.$$
(8.16)

For the sake of plugging in numbers, let us assume $\rho_p = 3 \text{ g/cm}^3$, a typical density for rock dust. The CTF temperature is about $18^{\circ}\text{C} = 291 \text{ K}$. For pseudocumene, $\rho = 0.875 \text{ g/cm}^3$. These figures yield a minimum particle radius of 11 nm. With the less conservative figures for the column spread in the *xy*-plane, the minimum particle radius becomes 20 nm. It is worth noting that the smallest filter pore size in any of the CTF purification modules is $0.05 \,\mu\text{m} = 50 \text{ nm}$, explaining how these particles might remain present in the scintillator.

The viscosity of pseudocumene at 20°C is $0.960 \text{ cP} = 9.6 \times 10^{-4} \text{ Pas}$ [167]; therefore from Equation (8.11), a particle of density 3 g/cm^3 and radius 11 nm (20 nm) would take 11 yr (3 yr, respectively) to fall the 2 m height of the CTF vessel. This is inconsistent with the appearance of the column over a period of a few months; even the less conservative lower

limit yields a falling time that is wrong by an order of magnitude. We must conclude that the average particle size is at least $\sqrt{10} \approx 3$ times larger than a 20 nm radius, implying a minimum radius of about 60 nm. This is consistent with the value of r = 140 nm derived in Section 8.3.3; however, there is insufficient evidence to decide whether the particulate contamination hypothesized in that section is identical to that making up the more recently observed column of high ²¹⁰Po activity.

8.5 Lighter radioactive isotopes

The bulk of this chapter has focused upon CTF internal contamination due to heavy isotopes in extensive decay chains. Of course, other radioactive isotopes are also of great concern in the CTF and Borexino. Limited time and space, however, prevented them from being studied extensively for this work. A review of previous results is nevertheless provided in this section for completeness, as is a new analysis of 40 K.

It is useful to mention here, for reference, the total internal event rate within the neutrino energy window (250–800 keV). For Runs 2300–2563, the average value of that rate, determined using a 65-cm radial cut, is 33 events/day/ton. For Runs 2350–2447, the period of lowest internal contamination, the average neutrino window rate is "only" 25 events/day/ton. A quick comparison to the numbers in Tables 8.3 and 8.7 makes it clear that uranium daughters can account for only a small fraction of this event rate.

With the exception of the coincidence decay of 85 Kr (which has a very small branching ratio, so measuring it is not a very sensitive technique) and the electron capture photopeak of 40 K, the lighter radioactive isotopes do not exhibit any strongly distinguishing features permitting their concentrations to be easily determined. Analyses of these isotopes are therefore generally restricted to a fit of the observed energy spectrum to the superposition of a large number of theoretical curves with parameterized amplitudes. One of the most sophisticated such fits, from reference [168], is shown in Figure 8.25. It is the opinion of



Figure 8.25: A sample of a CTF spectral fit; this one comes from reference [168]. The fit is against the set of events in Runs 2316–2481 having a reconstructed radial position (using the Milan reconstruction code) of r < 60 cm. Some parameters are constrained using the concentrations of contaminants already known from coincidence and α/β analyses. This fit, for additional sophistication, is also combined with a spectral fit over all radii (not shown) to get better measurements of the external γ ray contribution, which was modeled using Monte Carlo methods.

the author that spectral fits to the continuous smooth curves of β -decay spectra must in general have shallow χ^2 functions and are not too much to be trusted. The utility of such fits is made even more questionable by uncertainties in the spatial and energy distributions of external γ rays. This energy spectrum has in the past been modeled by elementary functions such as exponentials, linear functions, and even constants, in addition to various histograms produced in Monte Carlo simulations.

We here mention the two important carbon isotopes, the cosmogenic ¹¹C and the high-rate but low-energy scintillator contaminant ¹⁴C, only briefly. The former has, with the spatial and time cut technique described in Section 2.2.2, been measured to have a rate in CTF of $0.135 \pm 0.024 \ (stat.) \pm 0.014 \ (syst.)$ events per day per ton [22, 41], compatible with the prediction of 15 ± 2 events per day in the 100-ton Borexino Fiducial Volume. The latter has been measured to have a mass fraction in the CTF scintillator of $(5 \pm 2) \times 10^{-18}$ [44], 3.5 times higher than the scintillator of the original CTF 1. The large error in this number is not statistical; it comes from the observation of four different batches of pseudocumene added to the CTF 3 at different times, which had measured ¹⁴C mass fractions ranging from $(3.7-7.1) \times 10^{-18}$ [44]. It is hoped that the scintillator used in Borexino does not present still more such surprises.

8.5.1 The noble gases: ⁸⁵Kr and ³⁹Ar

A considerable concern for Borexino is the possibility of radioactive noble gas isotopes entering the detector via the nitrogen gas supply. For this reason, a special source of nitrogen low in argon and krypton gases will be used. This nitrogen was not available in time to be tested with the CTF, however. As a result it is thought that a significant number of events in the CTF detector are due to the β decays of ³⁹Ar and ⁸⁵Kr.

The more tractable of these is ⁸⁵Kr. The reason for its tractability, as mentioned in Section 2.2.2, is a decay path with a small branching ratio that permits use of the method of delayed coincidences. Usually ⁸⁵Kr undergoes the pure β decay ⁸⁵Kr \rightarrow ⁸⁵Rb. However, with a branching ratio of 0.43%, it may instead decay to an excited state known as ^{85m}Rb with a half-life of 1.01 μ s. The excited state decays to the ⁸⁵Rb nuclear ground state (a stable nucleus) by emitting a γ ray of energy 514 keV. That is, a β decay with maximum energy 173 keV is followed within a few μ s by a monoenergetic γ ray.

These coincidences have been sought in several different analyses. That of A. Pocar (2003), for instance, finds a total rate for 85 Kr decays of 7 ± 5 events/day/ton [44]. The study

covered the period of Runs 2125–2300, a total of 259.5 days of live time, and was restricted to a radial cut of events with r < 70 cm. (Lifting the radial cut led to a significant sample contamination from ²¹²BiPo coincidences at the vessel surface.) A later study by D. Franco (2005) investigated several periods, of which the longest consisted of Runs 2346– 2434 (165.8 days live time). This investigation, using a radial cut of 80 cm, found a total ⁸⁵Kr activity of 13 ± 4 events/day/ton [42]. Two other periods studied, respectively contained by and overlapping the end of the study of A. Pocar, yielded activities of 5 ± 2 and 16 ± 7 events/day/ton [42]. A third analysis by Aldo Ianni (2005) used Runs 2180–2474 (555 days live time). Although no radial cut was made, this study took care to account for background on the vessel surface from the ²¹²BiPo coincidence. The result obtained was a total ⁸⁵Kr activity of $9.4^{+3.2}_{-2.4}$ events/day/ton [169]. The wide scatter between these results indicates the low sensitivity of the method of coincidences for this case.

The isotope ³⁹Ar has no tagging method available at all. It can be studied only through spectral fits to the observed CTF energy spectrum. Because the spectra of ³⁹Ar and ⁸⁵Kr are quite similar above the ¹⁴C end-point, it is not really possible to determine the rate of the isotopes separately, but only their total rates. Values for this total in the neutrino energy window have been reported in the ranges of 7–13 events/day/ton. The fractions of the ³⁹Ar and ⁸⁵Kr spectra above 250 keV are roughly $\frac{1}{3}$, giving an estimated total noble gas activity around 20–40 events/day/ton. From the coincidence analyses, about 10 events/day/ton are due to krypton, so the remaining 10–30 events/day/ton must come from argon.

This estimated rate of argon decays is surprisingly high, since in normal air, activity from krypton is nearly 100 times higher than that from argon. In looking through old CTF 3 logbooks, Aldo Ianni has apparently found the answer to the mystery [169]. When the CTF 3 vessel was originally installed in 2001, it was (like the Borexino vessels) first inflated with nitrogen before being filled with fluid. However, near the end of the inflation process, the underground N₂ supply was exhausted. It was decided to finish the inflation with pure argon gas. With some reasonable assumptions (a 0.1 m^3 volume of argon added; a reduction factor of 20 after nitrogen stripping) and the partitioning constants of argon from nitrogen

to water (31) and nitrogen to pseudocumene (4.1), such an error in judgment was predicted to result in about 35 events/day/ton from ³⁹Ar in the CTF [170].

8.5.2 Long-lived natural radioisotopes: ⁸⁷Rb and ⁴⁰K

Other than the heavy element decay chains, there are two relatively common, naturally occurring, long-lived radioisotopes. One is ⁸⁷Rb (half-life 47.5 Gyr) and the other is ⁴⁰K (half-life 1.277 Gyr). ⁸⁷Rb, however, is a pure β emitter with a Q-value of 273 keV, near the low end of the neutrino energy window. In CTF, detection of ⁸⁷Rb is infeasible due to the large number of β events caused by ³⁹Ar and ⁸⁵Kr. Nor does it seem like a very problematic contaminant for Borexino.

 40 K, on the other hand, has a Q-value of 1311 keV, placing its decays throughout the neutrino energy window. Of all the natural, long-lived radioisotopes, 40 K has the highest activity in rock or dirt, making it a serious concern. For this reason it is important to understand the 40 K contamination seen in CTF 3.

It is possible to estimate the rate of ⁴⁰K β -decay events, since the isotope has a second decay branch with amplitude 10.7% in which it decays via electron capture to ⁴⁰Ar. The electron capture produces an excited ⁴⁰Ar nucleus which immediately emits a 1461 keV γ ray. This γ ray, being monoenergetic, can in principle be detected as a "bump," or photopeak, in the CTF energy spectrum rising out of the relatively uniform β and γ background. The peak should be Gaussian in shape with a predicted σ of 101 keV.

As this photopeak has not been studied in detail in CTF 3 before, some simple analyses were done on it for this work. Runs 2300–2563 (474 days livetime) were taken as the data set. The CTF scintillator volume was split up by several radial cuts. Two concentric regions of equal volume were defined, each containing 229 kg of scintillator: region Ia, $r < 2^{-1/3} \times 50 \approx 40$ cm; and region Ib, $2^{-1/3} \times 50$ cm < r < 50 cm. This was done in order to check for a radial dependence of the ⁴⁰K event activity; if the activity has a strong



Figure 8.26: Division of the CTF scintillator into several radially symmetric volumes, shown here in cross-section, for the purpose of 40 K analysis. The CTF vessel is represented by the largest (bold) circle. Regions I, II, and III each include $\frac{1}{8}$ of the CTF scintillator. Regions Ia and Ib each have a volume of half that. The volume outside region III is not analyzed in order to avoid edge effects on the perceived spatial distribution of internal γ rays. See text for details.

component due to potassium in the vessel film or outside the vessel, it should grow rapidly as a function of distance from the detector's center. Similarly, three larger regions of equal volume, each containing 458 kg of scintillator, were analyzed (Figure 8.26): region I, with r < 50 cm (the union of regions Ia and Ib); region II, having 50 cm $< r < 2^{1/3} \times 50 \approx 63$ cm; and region III, with $2^{1/3} \times 50$ cm $< r < 3^{1/3} \times 50 \approx 72.1$ cm.

In every case, the analysis was done by observing the energy spectrum for each volume in the range 1150–2200 keV, and fitting these spectra to the function

$$\rho(E) = A \, e^{-(E-1150 \,\mathrm{keV})/\varepsilon} + \frac{N_{\mathrm{K40}}}{\sqrt{2\pi} \,\sigma_0} \, e^{-\frac{(E-E_0)^2}{2\sigma_0^2}} + \frac{N_{\mathrm{Bi214}}}{\sqrt{2\pi} \,\sigma_1} \, e^{-\frac{(E-E_1)^2}{2\sigma_1^2}}.$$
(8.17)

The additional Gaussian term labeled $N_{\rm Bi214}$ takes into account a second peak seen in the energy spectrum, which presumably results from the 1.76 MeV γ ray emitted (with 16% probability) in the decay of ²¹⁴Bi. Internal ²¹⁴Bi decays are generally seen in coincidence

	40 K γ peak			$^{214}\mathrm{Bi}\;\gamma\;\mathrm{peak}$			Background		
Volume cut	E_0	σ_0	Events	E_1	σ_1	Events	A	ε	
[cm]	$[\mathrm{keV}]$	$[\mathrm{keV}]$		$[\mathrm{keV}]$	$[\mathrm{keV}]$		$[bin^{-1}]$	$[\mathrm{keV}]$	
Ia $r < 40$	1576	101	112 ± 38	1953	> 120	65 ± 33	30 ± 4	615 ± 270	
Ib $40 < r < 50$	1538	> 120	113 ± 29	1945	> 120	94 ± 30	38 ± 4	706 ± 193	
Subtotal: $Ia + Ib$			225 ± 48			159 ± 45			
I $r < 50$	1564	114	222 ± 68	1947	> 120	159 ± 46	71 ± 6	659 ± 170	
II $50 < r < 63$	1536	> 120	318 ± 48	1884	114	223 ± 86	116 ± 7	823 ± 144	
III $63 < r < 72$	1566	> 120	428 ± 53	1906	84	215 ± 58	176 ± 8	709 ± 71	

Table 8.9: Results for the fit to the 40 K and 214 Bi photopeaks in the various concentric volume cuts in CTF 3. All columns except the first are the fit parameters. For reasons of space, fit uncertainties in the parameters E_i and σ_i (i = 0, 1) are not shown, but were in the range of 15–40 keV. The subtotal for the individual analyses of volumes Ia and Ib (row 3) is compared with the analysis of all of volume I (row 4) as a consistency check.

with the decay of ²¹⁴Po immediately following. But when ²¹⁴Bi decays in the water buffer, outside the CTF vessel, only γ rays may penetrate into the scintillator and be recorded by the PMTs. Other background in the energy spectrum was (as usual) modeled by a decaying exponential. The peak width σ_0 was constrained to be in the range 70–120 keV; E_0 was constrained to be near the expected value 1461 keV. For the supposed ²¹⁴Bi peak, the width σ_1 and mean value E_1 were constrained similarly. All other parameters were constrained only to be non-negative. The energy spectra and the results of the fit for four of the defined volumes are shown in Figure 8.27. Results are given in Table 8.9.

