

Searching for the $0\nu\beta\beta$ Decay of ^{130}Te with the Ton-Scale CUORE Bolometer Array

by

Alexey Drobizhev

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Yury G. Kolomensky, Chair

Professor Daniel N. McKinsey

Professor Eric B. Norman

Spring 2018

Searching for the $0\nu\beta\beta$ Decay of ^{130}Te with the Ton-Scale CUORE Bolometer
Array

Copyright 2018
by
Alexey Drobizhev

Abstract

Searching for the $0\nu\beta\beta$ Decay of ^{130}Te with the Ton-Scale CUORE Bolometer Array

by

Alexey Drobizhev

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Yury G. Kolomensky, Chair

CUORE—the Cryogenic Underground Observatory for Rare Events—is an experiment searching for the neutrinoless double-beta ($0\nu\beta\beta$) decay of ^{130}Te , based at the Laboratori Nazionali del Gran Sasso in Italy. The detector consists of 988 $5\times 5\times 5\text{ cm}^3$ TeO_2 crystals operated as bolometers at temperatures of $\sim 10\text{ mK}$ inside the world’s largest and most powerful dilution refrigerator. CUORE began physics data collection in the spring of 2017, and has recently released its first limit on the $0\nu\beta\beta$ decay half-life of ^{130}Te from $24\text{ kg} \cdot \text{y}$ isotope exposure (~ 2 months live time). This result— $T_{1/2}^{0\nu} > 1.5 \cdot 10^{25}\text{ y}$ (Bayesian) and $T_{1/2}^{0\nu} > 2.3 \cdot 10^{25}\text{ y}$ (Frequentist) at 90% C.L.—is the most stringent to date and, together with two alternative analyses necessary to calculate it, forms the centerpiece of this thesis. In the future, with five years of live time, CUORE is projected to reach a median sensitivity of $9 \cdot 10^{25}\text{ y}$ on this half-life. Besides the main physics conclusions, in this work I present an analysis modeling the spectral line shape of the bolometer, which is used for constructing the region of interest fit PDF. Additionally, I discuss the major CUORE hardware projects to which I have contributed in a significant way. Specifically, these are our world-leading cryostat, a cryogenic feedback temperature control system, a radon-free detector installation environment, and a room temperature detector calibration system.

To my father, Mikhail Anatolyevich Drobizhev.

Contents

Contents	ii
List of Figures	v
List of Tables	ix
1 Introduction	1
2 The Neutrino	3
2.1 Early History of the Neutrino	3
2.2 The Standard Model Neutrino	4
2.3 Massive Neutrinos	6
3 Double-Beta Decay	11
3.1 Two-Neutrino Double-Beta ($2\nu\beta\beta$) Decay	11
3.2 Neutrinoless Double-Beta ($0\nu\beta\beta$) Decay	13
3.3 Scientific Implications	14
3.4 Experimental Considerations	16
4 The CUORE Experiment	20
4.1 Introduction	20
4.2 Bolometric Detection	20
4.3 CUORE: a State of the Art Ton-Scale Bolometer Array	22
4.4 The CUORE Assembly Line	29
4.5 Data Taking, Production, and Low-Level Analysis	33
4.6 The First Data from CUORE	38
5 The CUORE Cryostat	41
5.1 Introduction	41
5.2 Dilution Refrigerator Theory and Operating Principles	42
5.3 Renderings of the CUORE Cryostat	50
5.4 Precooling the CUORE Cryostat	52
5.5 The Dilution Unit	56

5.6	Unique Structural Features of the CUORE Cryostat	58
6	A Cryogenic PID Temperature Controller	63
6.1	Introduction	63
6.2	PID Control Principles	64
6.3	Heaters and Thermometry for Temperature Control in CUORE	66
6.4	Room Temperature Hardware for a LabView-based PID Setup	69
6.5	LabView PID Controller Software	70
6.6	Testing and Usage	75
7	A Radon-Free Environment for CUORE Installation	78
7.1	Introduction	78
7.2	Description of the Setup	79
7.3	System Performance	84
8	A Room-Temperature Calibration System for CUORE	87
8.1	Introduction	87
8.2	External DCS Description	89
8.3	Installation Work	93
8.4	System Use and Testing	96
9	CUORE Analysis: Spectral Line Shape Modeling	98
9.1	Introduction	98
9.2	A First Look at the 2615-keV Line in CUORE Calibration Data	99
9.3	Fit Function	100
9.4	Fitting Technique	105
9.5	Results	110
10	CUORE Analysis: Fitting the ^{130}Te $0\nu\beta\beta$ Decay ROI	115
10.1	Introduction	115
10.2	Fit Function	116
10.3	UEML ROI Fitting Technique with RooFit	119
10.4	MCMC ROI Fitting Technique with BAT	121
10.5	ROI Fit Results	123
11	CUORE Analysis: Limits on the $0\nu\beta\beta$ Half-Life of ^{130}Te	128
11.1	Introduction	128
11.2	Profile NLL Curves from the RooFit Result	129
11.3	Bayesian and Frequentist Limits from the Profile NLL	132
11.4	Bayesian Limits with BAT	135
11.5	Conclusions and Comparison with Published Result	138
12	Conclusions	140

Bibliography	142
A Thermal Conductivity Measurements of PTFE and Al_2O_3 Ceramic at Sub-Kelvin Temperatures	153
A.1 Introduction	153
A.2 Experimental Setup	154
A.3 Results and Discussion	160
A.4 Conclusion	165
A.5 Technical Details	166
B The $A = 130$ Isobar	169
C List of Abbreviations and Acronyms	170

List of Figures

2.1	Feynman diagrams of β decay via Fermi's interaction and the standard model weak interaction.	5
2.2	Summary schematic of the Standard Model of particle physics.	6
2.3	KamLAND's neutrino oscillation discovery.	8
2.4	The normal and inverted neutrino mass hierarchies.	9
3.1	Feynman diagrams of the $2\nu\beta\beta$ and $0\nu\beta\beta$ decay processes.	11
3.2	Schematic representation of the nuclear mass isobars for isotopes with even and odd mass numbers.	12
3.3	Simple energy level diagram: β decay of ^{130}Te to ^{130}I is forbidden, but rare $\beta\beta$ decay to ^{130}Xe occurs.	13
3.4	Constraining the effective Majorana mass of the neutrino ($m_{\beta\beta}$) from the current CUORE limit on $T_{1/2}^{0\nu\beta\beta}(^{130}\text{Te})$ and target sensitivity, as well as measurements from other isotopes.	15
3.5	Cartoon $\beta\beta$ sum energy spectrum, neglecting background and detector effects. .	16
3.6	Summary of $0\nu\beta\beta$ decay candidate isotopes: $Q_{\beta\beta}$ versus natural abundance. . .	17
3.7	Cartoon experimental $\beta\beta$ spectra, with $\Gamma_{0\nu\beta\beta}/\Gamma_{2\nu\beta\beta} = 10^{-6}$ and no background, as measured by a ^{136}Xe experiment with $\sim 2.9\%$ FWHM energy resolution, and by a bolometric ^{130}Te experiment with CUORE's 7.7-keV FWHM energy resolution. 18	
4.1	Schematic drawing of a bolometric detector.	21
4.2	(a) Photograph of an NTD Ge thermistor. (b) Example of NTD resistance v. temperature measurement and fit from UC Berkeley.	25
4.3	Photographs of fresh glue spots deposited on an NTD and Si heater from the CUORE assembly line.	26
4.4	An example of a CUORE pulse.	26
4.5	Renderings of the CUORE bolometers: one 52-crystal tower, with zoom to one floor of 4 crystals held in PTFE supports in a copper frame; the full 19-tower, 988-channel array.	27
4.6	Summary of the CUORE background budget.	30
4.7	Photograph of the interior of the CUORE gluing glovebox, used for attaching NTDs and Si heaters to TeO_2 crystals.	31