The number of ⁴⁰K events in each volume increases as one travels farther from the center of the CTF vessel. For this reason, it is clear that the data include γ rays originating in the nylon vessel, the water buffer, or external portions of the CTF (nylon pipes, collars, and so forth). To obtain better results for internal ⁴⁰K, it is necessary to examine the distribution of scintillation events caused by external γ rays; this will be done in the following chapter. For now we will simply quote the value obtained within the central 50-cm volume of the CTF as an upper limit: 222 ± 68 events in 474 days livetime, or 1.0 ± 0.3 events/day/ton, are observed. Taking into account the β/γ branching ratio, this number corresponds to an upper limit of 850 events/day of ⁴⁰K β emissions in the Borexino Fiducial Volume.



Figure 8.27: The energy spectra of several small volumes inside the CTF, fit to the presumed 40 K photopeak using Equation (8.17). At top, r < 40 cm and 40 cm < r < 50 cm; these regions have the same volume. At bottom, r < 50 cm and 50 cm < r < 63 cm; these regions also have the same volume, twice that of each of the regions shown at top. The spectrum for 63 cm < r < 72 cm is not shown, but is similar in nature (albeit with greater amplitude).

Chapter 9

CTF Surface Contamination and External Backgrounds

The preceding chapter included a thorough analysis of internal events in the CTF, those happening in the volume of scintillator. In this chapter we examine some characteristics of the surface and external contributions to events observed in the CTF. Recall that surface events are defined as those produced by radioactive materials embedded in the nylon vessel film, while external events are produced by radioactive decays outside the vessel.

Given that the range of an α particle in water or scintillator is negligible, while that of a β particle is on the order of millimeters, these particles will not be detected within the scintillator fluid if they originate outside the nylon vessel. They therefore make no contribution to the observed external activity, only to the surface activity. The observed rate of β emissions from the nylon film will be roughly half the true surface event rate: a β particle emitted from the nylon film has slightly less than a 50% chance of being directed inward into the scintillator, where it can be detected, instead of outward into the water buffer or sideways into the nylon where it is invisible. An α particle, with a range on the order of 100 μ m, is only likely to be detected if it originates from the thin skin of nylon film closest to the scintillator volume, as well as being directed inward. In addition, an α produced within the nylon film will lose energy before making its way into the scintillator (even assuming it can do so), leading to a low-energy tail of events in the α energy spectrum. In this chapter, we therefore presume that observed coincidences at the surface originate mostly from atoms adhering to the inner surface of the nylon vessel, not from those embedded within the vessel's nylon film.

The only radiation produced external to the nylon vessel that has a chance to be detected, and therefore the only contribution to the set of external events, is in the form of γ rays. These may travel many centimeters from their origins, permitting them to enter the volume of scintillator. Once inside the scintillator, a γ -ray photon will often scatter several times, producing more scintillation light each time, before being completely absorbed. A γ ray detected within the scintillator may also have first been scattered several times in the external water buffer. As a result, events that might be expected to produce a monoenergetic γ line are instead smeared out across lower energies. This characteristic makes an external γ -ray background intractable to analytic methods, and one must resort either to crude approximations or to Monte Carlo simulations in order to understand the observed energy spectrum and spatial distribution of γ rays.

9.1 Hot spots on the CTF vessel

Without yet making an attempt to understand the species present in the nylon film, we can observe that some regions of the CTF vessel are apparently more radioactive than others. It is worth taking a moment to discuss these hot spots. It should first be noted that in this chapter, the range of θ , the latitudinal coordinate in the spherical (r, θ, ϕ) coordinate system, is taken to be $-90^{\circ} < \theta < +90^{\circ}$, with $\theta = 0$ defining a great circle in the *xy*-plane (the CTF vessel "equator"). This convention is used in order to maintain the conceit of the CTF vessel as a globe having north and south poles at latitude $\pm 90^{\circ}$ and equator at latitude 0° . As in the previous chapter, the data used in this chapter come from Runs 2300–2563 unless stated otherwise.



Figure 9.1: Histogram of the sine of the latitude of neutrino energy window events reconstructed in the range 80 cm < r < 120 cm. Regions on a sphere bounded by circles of latitude whose sines differ by a constant amount have equal area. Thus a uniform distribution of surface events would produce a constant histogram. In this case, the histogram is mostly constant except for spikes at the south pole (left) and north pole (right), and a low-amplitude, broad bump near the equator (center). The bins of the histogram cover regions with $\Delta(\sin \theta) = 0.05$; therefore, each bin represents $\frac{1}{40}$ of the surface area of the CTF vessel.

Low-background detectors must be built with the utmost care in order to prevent radioactive contamination. For instance, an unknown mistake during construction of the acrylic vessel of the Sudbury Neutrino Observatory produced a hot spot nicknamed the "Berkeley blob," a point-like source on the vessel with an activity equivalent to about $9 \,\mu g$ of ²³²Th [171]. Though this hot spot had the serendipitous effect of providing a calibration source intrinsic to the SNO detector, its presence could have been problematic if SNO was intended to observe neutrinos at lower energy ranges. The CTF hot spots can tell us what to expect in the way of similar features on the Borexino Inner Vessel.
9.1.1 The north and south end regions

The most immediately obvious regions of higher radioactivity in the CTF are the north and south poles of the vessel. This is apparent in vertical cross-sections of the detector such as Figures 8.17a and 8.18a. We can also look directly at the distribution in θ of surface events. Selecting only those events in the neutrino window with a radial coordinate between 80 and 120 cm, we plot a histogram of the sines of their latitudes in Figure 9.1. It is easy to show that the region on a sphere bounded by two circles of constant latitude θ_1 and θ_2 has surface area $A = 2\pi r^2 |\sin \theta_1 - \sin \theta_2|$. Therefore, a uniform surface distribution should yield a constant-valued histogram. As seen in the figure, the histogram is fairly uniform between latitudes 55° S and 55° N (sin 55° = 0.82), but there are spikes at both poles.

The hotness of the poles comes as no surprise, since more material is present there (the two solid nylon end regions) than anywhere else on the vessel. Furthermore, there is evidence (described in the previous chapter, and in Section 9.2.2) that particulate material has settled at the south pole of the vessel, increasing the radioactivity further.

The average observed event rate per unit area during Runs 2300–2563, in the neutrino energy window (250–800 keV), is 75 events/day/m² for the north polar region ($\theta > 60^{\circ}$ N), 72 events/day/m² for the south polar region ($\theta < 60^{\circ}$ S), and 54 events/day/m² for the vessel film between 60° S and 60° N. These regions on the vessel have respective areas of 0.842, 0.842, and 10.88 m². The event rate includes events having a reconstructed radial coordinate between 80 and 120 cm; it therefore also incorporates a large number of internal events. Of course, it also cannot help but include many external γ -ray events. Nevertheless, the large difference between the quoted values shows that the poles are substantially more radioactive than the rest of the vessel film.

We may do somewhat better by subtracting away the estimated contribution due to internal events. From Section 8.5, we have an average value of 33 events/day in the innermost 65cm radius (or one ton) of scintillator volume. The expected ratio between the numbers of internal events with radial coordinates observed in the range 80–120 cm and in the range 0–65 cm is roughly 1.75, a figure obtained by performing the relevant integrals over the $f_d(r)$ function defined by Equation (5.48) and taking their ratio. (The value of the position resolution σ was assumed to be 15 cm, but using 10 cm instead changes the ratio only slightly.) This implies that the surface events are contaminated by internal events in the amount of 33 events/day ×1.75/12.56 m² = 4.6 events/day/m². This is small compared to the overall surface rate. Subtracting off the value for internal events gives respective values for the north polar region, south polar region, and remainder of the nylon film of 70, 67, and 49 events/day/m²: the poles are at least about 40% hotter than the rest of the vessel.

If we consider these same values only during the period of lowest contamination in the scintillator, roughly Runs 2350–2447 (live time 213.4 days), we obtain respective values (corrected to subtract off internal events) for the north polar region, south polar region, and remainder of the nylon film of 71 events/day/m², 67 events/day/m², and 47 events/day/m². The average internal contamination in the neutrino energy window during this period is 25 events/day/ton rather than 33. Despite this decrease in internal contamination, the combined surface and external contamination observed seems to be fairly constant and stable.

9.1.2 Hot spots on the nylon film

As shown in Figure 9.2a, the radioactivity of the CTF nylon vessel in the neutrino energy window is non-uniform even far from the poles. There are three roughly circular hot spots located on the vessel equator, distributed more or less at equal distances around it. These hot spots were observed to remain constant in position over time. The figure has been corrected to exclude an estimated contribution due to internal events of 4.6 events/day/m².

Including the internal event subtraction, the hot spots have maximum amplitudes of 75-85 events/day/m², compared to the minimum (background) value seen of 30 events/day/m².



Figure 9.2: "Maps" of radioactive hot spots on the surface of the CTF vessel. Bins are 5° on each side, and the colors correspond to the event rate observed in each bin during Runs 2300–2563, in units of events/day/m². (Bins near the poles of course have different area than bins near the equator.) The data shown here include only events with reconstructed radial coordinates between 80–120 cm, and reconstructed latitudinal coordinates (θ) between -70° and $+70^{\circ}$. In addition, the constant internal event contribution has been statistically subtracted away, although it is small compared to the observed rate of surface events. The upper map (part a) shows events in the neutrino energy range, while the lower map (part b) is for events in the energy range 1.2–2.0 MeV. Note the suppressed zero of the color scale in the upper map; the two color scales differ. Most interestingly, the pattern of hot spots varies depending upon the energy range observed.

This is almost a factor of three difference! Without the hot spots, the average vessel surface contamination away from the poles would be 330 events/day in the neutrino energy window instead of the observed 510. The origin of these hot spots is unknown. A rough estimate of the activity of the most radioactive spot, at $\phi = 240^{\circ}$, yields about 12 events/day over the background value. This rate is similar to the figure of 0.7 events/hr (17 events/day) made public for SNO's hot spot [171]. However, it must be noted that the low energy threshold for SNO is 30 PMT hits, corresponding in the case of that detector to a several-MeV event, so the total rate of the SNO hot spot is probably much higher than the CTF hot spots. This represents at least a small triumph for the method of manufacture of the CTF vessel. In Borexino, the vessels were assembled in a yet more careful manner; some details are given in reference [44].

A similar map for higher energy events, 1.2-2.0 MeV, is shown in Figure 9.2b. Some equatorial hot spots are still present, but they do not precisely match up with those on the other map. This implies that different hot spots are produced by a different mix of radioactive isotopes, and therefore have different origins. The higher energy events also seem to be more prominent in the northern hemisphere, likely due to the greater number of γ rays produced by the bulkier north end-cap and pipe.

9.1.3 Attempts to observe the hold-down ropes

Recall from Chapter 6 that the CTF vessel, which is suspended in water but filled with less dense organic scintillator fluid, is held down against the buoyant force by 16 nylon monofilament lines. Since each of these lines is strung over the top of the nylon vessel and tied down at the bottom of the water tank on each end, 32 nylon lines cross the CTF equator in all, or two per panel of nylon film.

We may ask whether it is possible to see a pattern of events due to these lines. A pattern could also be seen due to the seams between the nylon film panels if the glue used was sufficiently radioactive. In either case, the pattern should be most apparent in a histogram of the number of events in specific ranges of the longitudinal coordinate ϕ . Such a histogram is shown in Figure 9.3 for events with energies in the range 1.2–2.0 MeV (bracketing the 1.46 MeV ⁴⁰K electron capture photopeak and the 1.76 MeV ²¹⁴Bi γ ray). The events in this figure have latitudes between 60° S and 60° N, and radial coordinates between 80 and 120 cm. Though no pattern with a period of 11.25° (360°/32) is apparent, the hot spots on the equator are quite obvious.

It is possible that the Fourier transform of this histogram is more informative. The unnormalized discrete Fourier transform of a histogram with N intervals numbered $j = 0, \ldots, N-1$ is defined by

$$Y_k = \sum_{j=0}^{N-1} X_j e^{-2\pi i j k/N},$$
(9.1)

where X_j is the value for the j^{th} interval of the histogram. It is similar to the well-known continuous Fourier transform, but adjusted for the case in which only a finite amount of data is known. The k^{th} element of the result is a measure of the sinusoidal component of the original histogram which repeats k times over the histogram range.

The complex Fourier components Y_k for k = 0, ..., 50 of the longitudinal histogram were calculated numerically using the FFTW software package [172], version 3.0.1. Their magnitudes $|Y_k|$ are shown in Figure 9.4. Note the logarithmic scale on the y-axis. The first maximum of the plot (after the constant component k = 0) is at k = 2, corresponding to the two most obvious hot spots on the equator. No significant maxima are seen afterward; in particular, the expected large peaks at k = 16 and k = 32 are not present.

If the nylon ropes do contribute significantly to the spectrum, either they are in too much disarray to contribute a periodic component, or they are sufficiently close together that the finite spatial resolution of the detector smears their signature to the point of invisibility in Fourier space. This latter possibility at first seems unlikely, since the average separation between ropes at the equator should be 19.6 cm, rather greater than the detector spatial



Figure 9.3: Histogram of events in the energy range 1.2–2.0 MeV as a function of longitude (the CTF ϕ coordinate). Bins are 1° of longitude wide. The events selected for inclusion had reconstructed radial coordinates between 80–120 cm, and reconstructed latitudinal coordinates (θ) between -60° and +60°. The data are taken from Runs 2300–2563. The peaks of the graph correspond to the hot spots of the surface activity map in Figure 9.2b.



Figure 9.4: The discrete Fourier transform of Figure 9.3. Magnitudes are shown for Fourier components 0 through 50. Note the logarithmic scale on the vertical axis.

resolution. We must remember, though, that the ropes may only be observed via external γ rays entering the scintillator. The difficulty of reconstructing γ ray positions accurately, when they scatter multiple times at locations several cm apart, has already been mentioned. In any case, it is also quite probable that the ropes do not make a significant contribution to background; their total mass is very small.

9.2 Radon and thorium in or near the nylon film

We can search for heavy elements (the ²³⁸U and ²³²Th decay chains) on the nylon vessel in the same way as in the bulk of the scintillator: via the method of coincidences. The event loss due to geometric considerations, however, means that the method becomes unreliable for surface events. For more than $\frac{1}{4}$ of coincidences occurring on the inner surface of the nylon vessel, neither event in the coincidence will be detected. For another more than $\frac{1}{2}$, only one of the two events will be seen, implying that the set of supposed singles events is itself contaminated by coincidence events. (For instance, it is not necessarily true that the energy spectrum of α -like singles events is free of ²¹⁴Po or ²¹²Po.) Only in less than $\frac{1}{4}$ of cases can the full coincidence be detected. And this is the best-case scenario, for decays that occur on the inner surface of the nylon vessel, not embedded inside the nylon film. Any study of coincidences at the vessel surface must bear this fact in mind.

9.2.1 ²¹⁴BiPo coincidences and the implications for radon

We will use the same set of cuts defined in Section 8.2.1. In this case, we now consider all coincidences passing the cuts, not only those with a radial coordinate of less than 65 cm. Two basic phenomena may be studied with these data: the contamination of the nylon vessel film by uranium and radon, and the amount of light collected by the CTF as a function of the event radial coordinate.



Figure 9.5: Histogram of the sine of the latitude of 214 Po events reconstructed in the range 80 cm < r < 120 cm. A uniform distribution of surface events would produce a constant histogram. Here, the histogram is mostly constant except for spikes at the south pole (left) and north pole (right). The south pole is clearly more radioactive in 214 Po. Compare with the 210 Po distributions shown in Figures 8.17 and 8.18.

First, however, consider the distribution in θ of surface events. Selecting only those ²¹⁴Po events (from the ²¹⁴BiPo delayed coincidences) with a radial coordinate between 80 and 120 cm, we plot a histogram of the sines of their latitudes in Figure 9.5. As before, a uniform surface distribution should yield a constant-valued histogram. The histogram is indeed quite uniform between latitudes 55° S and 55° N (sin 55° = 0.82), but there are spikes at both poles. They indicate that, as with the total set of events in the neutrino window, there are concentrations of radon daughters near the top and bottom of the vessel. These may be indicative of several things. There could be leaks of radon entering from the two end pipes; the poles could be hotter due to the presence of the end pipes (and other masses of bulk nylon) themselves; or, as proposed before, particulate could have settled onto the bottom of the vessel. The last explanation seems perhaps the best one, given that the south pole has a higher amount of ²¹⁴Po than the north.

Radial distribution of the coincidences

In reference [42], D. Franco puts Equation (5.48), the expected distribution function $f_d(r)$ of internal events, to good use in fitting the radial distribution of ²¹⁴Po events from ²¹⁴BiPo's. The use of only the internal event distribution function was a reasonable simplification. In that case, the period under examination included Runs 2074–2090, which had a high internal radon contamination (over 100 ²¹⁴BiPo coincidences/day/ton).

During the recent period of Runs 2300–2563, the amount of radon in the detector was much lower, on the order of one coincidence/day/ton or less. Therefore, we have a reasonable hope of observing surface contamination by fitting the observed radial distribution of coincidences to the function

$$f(r) = N_i f_i(r;\sigma) + N_s f_s(r;\sigma), \tag{9.2}$$

where N_i is the number of internal events in a data sample, N_s is the number of surface events, and f_i , f_s are the functions of Equations (5.48) and (5.51), respectively. (For simplicity, we take σ to be a constant parameter of the model instead of a radially-dependent function.) Instead of using the radial coordinate of the average position of the two events, as in Section 8.2.1, we follow reference [42] and only use the radial coordinates of the ²¹⁴Po event; this avoids problems of reconstructing the γ -ray rich ²¹⁴Bi events near the vessel film. The free parameters of the fit are R, the CTF vessel radius; σ , the position resolution of the detector; and N_i , the number of internal events. R and σ are kept free since we do not know the precise shape of the vessel, nor the exact dependence of σ on the radial coordinate of an event. It is hoped that by letting them float, they will settle down to reasonable averaged values. The integral over f(r) is required to be N, the total number of ²¹⁴BiPo coincidences, so N_s is constrained to equal $N - N_i$.

The set of coincidences meeting all the cuts includes 1099 214 Po coincidence events. 197 of these were previously analyzed in Section 8.2.1. However, as one additional tweak to the analysis, all 214 Po events with a latitude coordinate θ whose absolute value is greater than 60° are excluded. (In other words, we exclude twin cones with apexes at the origin and symmetry axes on the z-axis.) This tweak removes from the data sample whatever may be the cause of the higher concentration of radon daughter events at the poles of the vessel, permitting us to consider only diffusion through the nylon film and emanation from it as sources of radon. The excluded region is shaped like a double cone, apexes at the vessel center, in order to keep the ratio of surface to volume included by the cut constant over all radii. The cut excludes only 13.4% of internal events if they are distributed homogeneously, so the statistics should not suffer. Once this volume cut is made, 787 candidate ²¹⁴Po events remain in the data sample (the excluded double cone volume thus contains 312 possible ²¹⁴Po events).

We now briefly consider the possibility of contamination of the data sample with accidental coincidences. If the time cut is relaxed to accept coincidences with delay times up to $\Delta t = 1000 \,\mu$ s instead of 500 μ s, there are 894 candidate coincidences instead of 787. A fit of the coincidence time for these 894 coincidences to an exponential decay curve (keeping the decay time constant fixed to the mean lifetime of ²¹⁴Po) plus constant term (Figure 9.6) yields an expected total of 951 ± 39 true ²¹⁴BiPo coincidences, and a value of -17 ± 19 accidental coincidences. The former value is consistent with the observed 787 events meeting the more restrictive time cut, divided by that cut's efficiency of 79.8% (resulting in 986 expected coincidences). The latter value gives an upper limit at 90% CL of 7 accidental coincidences in the range 20–1000 μ s, or 3.5 accidental coincidences among the set of coincidences meeting the more restrictive time cut. This value, happily, is negligible.

The fit to the radial distribution, Equation (9.2), for the ²¹⁴Po events meeting all the cuts (including the volume cut and the more restrictive coincidence time cut) is shown in Figure 9.7. This fit yields a value of 9.3 ± 1.4 cm for σ , while the value of the vessel radius parameter R is found to be 112 ± 4 cm. The unexpectedly large value of R may result from the distorted hot-air balloon shape of the vessel. The number of internal events was determined to be 653 ± 87 . Since the amount of scintillator in the CTF is 3.85 tons [169], and the volume cut excludes 13.4% of it, the number of observed internal events per ton



Figure 9.6: Histogram of the coincidence times for candidate ²¹⁴BiPo coincidences in Runs 2300–2563, with the double cone-shaped volume containing the vessel poles excluded. They were fit to an exponential decay plus constant term. Only candidates with $20 \,\mu s < \Delta t < 500 \,\mu s$ were used in the main analysis, although those having Δt up to 1000 μs are shown and fit in this histogram. The fit results give an expected total of 951 ± 39 true ²¹⁴BiPo coincidences, and a value of -17 ± 19 accidental coincidences with Δt in the range $20-1000 \,\mu s$. The ²¹⁴Po mean life of $237 \,\mu s$ was kept fixed in the fit.

is 196 ± 26 . This is to be compared with the 197 events observed in the central 65-cm radius of scintillator; multiplying that value by the 65-cm radial cut scale factor of 1.073 yields 211 events/ton. The two values are consistent, implying that the internal ²¹⁴BiPo coincidences are distributed uniformly within the vessel.

The derived surface contamination of $N_s = N - N_i = 134 \pm 87$ events implies a surface event rate, extrapolated to the entire vessel surface, of 0.33 ± 0.22 events/day. [In the excluded volume, on the other hand, the number of surface events may be estimated as $(1099 - 787) - (196 \text{ events/ton} \times 3.85 \text{ tons} \times 13.4\%) = 211 \pm 26$, implying a rate per unit area ten times greater near the poles!] We may correct the overall surface event rate by the efficiency of the coincidence time cut, $\epsilon = 79.8\%$, to obtain 0.41 ± 0.28 surface events/day. The efficiencies of the other cuts used in detecting coincidences are not well-known for surface events. For instance, about half of the γ -ray emissions from ²¹⁴Bi decays near the



Figure 9.7: Distribution of the reconstructed radial coordinates of the ²¹⁴Po events in the 787 observed ²¹⁴BiPo coincidences within the volume cut during Runs 2300–2563. The black curve is the best fit to Equation (9.2). The blue and green curves represent the internal and surface components of the fit, respectively. They have respective integrals of 653 events and 134 events.

vessel surface are lost to the water buffer, which as shown below shifts the spectrum of that isotope to lower energies, decreasing the efficiency of that energy cut. Let us estimate their combined efficiencies at 85%; then the overall surface event rate becomes 0.48 ± 0.33 surface events/day.

Surface events may be attributed both to radon diffusing inward from the CTF water buffer, to radon emanation from radium atoms embedded in the nylon film, and to radioactive atoms adsorbed on the inner film surface. If we suppose that the CTF scintillator fluid does not exhibit convection, radon atoms from these sources will be present within a thin (few cm) boundary layer just inside the vessel. The thickness of this boundary layer is effectively zero, given the finite position reconstruction resolution. In the worst case, all observed events are from atoms on the inner surface of the vessel or in "dirt" attached to it; in the best case, all radon atoms whose progeny we see have diffused into a boundary layer in the scintillator. (The fraction of coincidences detected that originate from atoms embedded *within* the nylon film will be much smaller than the total, due to the extreme short range of the α particle.) The figure of 0.48 events/day thus represents at least about $\frac{1}{4}$ of all surface events; the true value is presumably in the range 0–4.5 events/day at 2σ .

Following the discussion on radon in CTF 1 of Section 6.4.1, suppose that this event rate corresponds only to radon diffusion from the water buffer. The radon concentration in the external water implied by the surface coincidence rate has an upper limit of 75 mBq/m^3 , quite a bit higher than the known value for CTF 3. (Near the walls of the tank, the radon activity is about 30 mBq/m^3 ; in water inside the shroud, it is more like 2 mBq/m^3 [169].) On the other hand, if it results only from radon emanation of the nylon film, the ²²²Rn production rate in the film has an upper limit of 20 mBq/m^3 , similar to the value of 18.2 mBq/m^3 measured for Sniamid nylon film. Recall that the CTF 3 film consists of C38F, not Sniamid; however, there is no reason to believe that the radioactive contaminant levels in the two materials would be quite different. It seems, then, that most of the surface radon in CTF 3 is probably produced by radon emanation from the nylon itself or by "dirt" adhering to the nylon film, rather than entering the scintillator via diffusion.