4.8	Photographs of CUORE tower construction and NTD and Si heater wire bonding.	32
4.9	An example of a CUORE average pulse, before and after application of the Optimal Filter.	35
4.10	The decay chain of ^{232}Th —the source used in the CUORE DCS. We use six γ lines from ^{228}Ac , ^{212}Pb , and ^{208}Tl (bold) to calibrate the CUORE spectrum. . .	36
4.11	CUORE TeO_2 exposure versus date.	39
4.12	CUORE background (physics) spectrum from datasets 1 and 2, overlaid with the calibration spectrum scaled to the amplitude of the 2615-keV line.	40
4.13	CUORE blinded and unblinded background (physics) spectra near the ^{130}Te $0\nu\beta\beta$ decay ROI.	40
5.1	A phase diagram of a liquid ^3He – ^4He mixture, showing the separation between the ^3He -rich (concentrated) and -poor (dilute) phases below ~ 870 mK.	42
5.2	A schematic diagram of a single stage pulse tube refrigerator.	43
5.3	A schematic of a dilution refrigerator with “dry” pulse tube precool.	45
5.4	A schematic diagram of the Dilution Unit, showing the ^3He concentrated-dilute phase transition in the Mixing Chamber, the liquid/gas boundary in the Still, and the flow of ^3He	48
5.5	A three-dimensional rendering of the CUORE cryostat setup within the CUORE experiment “hut.”	50
5.6	A cutaway three-dimensional rendering of the CUORE cryostat, with key components and temperature stages/plates labeled.	51
5.7	A technical drawing of the Cryomech PT415 RM (remote motor) pulse tube cryocooler.	53
5.8	Photographs of one of CUORE’s Cryomech PT415 cold heads.	54
5.9	Photograph of the 10 helium lines for the CUORE cryostat’s five PTRs passing through the sandbox, for vibration isolation, on the ground floor of the CUORE hut.	55
5.10	(a) The cooling power of the CUORE DU, measured at the Mixing Chamber plate, as a function of temperature, from four different tests conducted at the Leiden facility and at LNGS. (b) A photograph of the CUORE DU with components and temperature stages labeled.	57
5.11	CUORE cryostat base temperature over a 12-hour period during a test run. . .	58
5.12	Side-view cutaway rendering showing the CUORE cryostat’s plate suspension and cold lead shields.	59
5.13	Side-view technical drawing and cutaway three-dimensional rendering of the CUORE cold lead shielding.	61
5.14	A comparison of the radioactivity of modern common lead, modern low-activity lead, and ancient Roman lead.	62

6.1	(a) A photograph of two PID heaters and one NTD thermometer, in their special copper holders, on the tower support plate (TSP). (b) A schematic of a CUORE thermometer or heater holder.	67
6.2	Mixing chamber and tower support plate (TSP) drawings with NTD thermometer and heater locations.	68
6.3	A block diagram of the room temperature hardware of the CUORE LabView-based backup PID system, with photographs of key components.	70
6.4	A flowchart schematic of the PID control software.	72
6.5	Screenshots of operating PID software in different configurations.	75
6.6	Temperature of the Still v. time during the brief impromptu PID test.	77
7.1	(a) Comparison of CUORE-0 and Cuoricino spectra, demonstrating background reduction. (b) Simulated CUORE energy spectrum from background budget study.	79
7.2	(a) Simplified diagram of the Radon Abatement System. (b) Schematic of the Radon Monitor.	80
7.3	(a) The decay of ^{222}Rn , from the ^{238}U decay chain, which is the primary source of the radon contamination we are mitigating. (b) Radon Monitor energy spectrum.	81
7.4	(a) A line drawing of the CUORE cleanroom facility showing the separate spaces. (b) Rendering of CR6, showing the soft walls, vestibule, seals in blue, two ULPA filter cabinets, as well as the cryostat and detector installation hardware inside.	82
7.5	Time response of the Radon Abatement System, CR6, and the Radon Monitor.	85
7.6	^{218}Po and ^{214}Po count rate v. time over the course of the entire ~ 4 month installation and cryostat closing procedure, with different operation phases labeled.	86
8.1	A simplified schematic diagram of the CUORE External DCS.	89
8.2	Simulated CUORE detection efficiency to the vertical sources from the External DCS.	91
8.3	(a) View of the OVC with PVC pipes installed. The curving path can be seen at the top. (b) Close up of one source string deployed in the “calibration” position in its pipe.	94
8.4	(a) 3-D rendering of winch locations above the cryostat and MSP. (b) Photograph of three winches in situ.	95
9.1	The CUORE calibration spectrum between 2530 and 2720 keV—the line shape fit domain.	99
9.2	The CUORE calibration spectrum between 2530 and 2720 keV—the line shape fit region, with the three regions we integrate to initialize the background shaded.	108
9.3	CUORE 2615 keV ^{208}Tl line shape fit result for one representative tower (number 14).	110
9.4	CUORE line shape fit results from two typical channels: 699 and 705.	111
9.5	Sum of the 19 tower-level line shape fit results overlaid with full (1811 channel-dataset pairs) data. The ratio of the data to the curve is plotted above.	112

9.6	Distributions of the full width at half maximum (FWHM) of the line shape fits.	113
10.1	CUORE effective energy resolution as a function of energy in datasets 1 and 2.	119
10.2	ROI RooFit UEML fit result. The blue vertical band shows $Q_{\beta\beta} = (2527.515 \pm 0.013)$ keV. $R_{\beta\beta}$ is allowed be negative. Pulls (residuals in units of σ) indicating the agreement between the fit and the data, are plotted above	124
10.3	ROI BAT MCMC fit result, allowing negative values of $R_{\beta\beta}$.	124
10.4	Fit agreement: comparing our BAT MCMC fit and the “official” result from [33] with our RooFit UEML fit.	125
10.5	ROI RooFit UEML fit result. $R_{\beta\beta}$ is constrained to positive values. Pulls (residuals in units of σ) indicating the agreement between the fit and the data, are plotted above.	126
10.6	ROI BAT MCMC fit result, constraining $R_{\beta\beta}$ to positive values.	126
11.1	The CUORE profile NLL curve, with and without systematic uncertainties, if $\Gamma_{0\nu\beta\beta}$ is allowed to take on negative values in the ROI fit.	129
11.2	The CUORE profile NLL curve, if $\Gamma_{0\nu\beta\beta}$ is constrained to positive values only in the ROI fit.	130
11.3	The profile NLL curves from the CUORE, CUORE-0, Cuoricino, and TTT experiments, together with their combination.	133
11.4	An illustration of integrating the profile likelihood of $\Gamma_{0\nu\beta\beta}$ (CUORE, stat. + syst.) to set a 90% CI Bayesian limit.	134
11.5	$\Gamma_{0\nu\beta\beta}$ posterior distribution from marginalization, with the fit parameter allowed to take on negative values.	136
11.6	$\Gamma_{0\nu\beta\beta}$ posterior distribution from marginalization, with the fit parameter allowed constrained to the positive, physical range.	137
11.7	Profile NLL curves derived from the posterior log-likelihoods of the two BAT fits—with and without constraining $\Gamma_{0\nu\beta\beta}$ to negative values.	138
12.1	The sensitivity of CUPID to the effective Majorana mass of the neutrino ($m_{\beta\beta}$), compared to CUORE and state of the art experiments using other isotopes.	141
A.1	Experimental setups for the alumina thermal conductivity measurements.	157
A.2	Experimental setup for the PTFE thermal conductivity measurements	158
A.3	PTFE thermal conductivity measurement results.	162
A.4	Alumina sample D single power law and Structure Scattering Model fits.	163
A.5	PTFE thermal conductivity measurement results.	164
A.6	Simulated temperature gradient along the alumina sample D at 100 mK.	167
A.7	Plot of applied power versus temperature as obtained from the 2-D thermal model and overlaid on the experimental data points for alumina sample D.	168
B.1	A detailed chart of the $A = 130$ isobar.	169

List of Tables

3.1	Experimentally significant $0\nu\beta\beta$ decay candidate isotopes: $Q_{\beta\beta}$ values, isotopic abundances, $0\nu\beta\beta$ $T_{1/2}$ limits, and $2\nu\beta\beta$ $T_{1/2}$ measurements.	18
4.1	Summary of CUORE efficiency for the first two datasets.	37
4.2	Summary of CUORE detector performance for the first two datasets.	38
7.1	CUORE RAS and cleanroom performance summary.	86
8.1	Summary of manufacturer-specified and measured thorium content of welding rods used for fabricating the calibration sources for the External DCS.	90
9.1	The relative amplitudes and energies of six Te X-rays that create our X-ray escape peak.	103
9.2	The ratios and energies of six Te X-rays that create our X-ray escape peak, as reconstructed by fitting the detector response in the CUORE-0 Monte Carlo model.	103
9.3	Summary of the CUORE spectral line shape fit parameters and their constraints.	109
10.1	Values of constant parameters in the ROI fit function.	118
10.2	ROI fit results, allowing negative values of $0\nu\beta\beta$ decay rate parameter.	123
10.3	ROI fit results, constraining the $0\nu\beta\beta$ decay rate parameter to positive values.	125
11.1	Systematic uncertainties in the CUORE $0\nu\beta\beta$ decay analysis.	130
11.2	90% CI/CL limits on the $0\nu\beta\beta$ half-life of ^{130}Te from RooFit profiled NLL.	135
11.3	Marginalized and profiled 90% CI limits on the $0\nu\beta\beta$ decay half-life of ^{130}Te from the BAT analysis.	137
A.1	Results for PTFE thermal conductivity.	161
A.2	Results for alumina thermal conductivity	165
A.3	Results for the PTFE scaled power fits.	166
A.4	Results for the alumina scaled power fits.	166

Acknowledgments

I would like to acknowledge my colleagues and the support staff at the University of California, Berkeley Department of Physics, Lawrence Berkeley National Laboratory, the Laboratori Nazionali del Gran Sasso, and the members of the CUORE Collaboration for their invaluable assistance and instruction. In particular, I want to thank Giovanni Benato, Laura Marini, and Jonathan Ouellet—the analysis presented in this thesis was a collaborative effort with them. I spent a long time getting invaluable skills and expertise in cryogenics working on the CUORE cryostat, and am very thankful to Carlo Bucci, Nic Chott, Paolo Gorla, Vivek Singh, and all of the cryostat team for everything I have learned from them and for the opportunity to contribute to that amazing machine. Furthermore, Jon, together with Tommy O'Donnell, played a big role in getting me started working on CUORE and bolometer R&D during their time at UC Berkeley.

I am thankful to the members of my thesis committee, Professors Daniel N. McKinsey and Eric B. Norman, for their valuable feedback and advice during the writing of this dissertation.

I am especially grateful to my advisor and mentor, Professor Yury G. Kolomensky, who has guided me through just about the entirety of my physics education and career thus far—since 2011, when I was a sophomore, and up to earning my PhD and getting a job in 2018. Equally, I am thankful to Joe Wallig, my colleague and friend who has taught me nearly everything I know about good engineering, and has always been my example of how to do honest and quality work.

I thank the late Professor Stuart J. Freedman for inspiring me to pursue my chosen profession of experimental nuclear physics.

This work was supported by the US Department of Energy (DOE) Office of Science under Contract No. DE-AC02-05CH11231, by the DOE Office of Science, Office of Nuclear Physics under Contract No. DE-FG02-08ER41551, and by the National Science Foundation under grant PHY-1314881. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or re-produce the published form of this manuscript, or allow others to do so, for United States Government purposes.

Chapter 1

Introduction

The CUORE (Cryogenic Underground Observatory for Rare Events) detector, which became operational in 2017 at the Laboratori Nazionali del Gran Sasso (LNGS) in Italy, is one of the newest devices searching for neutrinoless double-beta ($0\nu\beta\beta$) decay. CUORE looks for the $0\nu\beta\beta$ decay of ^{130}Te by means of a ton-scale array of $988\,5 \times 5 \times 5\,\text{cm}^3$ TeO_2 crystals operated as bolometers at temperatures near 10 mK. Via this theoretically postulated process, our experiment is among the most sensitive probes of the Majorana nature of the neutrino and lepton number violation currently active.

Over the course of the last five years, my doctoral research has primarily been devoted to this large-scale project. As such, CUORE also forms the core of this dissertation. Given that the experiment has been collecting physics data for less than a year as of this writing, my work on it has involved building and preparing the CUORE detector and the surrounding technical infrastructure, as well as performing analysis of our first data. On the whole, as tends to be the case with collaborations in nuclear and particle physics, these tasks involved dozens of people working on many diverse and important tasks. In this thesis, I try to focus specifically on those aspects of the CUORE effort to which I contributed directly in a meaningful way.

I begin by giving an overview of the history, as well as current open questions, of neutrino physics in Chapter 2. This is important for setting the context and conveying the significance of $0\nu\beta\beta$ decay research in general, and CUORE in particular. Next, in Chapter 3, I move on to double-beta decay itself, which needs a more detailed treatment. This chapter is primarily theoretical, describing the history and implications of the potential discovery of this process for physics beyond the Standard Model. I also introduce experimental considerations for a $0\nu\beta\beta$ decay search, which came into play when CUORE was designed and would be important for improving upon it in the future.

Chapter 4 is the first one explicitly discussing CUORE, and features a complete overview of the experiment. Its content contextualizes the work covered in detail by all of the following chapters, and explains how the detector functions, how it was built, how the data is processed, and introduces our first measurement. In the beginning of my work on CUORE, I was involved in characterizing the neutron transmutation doped (NTD) Ge thermistors (Sec-

tion 4.3). During detector assembly, I worked on coupling sensors and heaters to bolometer crystals and wire bonding those chips' electrical connections, as discussed in Section 4.4. The basic approach and software we use for data production, detailed in Section 4.5, is inherited from the predecessor CUORE-0 experiment. For CUORE, this had to be modified. In parallel with updating the low-level analysis, we also tune the detector hardware parameters and physical properties. Over the course of about three months prior to commencing our first physics runs, I served as the on- and off-site analysis expert.

In Chapter 5, I give a detailed description of the CUORE cryostat. It is the largest and most powerful dilution refrigerator that we know of in the world today, and also includes a wide range of novel technological solutions optimized for its unique role, including radiopure structural materials, cold lead shielding, and extensive vibration isolation. As a member of the CUORE Cryostat Team, I contributed to the installation and testing of a large portion of the hardware described in the chapter, particularly at later stages. I have also controlled and monitored the cryostat during cooldowns and base temperature operations. I implemented the cryogenic PID feedback temperature controller described in chapter 6 practically independently, during the testing phase in the fall of 2015.

Chapters 7 and 8 present two important pieces of the CUORE infrastructure: a radon-free cleanroom environment used during the installation of the detector in the cryostat, and an external calibration system, respectively. Both of these essentially self-contained projects are Lawrence Berkeley National Laboratory deliverables, headed by Joseph Wallig of the Engineering division who oversaw the designs and implementation. I played a significant role during the installation of both systems on-site at LNGS, including context-specific modifications and adaptations.

The subsequent three chapters contain my contributions to the main CUORE physics data analysis. Chapter 9 describes modeling and fitting the CUORE spectral line shape, which was implemented in collaboration with Laura Marini, an Italian PhD student from the University of Genoa. Chapters 10 and 11 describe fitting the ^{130}Te $0\nu\beta\beta$ decay region of interest in CUORE's first data, and extracting limits on the decay half-life from the fit results. I present results from two basic techniques: an unbinned extended maximum likelihood (UEML) fit, from which we set both Bayesian and Frequentist profiled limits, and a fully-Bayesian treatment with a Markov Chain Monte Carlo (MCMC) and marginalized limit. The former is a further development of my collaboration with Laura, while the latter was developed together with Giovanni Benato, who is a postdoctoral researcher at the University of California, Berkeley. All of these analyses form a major part of CUORE's first result published in.

While the majority of my PhD research contributes directly to CUORE, I have also put forth a significant amount of effort in next-generation bolometer R&D with Prof. Yury G. Kolomensky's group. This centers around our dilution refrigerator facility on the UC Berkeley campus, for which I was responsible for a number of years. I describe one of the measurements I performed in our cryostat—a low-temperature thermal conductivity measurement of Al_2O_3 ceramic—in Appendix A.

Chapter 2

The Neutrino

Пусть не поймает нейтрино
за бороду
И не посадишь в пробирку,-
Было бы здорово, чтоб
Понтекорво
Взял его крепче за шкуру!

В.С. Высоцкий

2.1 Early History of the Neutrino

The neutrino (ν) is probably the most exotic form of matter that has been experimentally confirmed to exist. The only fundamental neutral fermion, this “ghost particle” has led to unexpected discoveries. In the past twenty years, the neutrino sector has been the only one to violate the predictions of the otherwise highly successful “Standard Model.”