Light yield of events as a function of radial position

As mentioned above, the location of an event in the CTF vessel—energy of the event being taken as constant—has some effect on the number of photoelectrons detected. The most extreme cases are events whose energy spectrum includes a large number of γ rays, such as ²¹⁴Bi. For these events, the light loss is caused when a γ ray leaves the site of the event in a direction taking it out of the vessel into the water buffer. There it is invisible to the PMTs, as no scintillation light is produced. The probability of a γ ray escaping this way nears 50% for events near the vessel surface. Figure 9.8 shows the effect of such light loss on the ²¹⁴Bi spectrum.



Figure 9.8: The energy spectrum of ²¹⁴Bi coincidence events in Runs 2300–2563. Both spectra are cut off below 200 keV by one of the coincidence energy cuts. In part (a), at left, the spectrum is shown for events with a radial coordinate r < 65 cm. It nicely matches the theoretical prediction shown in Figure 7.18a. (This plot is not perfectly identical to Figure 8.6a; in that case, the data sample was selected by the average position of *both* events in the ²¹⁴BiPo coincidences.) In part (b), at right, the spectrum is shown for events with a radial coordinate r > 90 cm. The spectrum at right has a much higher percentage of events with a low observed energy, as many of its γ rays are lost to the water buffer.



Figure 9.9: Reconstructed energy of the ²¹⁴Po α decay as a function of the reconstructed event radial coordinate. For $r < 70 \,\mathrm{cm}$ or so, the observed event energy is constant. Beyond that, up to 20% of the event light is lost. The reasons for the light loss are unclear.

Nevertheless, even events that produce no γ rays are affected to some extent. The energy peak of any monoenergetic α decay, for instance, appears to be at lower energies as the distance of the event from the detector center increases. To examine this phenomenon, we use ²¹⁴Po events once again. In order to obtain better statistics, the window of runs considered is expanded to Runs 2180–2563; runs toward the beginning of this period exhibit a high radon concentration due to the then-ongoing purification tests. ²¹⁴Po events are once again selected with the standard set of ²¹⁴BiPo coincidence cuts; the volume cut used earlier is now dropped. The mean energy of the ²¹⁴Po α peak as a function of radius is shown in Figure 9.9. Up to about 70 cm from the vessel center, it is fairly constant; beyond that, it decreases rapidly by 20%, from 780 keV down to 630 keV. It is also noteworthy that the width of the peak is defined by $\sigma_E \approx 60-70$ keV up to 80 cm, but beyond that the peak becomes significantly more broad, $\sigma_E \approx 90$ keV or more.

The cause of this phenomenon, which has also been described in reference [166], is not entirely clear. Somehow, light is being lost for events near the surface. It is possible that only events very near the vessel surface lose light, and the apparent light loss in the outer region of scintillator results from the fraction of true surface events that are erroneously reconstructed at those radial positions. This hypothesis would also explain the larger values of σ_E seen for the α energy peak near the vessel; they result from the overlap of the "normal" peak from internal events, and the lower-energy peak from surface events reconstructed in the wrong place. Perhaps the effect occurs as a result of light trapping, a consequence of the same total internal reflection of light that causes the "dark zones" described in Section 7.6.3. The dark zones have an inner radius of about 89 cm. However, the deformations of the CTF vessel from sphericity make this hypothesis seem unlikely. Alternatively, it may be simply that the reflectivity of the light cones attached to the PMTs is degrading, reducing the PMT efficiency for detecting photons that enter at a shallow angle. But if this was the case, we would expect the light yield function not to be so flat in the central region of the detector out to about 70 cm. No current hypothesis really seems to explain the light loss well.

Whatever the cause of the light loss near the surface of the CTF vessel, the effect is important to keep in mind when searching for peaks in the energy spectrum in that region.

9.2.2 ²¹²BiPo coincidences: evidence for thorium

To search for thorium, we use the ²¹²BiPo coincidences, detected with the same set of cuts defined in Section 8.2.2. Instead of having a 65-cm radial cut, we now look at the full data sample in Runs 2300–2563. A total of 637 candidate coincidences are observed. Of these, we expect that about 36 (only 5.6%) are internal events, judging by the observation of nine candidates within the central 65-cm radius of the detector in Section 8.2.2. Dividing by the time cut efficiency (the probability that a ²¹²BiPo pair has a coincidence time within the required $0.1-2 \,\mu s$ limits) of 78.4% yields an expected total of 813 ²¹²BiPo coincidences, both surface and internal, before taking the other cut efficiencies into account.

Consider what happens when the coincidence time cut is relaxed to include coincidences meeting all the other ²¹²BiPo cuts with $\Delta t < 5 \,\mu$ s. We may fit a histogram of the coincidence time to an exponential decay curve plus constant term. In the fit, the decay time is fixed to that of ²¹²Po, 431 ns. The fit is made only to events having 200 ns $< \Delta t < 5 \,\mu$ s; as seen in Figure 9.10, a significant fraction of coincidences with a coincidence time less than about 200 ns are lost. The results are a predicted number of 844 ± 39 ²¹²BiPo coincidences in total (matching the value of 813 obtained above), and a data contamination of 25 ± 12 accidental coincidences with delay times in the above range. From this we conclude that the number of accidental coincidences in the pool of candidate ²¹²BiPo coincidences with the more stringent time cut is about 9 ± 5, or 1.4%.

The so-called accidental coincidences also include very fast 214 BiPo coincidences, as the 214 Po mean life is too long on this scale to distinguish their decay time distribution from a constant function. Statistically, the number of 214 BiPo coincidences included in the 212 BiPo sample should be 68.7% [energy cut] × 2.5% [less restrictive time cut] times the 1099



Figure 9.10: Histogram of the coincidence times for candidate ²¹²BiPo coincidences in Runs 2300–2563. They were fit to an exponential decay plus constant term. Only candidates with 100 ns $< \Delta t < 2 \,\mu$ s were used in the main analysis, although those having Δt up to $5 \,\mu$ s are shown and fit in this histogram. The lower limit of the fit is at 200 ns due to an apparent inefficiency in the Group 2 trigger for faster coincidences. The fit results give an expected total of 844 ± 39 true ²¹²BiPo coincidences, and a value of 25 ± 12 accidental coincidences with Δt in the range 200 ns–5 μ s. The ²¹²Po mean life of 431 ns was kept fixed in the fit.

coincidences observed in the previous section, or 19; this is consistent with the value of 25 ± 12 .

The energy spectrum of the ²¹²Po candidate events has a roughly Gaussian peak, shown in Figure 9.11. However, the lower-energy side of the peak exhibits some noise. It is not clear whether these events are really ²¹²Po shifted to lower energies by some effect, or some other species. Note that the data sample should be free of the ⁸⁵Kr \rightarrow ^{85m}Rb \rightarrow ⁸⁵Rb coincidences, whose first event is a β decay with a Q-value of only 173 keV. In any case, a Gaussian curve fit to the part of the spectrum above 700 keV has a peak value at 810±10 keV and a width defined by $\sigma = 112 \pm 7$ keV. This peak value is 18% lower than the predicted quenched energy of 989 keV, as might have been expected from the observed ~ 20% light loss for ²¹⁴Po events at the vessel surface.



Figure 9.11: Energy spectrum of the 212 Po candidate events in Runs 2300–2563. The Gaussian curve shown, with a peak value at 810 keV, was only fit to the portion of the histogram above 700 keV. The origin of the lower-energy noise to the left of the Gaussian is not known.

Since we saw that about 5.6% of all the 637 observed ²¹²BiPo coincidence events are internal, while about 1.4% of them are actually accidental (including ²¹⁴BiPo) coincidences, roughly 93% of all observed coincidences passing all cuts really are ²¹²BiPo coincidences at the vessel surface. The number of ²¹²BiPo surface events we expect is therefore about $93\% \times 844 \approx 785$ events, divided by the efficiency of the other cuts (we guess this to be again on the order of 85%) and the ²¹²Bi \rightarrow ²¹²Po branching ratio of 64%, and multiplied by the factor of four for the decay visibility. Assuming secular equilibrium, this gives an estimated total of about 5800 ²³²Th decays on the vessel surface during the time period being analyzed.

In the 474 days of livetime covered by Runs 2300–2563, this number yields an event rate of about 12 thorium decays per day, or a total ²³²Th mass on the inner vessel surface of 34 ng. (Compare to the thorium equivalent mass of 9 μ g estimated for the SNO hot spot [171].) For comparison, such a mass of thorium evenly distributed through the scintillator fluid would correspond to 9×10^{-15} g/g, much higher than the actual value of $(1.1 \pm 0.4) \times 10^{-16}$ g/g derived in Section 8.2.2.



Figure 9.12: Histogram of the sine of the latitude of 212 Po events reconstructed in the range 80 cm < r < 120 cm. If thorium were evenly distributed over the vessel surface, this histogram would be constant. This is far from the case; the vast majority of 212 Po events occur on the southern hemisphere of the vessel. Notice how broad the peak in the southern hemisphere is when compared with the corresponding peak in the 214 Po angular distribution, Figure 9.5.

If this mass of thorium were evenly distributed in the CTF nylon vessel film (total mass approx. 7.2 kg), the thorium contamination by mass would be 4.7 ppt. This is not far from the value of 3.9 ± 0.5 ppt obtained for C38F nylon film by Tama Chemicals in 1999 [12]. However, we must remember that the method of coincidences essentially sees only α particles produced in a thin skin of the nylon film closest to the scintillator, only a fraction of the total vessel mass. Thus the observed coincidence rate is in fact much higher than would be expected due solely to thorium contamination in the C38F nylon. This observation implies that most of the ²¹²BiPo coincidences have a source other than radioimpurities in the vessel material.

An even more suggestive finding is that the ²¹²BiPo coincidences are *not* evenly distributed over the vessel. (To measure their positions, we use only that of the better-defined ²¹²Po α decay.) A histogram of the sines of the latitudes of candidate events (Figure 9.12) shows that the vast majority of them occur on the southern hemisphere of the vessel. This is much different from the case of ²¹⁴BiPo coincidences (Figure 9.5), which are largely uniformly distributed over the vessel surface area except for peaks at either pole of the vessel. The pattern seen in the ²¹²BiPo distribution by latitude strongly suggests that a layer of thoriumrich material has settled down to the bottom of the nylon vessel. This pattern has been confirmed independently in other analyses such as reference [166]. It is nice to see an independent confirmation, outside the ²¹⁰Po data, of a probable consequence of particulate falling out from the volume of scintillator.

The difference in shape between the radon and thorium surface distributions seen in Figures 9.5 and 9.12 also implies that, if the radon distribution results from particulate as well, the two sets of particulate have significantly different properties. The peak in this histogram at the south pole is much broader than the analogous peak in the sample of 214 BiPo coincidences. This suggests that most of the 214 BiPo coincidences belong to the same population as the 210 Po coincidences in the vertical column described in Section 8.3.5. These atoms are distributed in a fairly narrow column along the *z*-axis, from which material would be deposited only very close to the vessel south pole, rather than all across the southern hemisphere of film surface as for the 212 Po atoms.

9.3 ²¹⁰Po on the vessel

We analyze the ²¹⁰Pb daughter ²¹⁰Po for Runs 2190–2346, a period which begins immediately after the second silica gel column test, and ends more than a year before the appearance of the vertical column of events in the scintillator. The reason for this selection of events is that the surface events are easy to disentangle from internal events only if a more-or-less uniform distribution of internal events can be assumed.

For each block of approximately 10 runs in this period, the number of ²¹⁰Po events was estimated by accepting only events having an α/β discrimination parameter $\gamma_1^* < 0$ (91.2% estimated cut efficiency). Candidate events were restricted to have reconstructed radial coordinates in the range 80–120 cm. The energy spectrum of these events between 200–800 keV was fit to the sum of a Gaussian curve and an exponentially decaying curve. The integral under the Gaussian was divided by the estimated α/β cut efficiency to yield the total number of ²¹⁰Po events within the radial cut for each time period. In order to have greater sensitivity to the presence of the Gaussian curve, its mean and width were fixed to specific values in each fit, defined below, instead of being allowed to float.

Meanwhile, for each set of about 10 runs, the number of 210 Po events seen within a 65cm radial cut was multiplied by the 65-cm radial cut scale factor (1.073) and by the ratio (1.75) of the numbers of events expected to be reconstructed within the respective ranges r < 65 cm and 80 < r < 120 cm. (The raw numbers of 210 Po events seen within the 65cm radial cut were determined as described in Section 8.3.3, and plotted in Figure 8.13.) This value, the predicted number of internal events within the 80–120 cm radial cut, was subtracted from the corresponding number described in the previous paragraph, giving the estimated number of 210 Po surface events.