Since the early years of nuclear physics, the study of the neutrino is firmly connected to the β decay process. This form of nuclear decay was observed to produce electrons, which were readily detected. In other types of decay— α and γ —the emitted helium nuclei or energetic photons are monoenergetic. Their energy corresponds neatly to the difference in the mass of the mother and daughter atoms. Conversely, β spectra are continuous distributions between zero and the “Q-value,” or full decay energy. This phenomenon puzzled physicists in the early 20th century because, assuming the electron and daughter nucleus were the only decay products, it appeared to be a violation of the fundamental laws of conservation of energy, momentum, and spin. Some, like Niels Bohr, suggested that energy may not be strictly conserved, but interestingly the β particles never had *extra* energy—only a lack thereof. Wolfgang Pauli suggested an explanation in the 1930, proposing the existence of an undetectable “neutron”: an electrically-neutral fermion carrying away the missing kinetic energy and momentum. Famously, Pauli was concerned by the idea of introducing a new

particle that seemed to be unobservable experimentally, saying “I have done something very bad today by proposing a particle that cannot be detected; it is something no theorist should ever do.”

As we know it today, the name “neutron” belongs to an uncharged nucleon slightly heavier than a proton and found in every nucleus heavier than ^1H . The model of the nucleus as a dense agglomeration of similar numbers of neutrons and protons was conceived of by Werner Heisenberg around the same time. Early on, Pauli was uncertain if this particle was distinct from the mysterious carrier of missing energy in β decays. It almost seemed preferable, in terms of simplicity and elegance, for these two fermions to be one and the same. The neutron was discovered experimentally by James Chadwick in 1932, as a heavy nucleon, while Pauli had pointed out that the particle emitted during β decay should be as light or lighter than the electron. The first largely complete theory of what are now known to be “weak” interactions was published by Enrico Fermi in 1934 [1]. Here, the process of β decay was defined as the conversion of a neutron to a proton in the nucleus, with the release of an electron and a tiny (*anti*)neutrino $\bar{\nu}$ —Italian for “little neutral one.” The theory incorporates corresponding antiparticles, including Paul Dirac’s positron (anti-electron) and a neutrino, and allows flipping the process in various ways—electron capture, β^+ decay, etc.

Fermi’s theory did not immediately gain significant traction, but went on to be strongly favored by experiments. Energy was, in fact, strictly conserved—excellent news! Significantly for experiments, the model provides a quantified interaction mechanism for the neutrino. Fermi’s interaction, still lacking the modern conception of gauge bosons as mediators, occurs at a point, and is described by an empirical coupling constant $G_F = 1.16 \cdot 10^{-5} \text{ GeV}^{-2}$ (see Fig. 2.1a). Contrary to Pauli’s initial concerns, the elusive particle could be detected. Having a sufficiently strong antineutrino flux and a target with a large number of protons, one only needs to reverse the β decay process: $\bar{\nu} + p \rightarrow n + e^+$. A quarter of a century after its theoretical prediction, Reines and Cowan made the first experimental observation of the neutrino, via inverse beta decay at the Savannah River nuclear reactor facility in 1956 [2]—a discovery that earned the Nobel Prize in 1995.

By the time of the discovery of the electron neutrino (ν_e), the muon (μ) had already been extensively studied and determined to be the second-generation charged lepton—a heavier, unstable brother to the electron. Though it would not be measured until the 1960s, the muon neutrino (ν_μ) was postulated as its neutral counterpart. The same would later happen upon the discovery of the third-generation τ and ν_τ .

2.2 The Standard Model Neutrino

The unification of previously separate particle theories into the highly successful “Standard Model” (SM) began in the 1960s. Of greatest importance to the neutrino, the electroweak theory of Sheldon Glashow combined the weak interaction with quantum electrodynamics (QED) [3]. A few years later, Weinberg [4] and Salam [5] merged this with the Higgs mass mechanism [6, 7, 8]. As a result, we have a scheme in which the charged leptons




Figure 2.1: Feynman diagrams of β decay via (a) Fermi's interaction and (b) the standard model weak interaction.

can interact electromagnetically, via the massless photon. Both the charged leptons and neutrinos interact via a similar non-contact (unlike Fermi's version) process, but only at very short-ranges, mediated by the heavy W (charged current) and Z (neutral current) bosons. The masses of all involved particles, including the gauge bosons, are generated from the Higgs, with the exception of the neutrinos, which are taken to be massless. Beta decay via a W boson exchange is illustrated in the Feynman diagram in Fig. 2.1b.

The SM would go on to include the strong force, mediated by gluons, via quantum chromodynamics (QCD). The modern SM is summarized in Fig. 2.2. Besides the fundamental leptons— e , μ , τ , ν_e , ν_μ , ν_τ —the SM incorporates quarks and antiquarks, which don't exist on their own, but rather assemble in doublets and triplets to make mesons and baryons, respectively. In this regard, the SM has been exceptionally successful. Not unlike Dmitri Mendeleev's periodic table accurately predicting the existence and chemical properties of dozens of elements a century earlier, the SM, quark model, and eightfold way anticipate as-of-yet undiscovered particle states. Heavier and heavier mesons and baryons have been found largely where expected. Though SM hadron mass prediction is not very accurate, there have been basically no unexpected discoveries or contradictions in the quark sector. The last missing SM particle, the Higgs boson itself, was discovered at CERN in 2012.

In the lepton sector, the original SM likewise prescribes the same particle species that we observe experimentally. The three generations of neutrinos differ in their flavor, which corresponds to the three charged leptons, but all of them are taken to be massless. In fact, the SM is inflexible on this: no coupling to the single Higgs, which serves as the sole generator of mass, is provided.




Figure 2.2: Summary schematic of the Standard Model of particle physics. Public domain image from Wikimedia Commons [9].

2.3 Massive Neutrinos

Neutrino Oscillations

The idea of a massive—and oscillating—neutrino actually predates the SM and Higgs mechanism theory. Bruno Pontecorvo suggested the possibility of neutrinos oscillating in an analogy with neutral kaons in 1957 [10, 11]. Mass is necessary for oscillations: neutrinos are created in a particular weak (flavor) state, and a particular energy-momentum state, within the Heisenberg uncertainty principle. An energy-momentum eigenstate, of which the mass is an eigenvalue, represents a *superposition* of the three weak ones. As the neutrino flies, its energy and momentum are obviously conserved, but the observed flavor may change with some probability depending on distance. A massless neutrino is invariable in its state unless it enters an interaction such as a capture. The theory of neutrino oscillations in vacuum was formulated by Maki, Nakagawa, and Sakata [12] and Pontecorvo [13] in the 1960s, simultaneously with the development of the SM. Mathematically, we write the superposition of three weak eigenstates ν_a into three mass eigenstates ν_i as follows:

$$|\nu_a\rangle = \sum_{i=1}^3 U_{ai} |\nu_i\rangle. \quad (2.1)$$

U_{ai} is the PMNS¹ 3×3 unitary mixing matrix:

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{\text{CP}}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{23}c_{13} \end{pmatrix} \quad (2.2)$$

$$\times \begin{pmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

where s_{ij} and c_{ij} stand for $\sin(\theta_{ij})$ and $\cos(\theta_{ij})$, respectively. θ_{12} , θ_{23} , and θ_{13} are the the neutrino mixing angles between the three weak eigenstates. The phase δ_{CP} quantifies CP violation in the mixing, as do the Majorana phases α_1 and α_2 (if the neutrino has a Majorana mass). The three neutrino mass eigenvalues, regardless of mass mechanism, are given by m_1 , m_2 , and m_3 . Later, oscillations of neutrinos in matter, which are stronger or weaker than in vacuum depending on said matter's density, were described with the MSW effect [14, 15, 16].

The SM vision of the massless neutrino was not inconsistent with earlier experimental results measuring neutrinos from β decays and antineutrinos from reactors at short baselines with relatively low sensitivity and poor understanding of the reactor spectrum. The first evidence for neutrino oscillations, and hence their being massive, came from Ray Davis' experiment at the Homestake Mine in South Dakota, at the present location of the Sanford Underground Research Facility (SURF). Davis used a massive tank of cleaning solution to measure solar neutrinos via the $^{37}\text{Cl}(\nu_e, e^-)^{37}\text{Ar}$ reaction between 1970 and 1994, and observed only $\sim 1/3$ of the expected ν_e flux [17, 18]. Through the 1990s, other experiments using different techniques saw similar results, though with different fractions that depended on their thresholds and sensitivities. Many in the physics community believed quite confidently that this phenomenon, known as the “solar neutrino problem,” could be explained away by flaws in the measurements. Davis' technique of counting individual ^{37}Ar atoms after bubbling them out of the chlorine solution with helium was particularly suspected.

In the end, the SNO heavy water Cherenkov experiment, thanks to its sensitivity to all neutrino flavors, confirmed the Homestake results for *electron* neutrinos. Furthermore, it observed quantities of ν_μ s and ν_τ s to account for 100% of the expected flux if three flavors were considered. The ratios were consistent with neutrino oscillations and the MSW effect [20, 21]. Unlike SNO, which was a single detector mostly detecting events from a single source (the sun), the KamLAND scintillation detector measured electron antineutrinos from all of Japan's nuclear power plants. Having different sources at varying distances, KamLAND measured disappearance as a function of energy (plotted in Fig. 2.3) [22, 19]. These discoveries cemented the existence of neutrino oscillations, and hence the fact that neutrinos are massive. Ray Davis was vindicated and received the Nobel Prize in 1995, while oscillation experimentalists were awarded in 2015.

¹Named after the theory's four authors: Pontecorvo, Maki, Nakagawa, Sakata.




Figure 2.3: KamLAND’s neutrino oscillation discovery. Figure from [19].

Subsequent experiments measured all of the neutrino mixing angles (see PMNS matrix, Eq. 2.2), as well as mass splittings, though not absolute mass. Furthermore, it appeared that two of the eigenstates were much closer to each other in mass than the third. As of this writing, the mass splitting measurements are not sufficiently precise to determine if that outlier is the heavy or the light state, though the latter seems to be favored. This is known as the “neutrino mass hierarchy problem,” and is illustrated in Fig. 2.4.

The ν SM

The confirmation of the neutrinos as massive particles through oscillation measurements necessitates modifications to the SM. Theoretically, there are three basic approaches to constructing a “ ν SM”:

1. Allow ν to couple to the existing Higgs boson, but very weakly to explain the observed extremely low masses.
2. Introduce a new Higgs boson, as of yet undiscovered, as a new source of electroweak symmetry breaking.

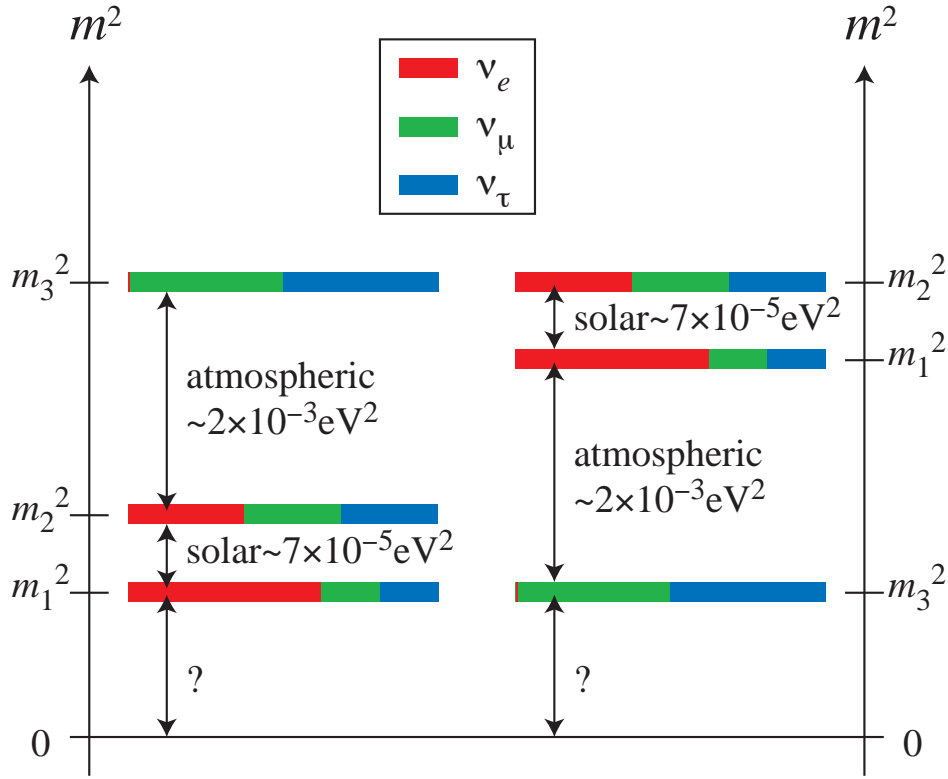


Figure 2.4: The normal (*left*) and inverted (*right*) neutrino mass hierarchies. Figure from [23].

3. Have a completely new mass mechanism, a new source of a different type of mass with a different energy scale that leads to the neutrinos' tiny masses.

In the first case, we would need a new (unnaturally small) Yukawa coupling for U(1) symmetry breaking, so this seemingly conservative approach would still require beyond-the-standard-model (BSM) physics. These neutrinos, like the charged leptons and the quarks, would be Dirac fermions. The Dirac equation gives us four degrees of freedom and necessitates that we have four states: left- and right-handed neutrinos and antineutrinos ν_L , ν_R , $\bar{\nu}_L$, $\bar{\nu}_R$. Problematically, the weak interaction has a chiral structure and we observe only left-handed W bosons, so the ν_R and $\bar{\nu}_L$ chirality states would be “sterile.”

The second and third cases both allow for “Majorana neutrinos.” In a fermion model proposed by Ettore Majorana in 1935, there is only one independent chiral field and hence two degrees of freedom [24]. Thus, the particle and antiparticle are identical up to a phase, differentiated only by their helicity state, which is not a fundamental property and can be flipped. Naturally, such a construction only works for neutral fermions, so it is not an option for the charged leptons and the quarks. Of the fundamental particles known to exist experimentally and in the SM, the neutrino is the only candidate. The right-handed neutrino

would be one and the same as the antineutrino, and vice versa, potentially solving the sterile neutrino question.

In [24], Majorana shows that the behavior of ordinary β decay would be unchanged were the neutrino indeed its own antiparticle, and we have gone on to determine that the remainder of experimentally observed weak processes would likewise be the same. As of this writing, there is no preference for one neutrino mass model over the other based on empirical evidence. In fact, the two types of masses can coexist, and you can have both sterile and Majorana neutrinos (such as the seesaw mechanism, a subset of scenario 3 above). Neutrinoless double-beta ($0\nu\beta\beta$) decay is a process that would only be possible with Majorana neutrinos, and probing this experimentally allows us to distinguish between option 1 and options 2 or 3, as we discuss in the next chapter.

Chapter 3

Double-Beta Decay

3.1 Two-Neutrino Double-Beta ($2\nu\beta\beta$) Decay

Single beta (β) decay, as discussed in Section 2.1 and illustrated by the Feynman diagram in Fig. 2.1, is a nuclear decay process wherein a neutron converts to a proton with the emission of an electron and an antineutrino. This type of decay is very ubiquitous. In 1935, Maria Goeppert-Mayer postulated the existence of *double* beta decay ($2\nu\beta\beta$), with two neutrons turning into protons simultaneously, releasing two $e^-\bar{\nu}_e$ pairs [25]:

$$(Z, A) \rightarrow (Z + 2, A) + 2e^- + 2\bar{\nu}_e . \quad (3.1)$$

The Feynman diagram in Fig. 3.1a depicts this channel in the electroweak/SM formalism, with two W boson-mediated interactions. As one can expect, this diagram is highly suppressed. In unstable nuclei where single β decay is possible, it is effectively nonexistent.




Figure 3.1: Feynman diagrams of (a) the rare $2\nu\beta\beta$ and (b) the theoretically-proposed $0\nu\beta\beta$ decay processes.

An interesting phenomenon occurs in a few of the “even-even” nuclei—those whose proton and neutron numbers are both even. A single- β decay $(Z, A) \rightarrow (Z + 1, A) + e^- + \bar{\nu}_e$ of an even-even nucleus would produce an odd-odd nucleus. Even numbers of protons and neutrons, however, have higher binding energies due to pairing effects. Even-even nuclei, consequently, tend to be more stable than their odd-odd counterparts, and those with a similar but odd mass numbers. As is illustrated by the nuclear mass isobars in Fig. 3.2b, isotopes with odd mass numbers can successfully undergo β decay if they are not stable. Likewise, those with even mass numbers can do so quite easily if both proton and neutron numbers are odd. A subset of unstable even-even isotopes, however, find themselves in a more energetically favorable state than their would-be daughters. Thanks to conservation of energy, this transition is forbidden, and the nucleus must undergo $2\nu\beta\beta$ decay to the next available even-even state, if it is to achieve an even more stable configuration. This scenario is illustrated in Fig. 3.2a.