Due to this convoluted procedure required to extract the number of ²¹⁰Po surface events, it was not feasible to create a histogram of the ²¹⁰Po events as a function of $\sin \theta$ of the event, or a map of the surface distribution of the ²¹⁰Po events. One can create such plots for "²¹⁰Po-like" events meeting the cuts 200 < E < 550 keV and γ_1^* < 0. However, these plots incorporate a fairly large number of other species as well, and can therefore be misleading.

Earlier sections in this chapter have shown that the CTF vessel end cap region tends to be much hotter than the rest of the nylon vessel. It is therefore worthwhile to divide up the surface into four regions: the two polar regions, with latitudes $|\theta| > 60^{\circ}$, and two equatorial regions, with latitudes $-60^{\circ} < \theta < 0^{\circ}$ and $0^{\circ} < \theta < 60^{\circ}$. The surface areas of the polar regions are each 0.842 m^2 , and those of the equatorial regions are each 5.44 m^2 . In the spectral fits of the energy histograms to a Gaussian curve plus exponential decay curve, the Gaussian mean and σ were kept fixed, with different values for each region:



Figure 9.13: The decay of surface 210 Po on the CTF vessel. The graph covers the same period as Figure 8.14, and each data point represents about ten runs. Internal contamination has been statistically subtracted away to produce this plot. The blue (asterisked) data points represent the number of 210 Po surface events in the northern hemisphere, between latitudes 0° and 60° N. The line drawn through them is an exponential decay curve with the mean lifetime of the isotope and an initial activity of 240 events/day. The line is not a fit (it is drawn for illustrative purposes only) but nevertheless matches the data nicely. The red (not asterisked) data points, on the other hand, represent 210 Po surface activity in the southern hemisphere (latitudes between 0° and 60° S). Another exponential curve with the mean life of 210 Po has been drawn through the first four points (its initial activity is 280 events/day), but after that the southern surface activity declines rather faster than the expected decay curve. This behavior is quite puzzling. Nevertheless, both curves reach an eventual plateau of about 40 events/day, or 7 events/day/m².



Figure 9.14: The decay of surface 210 Po near the north and south poles of the CTF vessel. This graph was constructed in the same way as Figure 9.13. The blue (asterisked) data points represent the number of 210 Po surface events near the north pole, above latitude 60° N. The line drawn through them is an exponential decay curve with the mean lifetime of the isotope and an initial activity of 38 events/day. The red (not asterisked) data points represent 210 Po surface activity near the south pole. An exponential curve with the mean life of 210 Po and an initial activity of 57 events/day has been drawn through the first four points, but afterwards the south pole activity declines much faster than the expected decay curve. Both curves reach an eventual plateau of 6–7 events/day, or 7–8 events/day/m².



Figure 9.15: A typical spectral fit to the ²¹⁰Po peak (Gaussian peak plus exponentially decaying background). This fit is for Runs 2200–2209, for events near the south pole of the CTF vessel. The number of events in the Gaussian peak is given by the fit as 801, and in the exponential background (within the range of the histogram) as 172.

- North polar region: $\bar{E} = 311 \text{ keV}, \sigma_E = 68 \text{ keV}$
- North equatorial region: $\bar{E} = 343 \text{ keV}, \sigma_E = 56 \text{ keV}$
- South equatorial region: $\bar{E} = 341 \text{ keV}, \sigma_E = 55 \text{ keV}$
- South polar region: $\bar{E} = 328 \text{ keV}, \sigma_E = 55 \text{ keV}$

These values were obtained by first allowing the parameters to float freely in fits of the first four data points (having the highest 210 Po activities) for each region. From the results of the light loss study on 214 Po, it is no surprise that the mean value of the 210 Po peak for surface events is about 15% less than the predicted quenched energy of 396 keV. The additional 5–10% light loss in the polar regions is most likely due to the opacity of the end caps. A typical spectral fit is shown in Figure 9.15.

The results of this analysis of surface polonium may be seen in Figures 9.13 and 9.14. The calculated ²¹⁰Po contamination in the two northern hemisphere regions (blue) closely follows the expected decay curve, with a mean lifetime of 200 days, for almost 300 days. After that, the curves plateau out to a constant value. In the two southern hemisphere regions (red), the contamination follows an exponential decay curve of the expected mean life for a short time (the first four data points, or 100 days). Afterwards, however, the number of 210 Po events in both data sets drops much faster than the isotope mean life. Despite this odd behavior, the southern hemisphere data also eventually reach constant plateaus. The equatorial regions tend towards constant values of 40 events/day, while the polar regions go to plateaus of about 6–7 events/day. These values all correspond to about 7 events/day per m² of nylon vessel inner surface.

These results have two surprising features. First, contrary to the pattern seen for other isotopes, the equilibrium contamination (eventual constant value) of ²¹⁰Po near the poles does not appear to be significantly higher than at other regions of the vessel. We may attribute this fact to the long lifetime of its progenitor ²¹⁰Pb: high radon levels near the two poles do not immediately translate into enhanced ²¹⁰Po concentrations. Instead, it seems that the surface of the nylon vessel is more or less uniformly covered with ²¹⁰Pb which eventually decays into ²¹⁰Bi and ²¹⁰Po at an approximately constant rate. This rate, a total of about 90 events/day over the entire inner surface, is to be compared with the observed ²¹⁴Po surface event rate of 0.33 ± 0.22 events/day. (Technically the former rate should be doubled, and the latter quadrupled, to take into account the geometric visibility factors.) The great disparity between the two rates means that the ²¹⁰Pb adhering to the vessel could not have originated from any source of radon in the CTF, but instead must have been carried in, for instance from contamination in the lines during scintillator filling.

The static nature of this contamination means that an adsorption equilibrium coefficient between nylon and pseudocumene may be deduced: k = S/Y, where Y is the internal contamination (per unit volume) and S is the surface contamination (per unit area) of ²¹⁰Po. The minimum internal rate of ²¹⁰Po was 1.6 ± 0.6 events/day/ton (from Section 8.3.4), or $Y = 1.4 \pm 0.5$ events/day/m³. When we double the observed surface contamination of 7 events/day/m² to account for the geometric visibility factor, we obtain a value for k (within about 50%) of 10 m. This is much greater than the value of 1.2 cm determined in a small-scale laboratory experiment [46], so we may conjecture that the two populations of ²¹⁰Po in the CTF are very far from reaching equilibrium. At equilibrium, a much greater amount of ²¹⁰Po should be present in the bulk of the scintillator. This further supports the hypothesis of little or no convective mixing in the scintillator fluid.

The second surprising feature is the transient period during which the two southern hemisphere regions of the CTF vessel lose their ²¹⁰Po contamination faster than the isotope can decay. This occurs roughly during the same period in which the internal ²¹⁰Po contamination falls faster than expected. This observation seems contradictory to the hypothesis of Section 8.3.3 that the fast rate of ²¹⁰Po decay seen within the central ton of scintillator was due to particulates falling out onto the vessel bottom. If that hypothesis was true, we would expect to see surface contamination on the southern part of the vessel increase with time, or at least decrease more slowly than the isotope's characteristic decay curve. Unfortunately we do not have a good model to explain these observations of surface decay; they currently remain one of the many mysteries of the CTF detector.

9.4 External gamma rays

Up to this point, a significant contribution to the CTF event rate has been mentioned only in passing. A large number of events in the CTF scintillator actually originate from γ rays produced outside the nylon vessel. These γ rays may travel dozens of cm in water or pseudocumene before being fully absorbed; the mean distance traveled by a 1.5 MeV γ in pseudocumene before it is scattered the first time, for instance, is 15–20 cm.

In this section, for specificity we investigate mainly the 40 K 1.46 MeV and 214 Bi 1.76 MeV γ rays. 214 Bi γ ray events detected in the CTF with an energy of about 1.76 MeV are almost guaranteed to have originated from the external water buffer. As seen in Figure 9.8a, the probability of seeing an internal 214 Bi decay with that specific energy is fairly small;

generally the energy of the electron emitted in the β decay is also seen. The probability that such an internal decay is in the pool of single events (not detected as a coincidence) is even smaller, due to the ²¹⁴Bi \rightarrow ²¹⁴Po \rightarrow ²¹⁰Pb chain. Finally, the rate of ²¹⁴Bi decays within the scintillator is already quite low, only a few per day. On the other hand, we are not guaranteed that a particular ⁴⁰K γ ray came from an external event. There very well may be internal potassium contamination. The ⁴⁰K decay by electron capture produces only the γ ray, with no additional particles to change the observed decay energy when the decay happens within the scintillator.

9.4.1 The spatial distribution of external γ rays

The spatial distribution of events in the CTF caused by a particular type of γ ray may be calculated in three ways. As mentioned in the introduction of this chapter, it may be approximated analytically by a crude derivation, or estimated through Monte Carlo simulations. With these two methods we will assume, for simplicity, that the source of all external ⁴⁰K γ rays is the nylon film of the CTF vessel. A third method is to use the observed distribution of one species of γ ray, and extrapolate it to another of similar energy. Below we will calculate, for each method, the fraction of external ⁴⁰K electron capture events expected to be reconstructed between distances r_i^{\min} and r_i^{\max} , for each of the radial cuts *i* defined in Section 8.5.2.

The simplistic approximation

Assume a very naive model of γ ray interaction with matter, as developed in Section 5.5.2, in which all the energy of each γ ray is deposited at the point where it first interacts with the scintillator. From Equation (5.57), in this model we expect that the fraction of surface γ ray events reconstructed to lie between distances r and $r + \delta r$ from the detector center should be equal to

$$f(r)\,\delta r \equiv \frac{1}{\sqrt{2\pi}} \frac{r\,\delta r}{\sigma\lambda R} \int_0^R \mathrm{d}r' \,\left[\mathrm{ei}\left(-\frac{R+r'}{\lambda}\right) - \mathrm{ei}\left(\frac{r'-R}{\lambda}\right)\right] \,e^{-\frac{r^2+r'^2}{2\sigma^2}} \,\sinh\frac{rr'}{\sigma^2}.$$
 (9.3)

The values R, σ and λ are parameters of the model. For the CTF, R = 1 m, and we suppose that σ and λ , the detector position resolution and γ ray attenuation length, respectively, are about 12 cm and 20 cm. This function, when integrated in r numerically over the radial cuts defined in Section 8.5.2, yields the following values:

- For region Ia (r < 39.7 cm): $\int f(r) dr \approx 0.00166$
- For region Ib (39.7 cm < r < 50.0 cm): $\int f(r) dr \approx 0.00298$
- For region I (r < 50.0 cm): $\int f(r) dr \approx 0.00464$
- For region II (50.0 cm < r < 63.0 cm): $\int f(r) dr \approx 0.01124$
- For region III (63.0 cm < r < 72.1 cm): $\int f(r) dr \approx 0.02163$
- For all $r \ge 0$: $\int f(r) dr \approx 0.45$

Graphs of this reconstructed radial function f(r) and the actual radial function $f_r(r)$ from which it is derived are shown in Figure 9.17a.

Monte Carlo simulations

For the Monte Carlo approximation, we simulate a uniform set of 1.46 MeV γ rays emitted isotropically from one point of the CTF nylon vessel. (It is only necessary to consider a single point source, since we are looking at radial distributions.) Each γ ray is followed on its path through water or scintillator until completely absorbed. Only those which deposit more than 95% of their energy (1.387 MeV) within the scintillator are considered further.



Figure 9.16: Monte Carlo simulation of the behavior of 40 K electron-capture γ rays originating on the vessel surface in the CTF. At left (part a) is a diagram of the trajectories of 100 of these simulated γ rays. The image is a projection onto the *xz*-plane, so some of the trajectories that appear to remain within the volume of scintillator actually are leaving it out of the plane of the page. At right (part b) is a histogram of the amount of energy deposited within the scintillator by each of 10⁶ simulated γ rays. From left to right, readily apparent features of the simulated energy spectrum are the back-scatter peak (50 keV), Compton edge (1.4 MeV), and photopeak (1.46 MeV). The spectrum shown does not include the effects of the finite CTF energy resolution, which would tend to smooth out these features. Note that the *y*-axis is logarithmic.