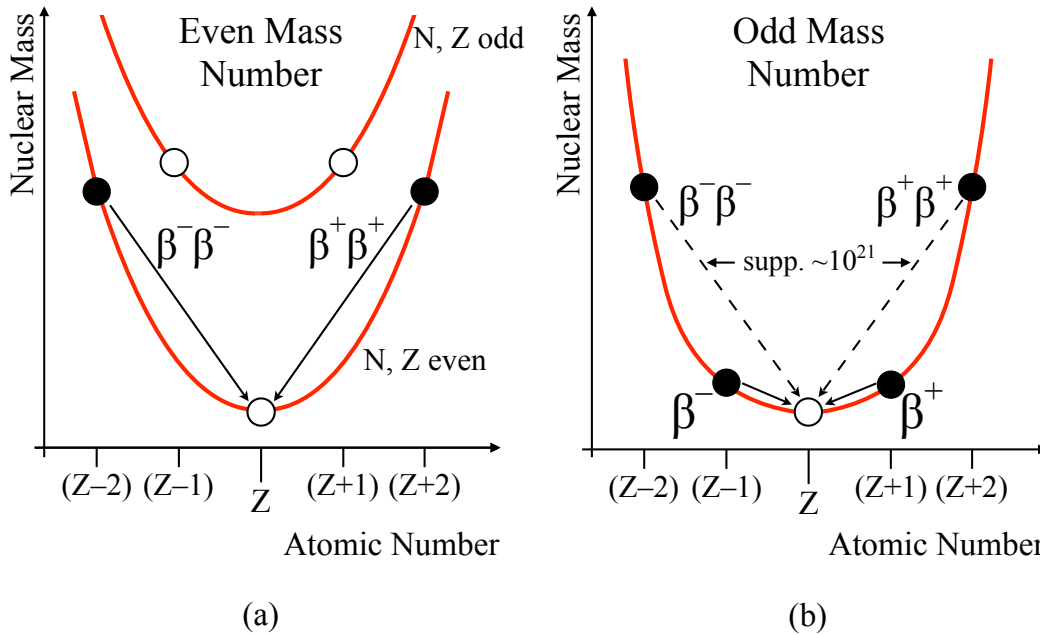


Figure 3.2: Schematic representation of the nuclear mass isobars for isotopes with (a) even and (b) odd mass numbers. Only the $\beta\beta$ decay channel is allowed for some (not all) nuclei with even proton and neutron numbers, as shown here. A detailed chart of the $A = 130$ isobar is given in Appendix B.

In nuclei with odd mass numbers, $2\nu\beta\beta$ decay is suppressed by a factor of $\sim 10^{21}$ with respect to the β decay. In those even-even isotopes where it is the only possible channel, however, the simultaneous conversion of two nucleons remains just as unlikely. Fig. 3.3 gives the particular example of ^{130}Te , the object of study of the CUORE experiment, which undergoes double-beta decay to ^{130}Xe . The measured half-life $T_{1/2}(2\nu\beta\beta)$ of ^{130}Te and three

other isotopes is given in Table 3.1. That list is not exhaustive, but all known $2\nu\beta\beta$ half-lives are all somewhere between 10^{19} and 10^{24} y. In fact, two-neutrino double-beta decay, first observed directly in 1987 in ^{82}Se by means of a time projection chamber [26], is the slowest natural process to have been measured experimentally.




Figure 3.3: Simple energy level diagram: β decay of ^{130}Te to ^{130}I is forbidden, but rare $\beta\beta$ decay to ^{130}Xe occurs. Mass excess values from [27], measured Q-value from [28, 29, 30]. A detailed chart of the $A = 130$ isobar is given in Appendix B.

3.2 Neutrinoless Double-Beta ($0\nu\beta\beta$) Decay

As we have discussed in Chapter 2, neutrinos are not massless as postulated originally in the Standard Model (SM). The mechanism from which these particles get their mass, however is not known—any scenario would require the introduction of new physics. A central question in this area of scientific investigation is whether or not some or all of this mass could come from a Majorana mass mechanism. That is, we must determine whether or not the neutrino is a Majorana fermion—a type of massive neutral fermion wherein the particle and antiparticle differ only in their helicity states [24]. Experimentally, we do this by attempting to either discover or exclude some physical process that would only be possible with a Majorana neutrino.

In 1939, building upon Majorana’s original model and further theoretical work done by Giulio Racah in 1938 [31], Wendell Furry showed that *if* the neutrino is a Majorana particle, it is possible for $\beta\beta$ decay to occur with the emission of only electrons [32]:

$$(Z, A) \rightarrow (Z + 2, A) + 2e^- . \quad (3.2)$$

This process—neutrinoless double-beta ($0\nu\beta\beta$) decay—is shown in the Feynman diagram in Fig. 3.1b. Given that the Majorana fermion is “its own antiparticle,” one can think of the neutrinos in this decay as annihilating virtually with each other. The SM forbids such a reaction because it violates lepton number—a conserved quantity in the model—by 2. If we are dealing with Majorana neutrinos, however, lepton number cannot be a natural symmetry by definition: the ν and $\bar{\nu}$ quantum states are only different up to a phase. A ν SM with Dirac neutrinos coupling weakly to the ordinary Higgs boson would conserve lepton number, and forbid $0\nu\beta\beta$ decay in the same way.

Considering the two $\beta\beta$ decays in Fig. 3.1 only in terms of phase space, one finds that the neutrinoless channel is actually expected to be *faster* than its two-neutrino counterpart. If this were the case, the fact that $2\nu\beta\beta$ decays have been measured on multiple occasions while $0\nu\beta\beta$ remains elusive would seem like evidence against the Majorana nature of the neutrino. Indeed, early experimental searches from the 1940s to the 1950s focused on $0\nu\beta\beta$ over $2\nu\beta\beta$ under precisely this assumption. However, we must also take into account the neutrino mass being exceptionally small. In $0\nu\beta\beta$ decay, the vector-axial vector weak current of the neutrino(s) is(are) in the role of a propagator, which suppresses said diagram. Alternatively, thinking again of this exchange as a virtual annihilation of two Majorana neutrinos, one of our $\bar{\nu}$ s must be transformed into a ν , which necessitates a “helicity flip,” which is more difficult with lighter particles. While it is not kinematically appropriate in this particular case, one can imagine the mathematically analogous question of what Lorentz boost would be necessary to transform the observed helicity state of a particle depending on its mass.

3.3 Scientific Implications

In terms of its scientific significance, the question of $0\nu\beta\beta$ decay and the Majorana nature of the neutrino is fascinating in a number of ways. As we have already made abundantly clear, either the discovery or exclusion of this channel will shed light on the source and nature of the neutrino mass. Either outcome would mean yet more exciting BSM physics from the neutrino sector.

Of course, measuring neutrinoless double beta decay would be more informative than ruling it out. Specifically, if the neutrino *does* have a Majorana mass, then the $0\nu\beta\beta$ decay rate $\Gamma_{0\nu\beta\beta}$ of a candidate nucleus depends on it via the “effective Majorana mass” $m_{\beta\beta}$:

$$\Gamma_{0\nu\beta\beta} = G|\mathcal{M}|^2|m_{\beta\beta}|^2, \quad (3.3)$$

where G is a phase space parameter, and \mathcal{M} is the appropriate nuclear matrix element. The effective mass is related to the true neutrino mass eigenvalues as follows:

$$m_{\beta\beta} = \left| \sum_i m_i \cdot U_{ie}^2 \right|. \quad (3.4)$$

U_{ie} are the ν_e elements of the PMNS matrix given in Eq. 2.2. Fig. 3.4 shows the current state of exclusion limits on $m_{\beta\beta}$, evaluated from limits on $0\nu\beta\beta$ half-lives of four different even-even isotopes. Matrix element values, at this point, are model-dependent.




Figure 3.4: Constraining the effective Majorana mass of the neutrino ($m_{\beta\beta}$) from the current CUORE limit on $T_{1/2}^{0\nu\beta\beta}(^{130}\text{Te})$ [33] and target sensitivity [34], as well as measurements from other isotopes [35, 36, 37]. Nuclear matrix elements from [38, 39, 40, 41, 42, 43, 44, 45, 46]. The red and green bands represent the normal and inverted hierarchy parameter spaces, smeared due to the unknown value of the Majorana phases.