Consider a simulated γ ray which deposits energy E_i at N different locations \mathbf{x}_i within the scintillator. Its "center-of-light" is defined to be

$$\bar{\mathbf{x}} \equiv \frac{\sum_{i=1}^{N} E_i \mathbf{x}_i}{\sum_{i=1}^{N} E_i}.$$
(9.4)

The histogram of the radial coordinates of all such centers-of-light is taken to represent $f_r(r)$, the radial distribution function of the event's "position." (The word "position" is within quote marks here due to its ill-defined nature for external γ ray events.) As in Section 5.5.2, we presume that the result of position reconstruction is to randomly reposition the centerof-light by convoluting its radial probability distribution function with a Gaussian function, yielding the reconstructed distribution function in r,

$$f_d(r) = \frac{r}{\sqrt{2\pi\sigma^2}} \int_0^R \mathrm{d}r' \, \frac{f_r\left(|r'|\right)}{r'} \, e^{-\frac{r^2 + r'^2}{2\sigma^2}} \, \sinh\frac{rr'}{\sigma^2}.$$
(9.5)



Figure 9.17: Two models of the radial distribution function of the center-of-light of ⁴⁰K electron-capture γ rays originating at the surface of the CTF nylon vessel. The center-of-light is defined in Equation (9.4) to include only energy deposits made by the γ ray within the scintillator. In each plot, the sharper peak represents the actual radial positions $f_r(r)$ of the simulated centers-of-light for the γ rays. The broader peak is the predicted radial distribution function f(r) of their reconstructed positions, calculated by convolution of the actual distribution function with a Gaussian resolution function of $\sigma = 12$ cm. At left (part a), these functions are shown for the naive model in which all the energy of each γ ray is deposited at one point, with a mean γ ray travel distance of $\lambda = 20$ cm. At right (part b), these functions are shown for a simulation made with the GEANT 4 Monte Carlo libraries of one million γ rays.

This assumption is slightly flawed. It would be a good assumption if the CTF reconstruction software only used the pattern of PMT hits to determine the location of an event. However, the reconstruction software also uses knowledge of the timing of the event. The different energy deposits made by the γ ray occur at different times as the photon is scattered around. A better simulation would also incorporate the scintillation light produced at each γ ray point of scattering, as well as the PMTs, and simulate the hit times registered by each TDC of the detector; these values would be fed into the reconstruction software as if they were real data. Such a complete simulation of the CTF detector would be a much more complicated undertaking than the present Monte Carlo.

The present simulation was done with GEANT 4, version 8.0 [173]. One million 1.46-MeV γ rays were generated, originating at the north pole of a 1-m radius sphere of pseudocumene suspended in water, and given randomly chosen isotropic momenta. They were tracked, and

for those that deposited more than 1.387 MeV of energy within the scintillator, a histogram was constructed of the radial positions of their centers-of-light. This histogram, $f_r(r)$, is shown in Figure 9.17b, together with the predicted reconstructed distribution function f(r) after convolution.

In the simulation, 65.6% of the γ rays deposited at least a little energy within the scintillator (Figure 9.16). However, only 25.4% of them deposited at least 95% of their energy in the pseudocumene. Even so, the photopeak in the simulated energy spectrum was still very sharp, far more prominent than any other feature in the spectrum.

When the function f(r) constructed using this model was numerically integrated over the radial cuts defined in Section 8.5.2, the following values were obtained:

- For region Ia: $\int f(r) dr \approx 0.00258$
- For region Ib: $\int f(r) dr \approx 0.00479$
- For region I: $\int f(r) dr \approx 0.00737$
- For region II: $\int f(r) dr \approx 0.01727$
- For region III: $\int f(r) dr \approx 0.02756$
- For all $r \ge 0$: $\int f(r) dr \approx 0.25$

The ${}^{214}\mathrm{Bi}\ \mathbf{1.76}\ \mathrm{MeV}\ \gamma$ ray

For reasons mentioned above, we expect that the observed ²¹⁴Bi 1.76 MeV γ -ray peak is produced only by external events. Because the energies and stopping distances of this γ ray and the 1.46 MeV ⁴⁰K γ ray in pseudocumene are similar, the amplitudes of this peak within the different radial cuts provide an approximation to the values of $\int f(r) dr$ for external ⁴⁰K events. These values are therefore given roughly by the following, tabulated earlier in Table 8.9, although the normalization constant k is unknown:

- For region Ia: $\int f(r) dr \approx 65k$
- For region Ib: $\int f(r) dr \approx 94k$
- For region I: $\int f(r) dr \approx 159k$
- For region II: $\int f(r) dr \approx 223k$
- For region III: $\int f(r) dr \approx 215k$

For specificity we will take $k = 6.45 \times 10^{-5}$, so that the integral of f(r) over regions I and II matches the value of 0.02464 obtained in the Monte Carlo simulation.

9.4.2 Determining the internal ⁴⁰K contamination

In order to determine the internal ⁴⁰K contamination, we assume that the observed numbers of events in the photopeak in each of the radial cuts of Section 8.5.2 correspond to the sum of internal and external event spatial distributions. The internal distribution is supposed to be uniform over internal volumes (except near the vessel, due to edge effects which may be ignored here). The spatial distribution of γ rays produced by external events is more difficult to determine. As above, let the fraction of external ⁴⁰K photopeak γ -ray events reconstructed to lie within the radial shell $[r, r + \delta r]$ be $f(r) \delta r$. Then, for each radial cut *i*, the rate of events R_i seen within the cut is

$$R_i \equiv N_i/t = \mathcal{A}V_i + R_{\text{surface}} \int_{r_i^{\text{min}}}^{r_i^{\text{max}}} \mathrm{d}r f(r).$$
(9.6)

In this equation, t is the livetime, V_i is of course the volume (or perhaps, depending on the units in which \mathcal{A} is desired, the mass) of the cut, and R_{surface} is the total rate of ⁴⁰K γ decays on the vessel surface. (Only 45% of these decays are actually seen in the scintillator.) The values of $\int_i dr f(r)$ for each of the three models under consideration, in each of the radial cuts, have already been calculated above.

External	External only		External + internal		
event	$R_{\rm surface}$	$\chi^2/(4-1)$	$R_{\rm surface}$	\mathcal{A}	$\chi^2/(4-2)$
distribution	[ev/day]		[ev/day]	[ev/day/ton]	
Naive	48.5 ± 4.4	3.19	25.2 ± 8.8	0.80 ± 0.26	0.08
Monte Carlo	36.3 ± 3.2	1.82	21.2 ± 7.5	0.67 ± 0.30	0.10
214 Bi data	$51.6\pm4.6^*$	3.17	$55.4\pm29.5^*$	-0.11 ± 0.84	2.37

Table 9.1: Fit to the rate of 40 K electron capture decays on the vessel surface and in the CTF scintillator, using the three models for external event distribution described above. The data points used were for the four regions Ia, Ib, II, and III (the regions were defined in Section 8.5.2). The columns headed "External only" give the fit results when only an external contribution is assumed; those labeled "External + internal" are the results of fits to the full Equation (9.6). The asterisked values in the row for 214 Bi data are based on the assumption that the normalization constant k defined previously is equal to 6.45×10^{-5} .

These data have been fit to Equation (9.6) using the numerical integrals over f(r) obtained for each of the three methods. The results are presented in Table 9.1, along with the results from a fit only to an external background $[R_i = R_{\text{surface}} \int_i dr f(r)]$. The errors shown in the fit results reflect only the errors reported in Table 8.9 for the numbers of ⁴⁰K γ events, originating with the energy spectral fits, not any systematic errors in any of the methods.

The results of Table 9.1 are technically consistent within the errors quoted, but they have rather different implications. Use of the ²¹⁴Bi γ ray peak to represent the presumed radial distribution of external ⁴⁰K γ rays (last row of the table) implies only an upper limit on the internal ⁴⁰K contamination. This limit is 0.73 events/day/ton at 1 σ , corresponding to an upper limit of 610 events/day of pure ⁴⁰K β decays in the Borexino Fiducial Volume. On the other hand, use of the naive radial distribution for γ rays, or of the distribution simulated by Monte Carlo, implies a real measurable value ranging from 560±250 to 670±220 events/day in the Fiducial Volume.

It is worth noting that when the amplitudes of the supposed ²¹⁴Bi γ peak at 1.76 MeV in the four volumes are fit to Equation (9.6) using the three different models, the naive and Monte Carlo models yield a non-zero value for internal contamination with ²¹⁴Bi γ rays as well. (In both cases, the specific result is roughly 0.65 ± 0.30 events/day/ton. In the third model, using the supposed ²¹⁴Bi peak amplitudes as a reference, the result is of course zero by assumption.) This is a bit odd because the maximum observed internal rate of ²¹⁴BiPo coincidences during the period of Runs 2300–2563 is never more than 2 events/day/ton. Very few internal ²¹⁴Bi decays should leak into the data sample in the 1.76 MeV peak, because coincidences were specifically excluded from the sample. The likely possibilities are as follows:

- The naive and Monte Carlo radial distribution functions do not model the actual radial distribution of the reconstructed positions of external γ rays very well.
- The amplitudes of the ²¹⁴Bi peaks are, just through an unfortunate chance, not in line with what would be expected statistically.
- The detection of coincidences in the CTF is much less efficient than generally supposed.
- The peak near 1.76 MeV is not actually the $^{214}\text{Bi}~\gamma$ ray, but instead is from some other isotope.

Of these hypotheses, the first two above are most probable (the other two are mentioned for the sake of completeness). The first hypothesis implies that the ²¹⁴Bi radial fit should be taken most seriously. The second, however, leads to a directly opposite conclusion. The second hypothesis is supported by the observation that the number of ²¹⁴Bi 1.76-MeV events measured in the next-to-outermost radial shell (223 ± 86 events) is greater than the number measured in the outermost shell (215 ± 58 events), whereas the situation for the ⁴⁰K peak is the other way around (318 ± 48 events versus 428 ± 53 events); refer to Table 8.9. Unfortunately, if the second hypothesis is true, contamination of the Borexino scintillator by ⁴⁰K may potentially be a large problem. Further studies are recommended.

Concluding Remarks

In this work, we have explored in depth many types of radioactive contamination that are present in the Counting Test Facility of Borexino, and may also appear to a lesser extent in the full-scale experiment. As we noted, the CTF is not sufficiently sensitive to say definitively that the scintillator it uses is pure enough for Borexino. Nevertheless, study of the CTF backgrounds is invaluable in its own right, as much for discovering and testing new techniques of background measurement as for setting radiopurity limits on individual isotopes.

Some of the techniques described in this work have been specific to the CTF—methods to deal with a two index of refraction system and with timing channel coverage of two PMTs during reconstruction of event positions, for instance. Even in this case, some lessons may be learned. Many times the tradeoff between cost and ease of detector construction versus the resulting complexity of data analysis, though it may seem better to favor the former at first, will yield data that are prohibitively difficult to analyze well. This restricts what can be discovered by the detector. Fortunately, most of these issues will not be present in the full-scale Borexino experiment.

Other techniques, mostly those at a higher level of analysis, are applicable to all scintillationbased detectors. This is true, for instance, of the analysis of the position reconstruction accuracy performed in Chapter 5. More importantly, it is the case for the techniques of particle identification described in Chapter 8. Many of these make it possible to remove background noise from the data sample individually, rather than using a statistical subtraction (which always worsens statistical errors in the final results). Perhaps the most important new technique developed in this work will prove to be the likelihood-based method for tagging radon daughter events, in particular the β decay of the isotope ²¹⁴Pb which would be difficult to distinguish otherwise.

Future improvements of the method may also be extended to isotopes in the thorium decay chain. In particular, the isotopes ²¹²Pb and ²⁰⁸Tl could conceivably be tagged through this technique, significantly reducing background in the ⁷Be and ⁸B neutrino energy domains, respectively. Given the 11-hour half-life of ²¹²Pb, tagging thorium daughter isotopes will also make it possible to detect scintillator convection with much better sensitivity than observations of the radon daughters (separated by on the order of one hour) can provide. It is hoped that such methods may be put to good use in Borexino when it comes online, as well as in later neutrino, dark matter, and neutrinoless double β decay experiments.

Relatively small-scale detectors such as the CTF may in the future act as radiopurity testing facilities, capable of measuring the radioactivity levels of materials to be used in larger experiments at unprecedented sensitivities. The CTF itself, in fact, performed the first measurement ever made of ¹⁴C in a petroleum derivative [149]—and this isotope makes by far the greatest contribution to the event rate measured in the facility! The ultimate sensitivity of a CTF-like detector to ¹⁴C, if a hypothetical organic scintillator with a much lower level of the isotope could be found, would be on the order of one event/day in the central ton of the detector, for a mass fraction of about 10^{-22} g (four atoms!) of ¹⁴C per gram of carbon. Sensitivities to uranium and thorium in a 5-kg sample, assuming secular equilibrium, are on the order of one part per trillion, on par with that of inductive coupled plasma mass spectroscopy. The sensitivity could be increased further with a larger sample or a longer counting period. Already there are ideas in the air to construct a ~1 ton liquid xenon scintillation detector to be used as a materials testing facility. Due to the increased scintillation yield of Xe with respect to organic liquids, energy and position resolutions are much improved, and its sensitivity could be even better.
Bibliography

- S. Eidelman, et al., "Review of particle physics", *Phys. Lett. B* 592 (2004) 1.
 URL http://pdg.lbl.gov
- [2] C. Tully, "The next energy frontier", seminar, Princeton University (September 2005).
- [3] Y. Grossman, H. J. Lipkin, "Flavor oscillations from a spatially localized source: a simple general treatment", *Phys. Rev. D* 55 (1997) 2760.
- [4] A. Strumia, F. Vissani, "Implications of neutrino data circa 2005", Nucl. Phys. B 726 (2005) 294.
- [5] J. N. Bahcall, A. M. Serenelli, S. Basu, "10,000 Standard Solar Models: a Monte Carlo simulation", arXiv:astro-ph/0511337 (2005).
- [6] L. Wolfenstein, "Neutrino oscillations in matter", Phys. Rev. D 17 (1978) 2369.
- [7] S. P. Mikheyev, A. Yu. Smirnov, "Resonant amplification of ν -oscillations in matter and solar-neutrino spectroscopy", *Nuovo Cimento C* **9** (1986) 17.
- [8] A. Yu. Smirnov, "The MSW effect and matter effects in neutrino oscillations", arXiv:hep-ph/0412391 (2004).
- [9] A. Yu. Smirnov, "The MSW effect and solar neutrinos", arXiv:hep-ph/0305106 (2003).
- [10] J. N. Bahcall, A. M. Serenelli, S. Basu, "New solar opacities, abundances, helioseismology, and neutrino fluxes", Astrophys. J. Let. 621 (2005) 85L.