Very deep physical significance may be drawn from the discovery of a lepton number violating process. In reality, we observe an asymmetrical universe with an excess of matter. Without the “baryon asymmetry,” the mutual annihilation of all the particles and antiparticles would prevent the forming of more complex structures, from heavy nuclei to biological life. Soviet physicist Andrei Dmitrievich Sakharov, who concerned himself in more ways than one about problems in nuclear physics with significance to human existence, determined three conditions necessary for baryogenesis [47, 48]:

1. Baryon number violation,
2. C- and CP-symmetry violation, and
3. Interactions out of thermal equilibrium.

The SM does not necessarily predict equal quantities of matter and antimatter in the universe, and it does meet all three of these criteria. However, it does so insufficiently: due

to very limited capacity for CP violation, the maximum allowed asymmetry in the SM is $\eta_B = \frac{n_B - n_{\bar{B}}}{n_\gamma} \sim 10^{-20}$ [49]. The cosmologically observed value is $\eta_B \simeq 6 \cdot 10^{-10}$ [50]. Via the process of leptogenesis, processes with *lepton* number violation can create baryogenesis. A net violation of the baryon number can stem from that of the lepton number—that is, from $0\nu\beta\beta$ and its equivalent reaction—and satisfy the first of the Sakharov conditions. For the second condition, the Majorana phases α (PMNS matrix, Eq. 2.2) let us violate CP.

3.4 Experimental Considerations

Neutrinoless double beta decay is the best available experimental tool for probing the Majorana nature of the neutrino and its mass, as well as lepton number and CP violation. To maximize our detection sensitivity to this process, we must optimize our experiment for its unique experimental signature. The $\beta\beta$ sum energy spectrum of ordinary $2\nu\beta\beta$ decay is analogous to the single β spectrum we discuss in Section 2.1—a bell-like continuum between zero and the full decay energy at $Q_{\beta\beta}$. The two undetected neutrinos (Fig. 3.1a, Eq. 3.1) carry away the “missing energy.” Neutrinoless double beta decay (Fig. 3.1b, Eq. 3.2) outputs two electrons only, so their sum energy must always be the full decay value. Hence, the spectral signature of $0\nu\beta\beta$ is a monoenergetic peak at the Q-value, as illustrated by the cartoon spectrum in Fig. 3.5.




Figure 3.5: Cartoon $\beta\beta$ sum energy spectrum, neglecting background and detector effects, showing the signatures of $2\nu\beta\beta$ decay (bell-like continuum) and $0\nu\beta\beta$ decay (monoenergetic peak at the Q-value). Note that the observed amplitude of this peak depends on $\Gamma_{0\nu\beta\beta}$, and its observed width is limited by detector energy resolution.

The spectrum in Fig. 3.5 is a representation of $\beta\beta$ energy at the *source*, and not at the detector. In real life, we must consider signal statistics and background. This translates into the interconnected choices of the candidate nucleus and the detector technology to employ. Fig. 3.6 summarizes experimentally studied $\beta\beta$ isotopes' Q-values and natural abundances. More detailed information about the four most experimentally significant isotopes are given in Table 3.1. Ideally, we would try to maximize both of these parameters. Having a high Q-value puts you in a region with lower background—below ~ 2.5 MeV there is a dense agglomeration of γ lines from the uranium and thorium series, as well as their combined Compton continuum. When searching for such a rare event, we may well have just one or two excess events in our $0\nu\beta\beta$ peak, and so background should be as low as possible. This parameter disfavors, for example, ^{128}Te . Again because event rates are low, and because experimental budgets tend to be limited, we want to maximize statistics with high isotopic abundance. Even if we pursue enrichment, the initial natural abundance must be high enough for it to be effective. ^{130}Te is a clear winner in this regard, while ^{48}Ca , with its amazingly high Q-value, is basically impractical.




Figure 3.6: Summary of $0\nu\beta\beta$ decay candidate isotopes: $Q_{\beta\beta}$ versus natural abundance. An experiment prefers high values of both parameters.

Detector performance is also very important. Real world devices all have finite resolution, for example, that smears out both the $2\nu\beta\beta$ continuum and the $0\nu\beta\beta$ peak, causing them to overlap. Solid state detectors—ionizing HPGe and cryogenic crystal bolometers—tend to be the best for this. Liquid scintillator experiments are noticeably worse. In Fig. 3.7, we see a pair of background-free $\beta\beta$ spectra from hypothetical liquid noble ^{136}Xe (Fig. 3.7a) and ^{130}Te bolometric (Fig. 3.7b) experiments. This example readily illustrates how energy resolution can make the difference between a prominent signal and one that is barely discernible. Some

Table 3.1: Experimentally significant $0\nu\beta\beta$ decay candidate isotopes: $Q_{\beta\beta}$ values, isotopic abundances, $0\nu\beta\beta$ $T_{1/2}$ limits, and $2\nu\beta\beta$ $T_{1/2}$ measurements (from [51]).

Isotope	$Q_{\beta\beta}$ (keV)	a (%)	$T_{1/2}(0\nu\beta\beta)$ (10^{25} y)	$T_{1/2}(2\nu\beta\beta)$ (10^{21} y)
^{76}Ge	2039	7.61	>8.0 [35]	$1.65^{+0.14}_{-0.12}$
^{100}Mo	3034	9.63	>0.11 [36]	0.0071 ± 0.0004
^{130}Te	2528	34.17	>1.5 [33]	0.69 ± 0.13
^{136}Xe	2458	8.87	>11 [37]	2.19 ± 0.06

detector technologies also allow for active background rejection—scintillators are great for this, and HPGe detectors also have some capacity. Pure bolometers, measuring only energy through thermal signatures, are weak in this regard.




Figure 3.7: Cartoon experimental $\beta\beta$ spectra, with $\Gamma_{0\nu\beta\beta}/\Gamma_{2\nu\beta\beta} = 10^{-6}$ and no background, as measured by (a) a ^{136}Xe experiment with $\sim 2.9\%$ FWHM energy resolution, and (b) by a bolometric ^{130}Te experiment with CUORE’s 7.7-keV FWHM energy resolution.

In the end, $0\nu\beta\beta$ search experimental design, for background-limited cases, comes down to optimizing the equation

$$\text{Sensitivity} \propto a \sqrt{\frac{Mt}{b\Delta E}}, \quad (3.5)$$

where a is isotopic abundance, M is total detector mass, t is experiment live time, b is background rate, and ΔE is energy resolution. Detector technology and isotope choice are inextricably linked, and thus no experiment is ideal. Xenon, for example, is a noble gas, so a solid state ^{136}Xe experiment would be highly challenging. It works well in scintillators

like KamLAND-Zen and nEXO, which are also easier to make large, but energy resolution is poor. HPGe detectors combine excellent energy resolution with low backgrounds, but the $\sim 10\%$ isotopic abundance is just adequate, and the low Q-value is also a downside. Furthermore, HPGe devices are expensive and difficult to scale up. A ^{130}Te bolometer like CUORE, the subject of this thesis, can have good isotope mass and energy resolution, but, in addition to lacking active background rejection, requires complicated sub-Kelvin cryogenic equipment that may be challenging to operate for long live times. In the end, we do not know in what portion of the available parameter space we will encounter neutrinoless double beta decay (if at all). As such, we do not know if we will be setting an exclusion limit, or making an actual half-life measurement. It can be argued that a low-resolution detector with high mass, long live time, and low background, is better for limit setting—low resolution is not a big problem if you don't see a signal. A high-resolution machine like CUORE could be better suited for declaring evidence or discovery.

Chapter 4

The CUORE Experiment

4.1 Introduction

CUORE—the Cryogenic Underground Observatory for Rare Events—is a currently operational experiment searching for the neutrinoless double-beta ($0\nu\beta\beta$) decay of ^{130}Te . Located in the Apennine Mountains of central Italy, at the Laboratori Nazionali del Gran Sasso (LNGS) underground laboratory of the Italian Istituto Nazionale di Fisica Nucleare (INFN), the CUORE detector is a one-of-a-kind ton-scale bolometer array. Comprising 988 $5 \times 5 \times 5 \text{ cm}^3$ TeO_2 crystals operated at $\sim 10 \text{ mK}$ in the world’s largest and most powerful dilution refrigerator, CUORE contains $\sim 206 \text{ kg}$ of the isotope of interest and lays a solid claim to being the coldest contiguous cubic meter in the known universe [52]. As we discussed in the previous chapter, a $0\nu\beta\beta$ experiment must balance several characteristics while making certain sacrifices. CUORE’s excellent energy resolution of $\sim 7.7 \text{ keV FWHM}$ at $Q_{\beta\beta} \simeq 2527.5 \text{ keV}$, $\sim 34\%$ isotopic abundance, high total mass, and the reasonably high ^{130}Te Q-value (above the Compton edge) make this technically challenging probe highly competitive.