- [11] N. Grevesse, A. J. Sauval, "Standard solar composition", Space Science Rev. 85 (1998)
 161.
- [12] L. Cadonati, The Borexino Solar Neutrino Experiment and its Scintillator Containment Vessel, Ph.D. thesis, Princeton University (2001).
- [13] L. Bornschein, "The KATRIN experiment a direct measurement of the electron antineutrino mass in the sub-eV region", Nucl. Phys. A 752 (2005) 14c.
- [14] C. Kraus, et al., "Final results from phase II of the Mainz neutrino mass search in tritium β decay", *Eur. Phys. J. C* **40** (2005) 447.
- [15] K. Assamagan, et al., "Upper limit of the muon-neutrino mass and charged-pion mass from momentum analysis of a surface muon beam", *Phys. Rev. D* 53 (1996) 6065.
- [16] R. Barate, et al., "An upper limit on the τ neutrino mass from three- and five-prong tau decays", *Eur. Phys. J. C* **2** (1998) 395.
- [17] S. Hannestad, J. Madsen, "Neutrino decoupling in the early universe", *Phys. Rev. D* 52 (1995) 1764.
- [18] W. L. Freedman, et al., "Final results from the Hubble Space Telescope key project to measure the Hubble constant", Astrophys. J. 553 (2001) 47.
- [19] D. N. Spergel, et al., "First-year Wilkinson Microwave Anisotropy Prove (WMAP) observations: Determination of cosmological parameters", Astron. J. Suppl. 148 (2003) 175.
- [20] P. Crotty, J. Lesgourgues, "Current cosmological bounds on neutrino masses and relativistic relics", *Phys. Rev. D* 69 (2004) 123007.
- [21] G. L. Fogli, et al., "Observables sensitive to absolute neutrino masses: constraints and correlations from world neutrino data", *Phys. Rev. D* 70 (2004) 113003.

- [22] D. D'Angelo, Towards the detection of low energy solar neutrinos in BOREXino: data readout, data reconstruction and background identification, Ph.D. thesis, Technische Universität München (2006).
- [23] E. Akhmedov, M. Frigerio, A. Smirnov, "Probing the seesaw mechanism with neutrino data and leptogenesis", *JHEP*09 (2003) 021.
- [24] R. Fardon, A. Nelson, N. Weiner, "Dark energy from mass varying neutrinos", arXiv:astro-ph/0309800 (2003).
- [25] V. Barger, P. Huber, D. Marfatia, "Solar mass-varying neutrino oscillations", arXiv:hep-ph/0502196 (2005).
- [26] V. Barger, D. Marfatia, K. Whisnant, "Confronting mass-varying neutrinos with Mini-BooNE", arXiv:hep-ph/0509163 (2005).
- [27] A. Aguilar, et al., "Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_{\mu}$ beam", arXiv:hep-ex/0104049 (2001).
- [28] M. Maltoni, et al., "Ruling out four-neutrino oscillation interpretations of the LSND anomaly?", Nucl. Phys. B 643 (2002) 321.
- [29] A. Strumia, "Interpreting the LSND anomaly: sterile neutrinos or CPT violation or...?", Phys. Lett. B 539 (2002) 91.
- [30] H. A. Tanaka, "Update on MiniBooNE", Nucl. Phys. B (Proc. Suppl.) 149 (2005)
 122.
- [31] S. Jullian, "Experimental status on neutrinoless double beta decay searches", Comptes Rendus Physique 6 (2005) 778.
- [32] H. V. Klapdor-Kleingrothaus, "First evidence for lepton number violation and of the Majorana character of neutrinos", Nucl. Phys. B (Proc. Suppl.) 145 (2005) 219.
- [33] D. O. Caldwell, P. A. Sturrock, "Evidence for solar neutrino flux variability and its implications", Astropart. Phys. 23 (2005) 543.

- [34] C.-S. Lim, W. J. Marciano, "Resonant spin-flavor precession of solar and supernova neutrinos", *Phys. Rev. D* 37 (1988) 1368.
- [35] A. Friedland, C. Lunardini, C. Peña-Garay, "Solar neutrinos as probes of neutrinomatter interactions", *Phys. Lett. B* 594 (2004) 347.
- [36] The Borexino Collaboration, C. Arpesella, et al., Proposal of Borexino: a Real Time Detector for Low Energy Solar Neutrinos, Vol. 1, University of Milano and Istituto Nazionale di Fisica Nucleare, 1991, initial Borexino proposal.
- [37] M. Deutsch, "Proposal for a cosmic ray detection system for the Borexino solar neutrino experiment", proposal submitted to the National Science Foundation (1996).
- [38] L. Oberauer, S. Schönert, "Status report on the Borexino Muon Identification System", Borexino Collaboration internal report (February 1997).
- [39] T. Hagner, et al., "Muon-induced production of radioactive isotopes in scintillation detectors", Astropart. Phys. 14 (2000) 33.
- [40] C. Galbiati, et al., "Cosmogenic ¹¹C production and sensitivity of organic scintillator detectors to *pep* and CNO neutrinos", *Phys. Rev. C* 71 (2005) 055805.
- [41] M. Balata, et al., "CNO and *pep* neutrino spectroscopy in Borexino: measurements of the cosmogenic ¹¹C background with the Counting Test Facility", arXiv:hepex/0601035; to be submitted to *Phys. Rev. C* (2006).
- [42] D. Franco, The Borexino Experiment: Test of the Purification Systems and Data Analysis in the Counting Test Facility, Ph.D. thesis, Università degli Studi di Milano (2005).
- [43] G. Zuzel, H. Simgen, G. Heusser, "Ar and Kr concentrations in nitrogen as a measure of the ³⁹Ar and ⁸⁵Kr activities in connection with the solar neutrino experiment BOREXINO", Applied Rad. and Isotopes **61** (2004) 197.

- [44] A. Pocar, Low Background Techniques and Experimental Challenges for Borexino and its Nylon Vessels, Ph.D. thesis, Princeton University (2003).
- [45] M. Wójcik, G. Zuzel, "Radon permeability through nylon at various humidities used in the BOREXINO experiment", Nucl. Instr. Meth. A 524 (2004) 355.
- [46] M. Leung, The Borexino Solar Neutrino Experiment: Scintillator Purification and Surface Contamination, Ph.D. thesis, Princeton University (2006).
- [47] M. Leung, "Surface contamination from radon progeny", AIP Conf. Proc. 785 (2005) 184.
- [48] L. Cadonati, J. C. Maneira, "External background from the light guides", Borexino Collaboration internal report (November 2000).
- [49] J. Benziger, et al., "The nylon scintillator containment vessels for the Borexino solar neutrino experiment", to be submitted to Nucl. Instr. Meth. A.
- [50] A. Ianni, D. Montanino, F. L. Villante, "How to observe ⁸B solar neutrinos in liquid scintillator detectors", *Phys. Lett. B* 627 (2005) 38.
- [51] C. G. Rothschild, M. C. Chen, F. P. Calaprice, "Antineutrino geophysics with liquid scintillator detectors", *Geophys. Research Let.* 25 (1998) 1083.
- [52] R. S. Raghavan, et al., "Measuring the global radioactivity in the Earth by multidetector antineutrino spectroscopy", *Phys. Rev. Let.* 80 (1998) 635.
- [53] L. Cadonati, F. P. Calaprice, M. C. Chen, "Supernova neutrino detection in Borexino", Astropart. Phys. 16 (2002) 361.
- [54] J. F. Beacom, W. M. Farr, P. Vogel, "Detection of supernova neutrinos by neutrinoproton elastic scattering", *Phys. Rev. D* 66 (2002) 033001.
- [55] The Borexino Collaboration, G. Alimonti, et al., "Science and technology of Borexino: a real-time detector for low energy solar neutrinos", *Astropart. Phys.* **16** (2002) 205.

- [56] O. Smirnov, personal communication.
- [57] I. Kaplan, "On the systematics of even-even alpha-emitters", Physical Review 81 (1951) 962.
- [58] P. Venkataramaiah, et al., "A simple relation for the Fermi function", J. Phys. G: Nucl. Phys. 11 (1985) 359.
- [59] J. P. Davidson, Jr., "The first forbidden shape factor and the f_nt products for betadecay", Phys. Rev. 82 (1951) 48.
- [60] H. Simgen, G. Heusser, G. Zuzel, "Highly sensitive measurements of radioactive noble gas nuclides in the BOREXINO solar neutrino experiment", *Applied Rad. and Isotopes* 61 (2004) 213.
- [61] W. C. Butterman, R. G. Reese, Jr., "Mineral commodity profiles: Rubidium", U. S. Geological Survey Open-File Report.
 URL http://pubs.usgs.gov/of/2003/of03-045/of03-045.pdf
- [62] R. W. Kay, "Elemental abundances relevant to identification of magma sources", Phil. Trans. R. Soc. London A 310 (1984) 535.
- [63] The Borexino Collaboration, C. Arpesella, et al., "Measurements of extremely low radioactivity levels in BOREXINO", Astropart. Phys. 18 (2002) 1.
- [64] J. B. Benziger, et al., A Proposal for Participation in the Borexino Solar Neutrino Experiment, Princeton University, 1996.
- [65] T. Araki, et al., "Experimental investigation of geologically produced antineutrinos with KamLAND", *Nature* 436 (2005) 499.
- [66] T. Araki, et al., "Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion", *Phys. Rev. Let.* 94 (2005) 081801.
- [67] L. S. Niedermeier, High Efficiency Purification of Liquid Scintillators for the Solar Neutrino Experiment Borexino, Ph.D. thesis, Technische Universität München (2005).

- [68] D. Manuzio, Towards the detection of subMeV solar neutrinos in Borexino: data reconstruction and analysis tools, Ph.D. thesis, Università degli Studi di Genova (2004).
- [69] J. C. Maneira, Calibration and Monitoring for the Borexino Solar Neutrino Experiment, Ph.D. thesis, Universidade de Lisboa (2001).
- [70] "The Merck Index: 13th edition". URL http://themerckindex.cambridgesoft.com/
- [71] "Chemexper Chemical Database". URL http://www.chemexper.com/ccd/index.html
- [72] M. Göger-Neff, Development of a Liquid Scintillator and of Data Analysis Methods for BOREXINO, Ph.D. thesis, Technischen Universität München (2001).
- [73] The Borexino Collaboration, H. O. Back, et al., "Phenylxylylethane (PXE): a highdensity, high-flashpoint organic liquid scintillator for applications in low-energy particle and astrophysics experiments", arXiv:physics/0408032 (2004).
- [74] J. B. Birks, *Photophysics of Aromatic Molecules*, Wiley-Interscience, London, 1970.
- [75] M. C. Johnson, Scintillator Purification and Study of Light Propagation in a Large Liquid Scintillator Detector, Ph.D. thesis, Princeton University (1998).
- [76] F. D. Brooks, "Development of organic scintillators", Nucl. Instr. Meth. 162 (1979) 477.
- [77] D. Philips, "Substituent effects in the photochemistry of benzene vapor", J. Photochem. 1 (1972) 97.
- [78] D. L. Horrocks, "Scintillation efficiencies at high solute concentrations: Possible energy transfer from S₃ states of excited aromatic solvents", J. Chem. Phys 52 (1970) 1566.
- [79] I. B. Berlman, Handbook of Fluorescence Spectra of Aromatic Molecules, 2nd Edition, Academic Press, New York, 1971.

- [80] The Borexino Collaboration, G. Alimonti, et al., "Light propagation in a large volume liquid scintillator", Nucl. Instr. Meth. A 440 (2000) 360.
- [81] G. Ranucci, et al., "A sampling board optimized for pulse shape discrimination in liquid scintillator applications", *IEEE Trans. Nucl. Sci.* 51 (2004) 1784.
- [82] T. A. King, R. Voltz, "Time dependence of scintillation intensity in aromatic materials", Proc. Royal Soc. London A 289 (1966) 424.
- [83] G. Ranucci, A. Goretti, P. Lombardi, "Pulse-shape discrimination of liquid scintillators", Nucl. Instr. Meth. A 412 (1998) 374.
- [84] J. C. Maneira, "On the need for a quencher in the Borexino buffer", Borexino Collaboration internal report (September 2000).
- [85] M. Chen, et al., "Quenching of undesired fluorescence in a liquid scintillator particle detector", Nucl. Instr. Meth. A 420 (1999) 189.
- [86] H. O. Back, Internal Radioactive Source Calibration of the Borexino Solar Neutrino Experiment, Ph.D. thesis, Virginia Polytechnic Institute and State University (2004).
- [87] D. Gaiotto, "Light sources on the Borexino scintillator vessels", experimental project report, Princeton University (2001).
- [88] G. Vogel, "Gran Sasso Laboratory sees light at the end of the tunnel", Science 304 (2004) 1101.
- [89] E. Harding, C. Galbiati, T. Shutt, "Nylon vessels installation", Borexino Collaboration official procedure (September 2003).
- [90] G. Ranucci, personal communication.
- [91] V. Lagomarsino, G. Testera, "A gateless charge integrator for Borexino energy measurement", Nucl. Instr. Meth. A 430 (1999) 435.

- [92] M. E. Monzani, Characterization and Calibration of the Borexino Detector for Solar and Supernova Neutrinos, Ph.D. thesis, Université Denis Diderot—Paris VII and Università degli Studi di Milano (2005).
- [93] O. Ju. Smirnov, P. Lombardi, G. Ranucci, "Precision measurements of time characteristics of ETL9351 photomultipliers", *Instr. and Exp. Tech.* 47 (2004) 69.
- [94] R. Dossi, A. Ianni, "Thinking about the photomultiplier EMI9351 μ-metal shielding", Borexino Collaboration internal report (January 1998).
- [95] J. X. Prochaska, The Design and Fabrication of Optimal Light Collectors for the CTF Upgrade, Senior thesis, Princeton University (1993).
- [96] L. Oberauer, et al., "Light concentrators for Borexino and CTF", Nucl. Instr. Meth. A 530 (2004) 453.
- [97] M. Deutsch, et al., "The Borexino muon detector and muon induced backgrounds", in: E. Norman, L. Schroeder, G. Wozniak (Eds.), Nuclear Physics in the 21st Century, Vol. 610 of AIP Conference Proceedings, 2001, p. 973.
- [98] K. McCarty, "Testing and programming the Borexino Outer Muon Detector", general project report, Princeton University (2002). URL http://www.princeton.edu/~kmccarty/papers/advproject.pdf
- [99] A. Razeto, Events Readout for a Real-Time Neutrino Detector, Ph.D. thesis, Università degli Studi di Genova (2002).
- [100] F. Gatti, et al., "The Borexino read out electronics and trigger system", Nucl. Instr. Meth. A 461 (2001) 474.
- [101] F. Gatti, et al., "DAQ electronics for Borexino experiment", Nucl. Phys. B (Proc. Suppl.) 78 (1999) 111.
- [102] ROOT: an object-oriented data analysis framework. URL http://root.cern.ch/

- [103] M. I. Kohan (Ed.), Nylon Plastics Handbook, Hanser/Gardner Publications, Inc., 1995.
- [104] E. de Haas, "The haze in nylon film: notes on telephone conversation with Dr. S. Porter of Honeywell", Borexino Collaboration internal report (May 2000).
- [105] Bayer Corporation. URL http://www.bayer.com/
- [106] E. de Haas, E. Harding, "Nylon extrusion in Pottsville, PA", Borexino Collaboration internal report (February 2001).
- [107] Honeywell Plastics. URL http://www.honeywell-plastics.com/
- [108] DuPont. URL http://www.dupont.com/
- [109] E. Harding, "Nylon film: Extrusion and cleanliness", Borexino Collaboration internal report (March 2001).
- [110] E. de Haas, "Visit to MF-Folien, Kempten, Bavaria", Borexino Collaboration internal report (March 2001).
- [111] E. de Haas, "Notes on telephone call with Dr. Walter Goetz of BASF, Ludwigshafen: nylon B4 vs B3", Borexino Collaboration internal report (October 2000).
- [112] BASF Group. URL http://www.basf.com/
- [113] mf-folien. URL http://www.mf-folien.de/
- [114] Military Standard 1246C, "Product cleanliness levels and contamination control program" (1994).

- [115] F. Rodriguez, Principles of Polymer Systems, 4th Edition, Francis & Taylor, 1996.
- [116] C. Galbiati, et al., "Glass transition, moisture content and mechanical characteristics of nylon", Borexino Collaboration internal report (2001).
- [117] ASTM D882, "Standard test method for tensile properties of thin plastic sheeting", ASTM International.
- [118] ASTM D638, "Standard test method for tensile properties of plastics", ASTM International.
- [119] D. J. Green, An Introduction to the Mechanical Properties of Ceramics, Cambridge University Press, 1998.
- [120] Arizona Instruments. URL http://www.azic.com/
- [121] OriginLab. URL http://www.originlab.com/
- [122] "Inner vessel for scintillator containment", Appendix 5 to an unspecified 1993 CTF report.
- [123] G. Zuzel, et al., "Ultra-traces of ²²⁶Ra in nylon used in the BOREXINO solar neutrino experiment", Nucl. Instr. Meth. A 498 (2003) 240.
- [124] Tama Chemicals Co. URL http://www.tama-chem.co.jp/english/
- [125] M. Wójcik, et al., "Radon diffusion through polymer membranes used in the solar neutrino experiment Borexino", Nucl. Instr. Meth. A 449 (2000) 158.
- [126] "Radon: Fate and transport of the agent in the environment". URL http://www1.umn.edu/eoh/hazards/hazardssite/radon/radonfate.html

- [127] "Effective radon diffusion coefficient". URL http://web.ead.anl.gov/resrad/datacoll/radon.htm
- [128] C. Galbiati, Data taking and analysis of the Counting Test Facility of Borexino, Ph.D. thesis, Università degli Studi di Milano (1998).
- [129] "Faq about uranium and radiation". URL http://www.cogema-inc.com/FAQs/uranium_and_radiation.htm
- [130] T. Shutt, "Radon plateout program", Borexino Collaboration internal report (2000).
- [131] T. Shutt, personal communication.
- [132] A. Pocar, "Low background techniques for the Borexino nylon vessels", AIP Conf. Proc. 785 (2005) 153.
- [133] C. Tsonopoulos, G. M. Wilson, "High-temperature mutual solubilities of hydrocarbons and water part I: benzene, cyclohexane and *n*-hexane", AIChE Journal 29 (1983) 990.
- [134] J. L. Heidman, et al., "High-temperature mutual solubilities of hydrocarbons and water part II: ethylbenzene, ethylcyclohexane, and *n*-octane", AIChE Journal 31 (1985) 376.
- [135] S. Fukuda, et al., "Determination of solar neutrino oscillation parameters using 1496 days of Super-Kamiokande-I data", *Phys. Lett. B* 539 (2002) 179.
- [136] J. Boger, et al., "The Sudbury Neutrino Observatory", Nucl. Instr. Meth. A 449 (2000) 172.
- [137] O. Ju. Smirnov, "Energy and spatial resolution of a large-volume liquid-scintillator detector", Instr. and Exp. Tech. 46 (2003) 327.
- [138] The Borexino Collaboration, G. Alimonti, et al., "Ultra-low background measurements in a large volume underground detector", Astropart. Phys. 8 (1998) 141.

- [139] D. N. McKinsey, K. J. Coakley, "Neutrino detection with CLEAN", Astropart. Phys. 22 (2004) 355.
- [140] Xmass Collaboration, "Xmass experiment II", in: Technique and Application of Xenon Detectors, 2001, p. 136.
- [141] K. J. Coakley, D. N. McKinsey, "Spatial methods for event reconstruction in CLEAN", Nucl. Instr. Meth. A 522 (2004) 504.
- [142] M. G. Boulay, A. Hime, J. Lidgard, "Design constraints for a WIMP dark matter and pp solar neutrino liquid neon scintillation detector", arXiv:nucl-ex/0410025 (2004).
- [143] G. Alimonti, et al., "A large-scale low-background liquid scintillator detector: the counting test facility at Gran Sasso", Nucl. Instr. Meth. A 406 (1998) 411.
- [144] G. Ranucci, "Time statistics of the photoelectron emission process in scintillation counters", Nucl. Instr. Meth. A 335 (1993) 121.
- [145] H. O. Back, et al., "Search for electron decay mode $e \rightarrow \gamma + \nu$ with prototype of Borexino detector", *Phys. Lett. B* **525** (2002) 29.
- [146] H. O. Back, et al., "New limits on nucleon decays into invisible channels with the BOREXINO counting test facility", *Phys. Lett. B* 563 (2003) 23.
- [147] H. O. Back, et al., "New experimental limits on heavy neutrino mixing in ⁸B-decay obtained with the Borexino Counting Test Facility", *Jour. Exp. Theor. Phys. Lett.* 78 (2003) 261.
- [148] H. O. Back, et al., "New experimental limits on violations of the Pauli exclusion principle obtained with the Borexino Counting Test Facility", *Eur. Phys. Journal C* 37 (2004) 421.
- [149] The Borexino Collaboration, G. Alimonti, et al., "Measurement of the ¹⁴C abundance in a low-background liquid scintillator", *Phys. Lett. B* 422 (1998) 349.

- [150] L. Cadonati, "Potassium peak in the CTF2 data", Borexino Collaboration internal report (2000).
- [151] E. Resconi, Measurements with the Upgraded Counting Test Facility (CTF-2) of the Solar Neutrino Experiment Borexino, Ph.D. thesis, Università degli Studi di Genova (2001).
- [152] M. E. Monzani, Metodi di Identificazione delle Particelle α/β in Grandi Volumi de Scintillatore Liquido, per la Reiezione del Fondo Radioattivo nell'Esperimento Borexino, Tesi di laurea, Università degli Studi di Milano (2000).
- [153] E. Gatti, F. D. Martini, "A new linear method of discrimination between elementary particles in scintillation counters", *IAEA Wien* (1962) 265.
- [154] T. Beau, Mesure des neutrinos solaires de la raie du béryllium dans l'expérience Borexino, Ph.D. thesis, Université Denis Diderot—Paris VII (2002).
- [155] O. Dadoun, Mesure des neutrinos de réacteurs nucléaires dans l'expérience Borexino, Ph.D. thesis, Université Denis Diderot—Paris VII (2003).
- [156] ANSI/IEEE Standard 754-1985, "Standard for binary floating point arithmetic".
- [157] The Borexino Collaboration, C. Arpesella, et al., Proposal of Borexino: a Real Time Detector for Low Energy Solar Neutrinos, Vol. 3, University of Milano and Istituto Nazionale di Fisica Nucleare, 1991, initial CTF proposal.
- [158] J. Benziger, F. P. Calaprice, et al., Borexino: a Real Time Detector for Low Energy Solar Neutrinos, Princeton University, 1992, CTF proposal submitted to the National Science Foundation.
- [159] J. Benziger, personal communication.
- [160] F. X. Hartmann, personal communication.
- [161] I. Manno, "Time of flight in Borexino", Borexino Collaboration internal report (September 1996).

- [162] CERN Program Library. URL http://cern.ch/cernlib
- [163] F. James, "MINUIT function minimization and error analysis: Reference manual".
- [164] B. Caccianiga, personal communication.
- [165] M. Neff, Untersuchungen zum Untergrund durch Radioaktivität und kosmische Strahlung im Sonnenneutrinoexperimentes BOREXINO, Diploma thesis, Technischen Universität München (1996).
- [166] O. Ju. Smirnov, "Analysis of slow decays in ²³²Th and ²³⁵U chains with CTF-III", Borexino Collaboration internal report (July 2005).
- [167] C.-H. Tsai, S. S. Q. Hee, "Permeation of alkylbenzene isomers of molecular weight 120 through nitrile gloves", *Journal Appl. Polymer Sci.* 60 (1996) 833.
- [168] O. Ju. Smirnov, "CTF-3 background calculator", Borexino Collaboration internal report (April 2005).
- [169] A. Ianni, personal communication.
- [170] A. Pocar, personal communication.
- [171] N. K. McCauley, Producing a Background Free Data Set for Measurement of the Charge Current Flux and Day-Night Asymmetry at the Sudbury Neutrino Observatory, Ph.D. thesis, The Queen's College, Oxford (2001).
- [172] M. Frigo, S. G. Johnson, "FFTW" (2004).
- [173] Geant 4. URL http://geant4.web.cern.ch/geant4/