4.2 Bolometric Detection

A bolometer¹ is a type of detector that functions by registering thermal signatures from particles. A simple schematic of a bolometer is shown in Fig. 4.1. The centerpiece of the device is the absorber. Its size, shape, and material is chosen depending on the detector’s intended application, and the types of radiation it must be most sensitive to. Bolometers built for underground searches for rare event physics typically use dielectric crystals of various types. A particle stops in the absorber, sitting at a base temperature T_b , and the kinetic

¹“Bolometer” is the commonly accepted term, though technically “calorimeter” is more appropriate: we are measuring energy through heat, while a true bolometer measures *power*. This is the case because “calorimeter,” in particle physics, is strongly associated with the electromagnetic and hadronic calorimeters of accelerator-based experiments. One sometimes encounters the designation “microcalorimeter,” though it is questionable to what extent a device millimeters or centimeters in size can be referred to as “micro.”

energy dissipates within its bulk. The lattice structure of most bolometers' single-crystal absorbers facilitates nearly instantaneous and uniform heating up to some temperature T related to event energy.




Figure 4.1: Schematic drawing of a bolometric detector.

Assuming we collect the entirety of the detected particle's energy, the magnitude of the thermal fluctuation depends inversely on the heat capacity of the absorber:

$$\Delta T(E) = T - T_b = \frac{E}{C(T)} . \quad (4.1)$$

For ΔT to be detectable, $C(T)$ must be very small. We achieve this by cooling the bolometer down to ultra-low temperatures: existing devices are operated at temperatures on the order of 10 – 100 mK. In this regime, the crystalline absorber approximately follows the Debye model, and its heat capacity has a cubic temperature dependence given in Eq. 4.2 [53].

$$C(T) \sim k_B \left(\frac{T}{\Theta_D} \right)^3 , \quad (4.2)$$

where k_B is Debye's constant and Θ_D is the Debye temperature—a property of the absorber material.

The crystal is coupled to a heat sink. This attachment must be designed and fabricated in such a way as to facilitate the cooling of the absorber back down to its original temperature T_b , but not too rapidly. Otherwise, if the heat dissipates immediately without building up,

the temperature change ΔT will not be enough for detection with sufficient sensitivity. Our ideal bolometric signal is in the form of a pulse with a nearly instantaneous rise and a fall-off in the form of an exponential decay with the time constant

$$\tau = RC \ , \quad (4.3)$$

where R is the thermal resistance of the weak coupling and C is the heat capacity of the absorber.

Herein lies the bolometric detection approach's main advantage: its excellent energy resolution. Theoretically, this quantity is limited only statistically—by random thermal fluctuations from the crystal across the ideal weak coupling:

$$\langle \Delta E^2 \rangle \sim k_B T^2 C(T) \ . \quad (4.4)$$

$C(T)$ is given in Eq. 4.2, and *resolution* goes as the square root of the above quantity, so we get

$$\sqrt{\langle \Delta E^2 \rangle} \sim k_B T \left(\frac{T}{\Theta_D} \right)^{3/2} \ . \quad (4.5)$$

Of course, couplings (and all the other components) are non-ideal, have heat capacities of their own, and this performance is not achievable in practice. Nonetheless, energy resolution is highly competitive, which is very useful when attempting to discover rare processes with distinct energy signatures.

The measurement of the temperature fluctuation is done by means of a thermistor—a “thermal resistor,” the electrical resistance of which depends on temperature. Thermistors come in various forms, usually chips fabricated with semiconductors or superconducting metals, which have the necessary strong and reproducible temperature dependence in the bolometer's sub-Kelvin domain of operation. Electronics and the readout technique depend on the type of device used. The ideal bolometer's thermistor has negligible heat capacity and a perfect thermal coupling to the absorber crystal, such that their temperatures are always equal. Real-world devices attempt to approach this goal as much as possible, but naturally there are practical limitations.

4.3 CUORE: a State of the Art Ton-Scale Bolometer Array

TeO₂ Crystals

The CUORE detector comprises an array of 988 channels, each of which is essentially a separate bolometer. We use radiopure TeO₂ [54, 55] crystals as our absorber. This selection is dictated by the practical considerations of $0\nu\beta\beta$ decay searches. Namely, because double-beta decays occur so rarely, we need a very significant quantity of the isotope of interest to

achieve meaningful sensitivity. Furthermore, since electrons are relatively short-ranged, the isotope must be distributed very evenly throughout the detection medium. Basically the only practical solution here is to design an experiment where the detector and source are one and the same. With CUORE, as discussed in Chapter 3, we search for $0\nu\beta\beta$ in ^{130}Te , which has a high natural abundance of $\sim 34\%$ and a decently energetic Q-value of (2527.515 ± 0.013) keV [28, 29, 30]. Telluride is mined industrially, and mass-produced single-crystal TeO_2 has already found extensive use in science thanks to its birefringent properties. Thus, it is a natural and practical choice for our bolometric absorber.

The CUORE crystals are $5 \times 5 \times 5 \text{ cm}^3$ cubes, weighing $\sim 750 \text{ g}$ each. With a Debye temperature of $\sim 232 \text{ K}$ [56] and operated at $\sim 10 \text{ mK}$, these have an inverse heat capacity of

$$C^{-1} \simeq 0.1 \text{ mK/MeV} . \quad (4.6)$$




The shape and size were selected semi-empirically, while balancing various characteristics. As in our “ideal bolometer” in Section 4.2, our sensitivity improves if absorber heat capacity is lower (Eq. 4.1), and apart from cooling the crystals to lower sub-Kelvin temperatures the best way to do so is to decrease their physical size. On the other hand, seeking to detect *both* of the electrons from $2\nu\beta\beta$ and $0\nu\beta\beta$ events in a single crystal, we must choose absorbers that are several centimeters in each dimension. Besides the strict maximum imposed by crystal-growing technology, the compromise between bolometer performance and detection efficiency essentially determines the CUORE crystals’ size. The greater the detector array’s segmentation, the greater the surface area to volume ratio. This translates directly into the issue of backgrounds from surface-versus-bulk contamination. The cube shape is most conducive to close packing, and therefore surface background rejection through anti-coincidence. The CUORE-type TeO_2 crystals are produced by Shanghai Institute of Ceramics, Chinese Academy of Sciences (SICCAS) in China, and were competitively selected due to their excellent radiopurity [54, 55, 57]. Internal contamination is quite small, especially since the crystal growing process tends to push impurities outwards. In the end, surface events from external sources dominate our background (see Fig. 4.6), and we benefit from cube-shaped absorbers over flatter rectangular prisms. The Cuoricino detector featured both $5 \times 5 \times 5 \text{ cm}^3$ channels, which would become standard, and $6 \times 3 \times 3 \text{ cm}^3$ ones. The former were confirmed to be superior [58].

NTD Readout

Semiconductors have a long history of use as thermal sensors at low temperatures. In these materials, charge carriers in the valence band are not free to move between atoms and conduct electricity—they must be excited to the conduction band, which is separated by an energy bandgap. Temperature directly affects how many carriers are available for conduction, and at low temperatures this number is much smaller. In the cryogenic regime, even small changes in excitation become relatively significant, and semiconductors have very high electrical resistances that depend strongly on temperature.

Eventually, if they are cold enough, intrinsic semiconductors become impossible to use as sensors. Germanium is better than silicon in this sense, because with its smaller band gap the failure temperature is also lower [53]. Hence, it is preferred for cryogenic thermometers. Even these, however need to be “doped” with desirable impurities. Group IV semiconductors, including germanium, are typically augmented with electron donors from Group III to create n-type semiconductors with an excess of free electrons, or acceptors from Group V to make p-type semiconductors where the charge carriers are electron holes. In most applications dopants are introduced via processes such as ion implantation, diffusion, or vapor-phase epitaxy. These, however, tend to concentrate impurities near the surface of a thin wafer and not penetrate into the bulk. Industrial uniform doping is done at the level of growing the semiconductor boule (crystal), but this tends to be limited to relatively low concentrations.

To operate germanium sensors in the sub-Kelvin domain of a bolometric detector like CUORE, the standard techniques are not workable. Our doping must be uniform, but also at very high levels. The answer lies in neutron transmutation doping, or NTD. In the context of precision bolometric measurements, this technique was pioneered by E. E. Haller at LBNL [59]. In the NTD process, Ge wafers are irradiated with neutron flux from a nuclear reactor. Instead of introducing dopant atoms, we create them when Ge nuclei absorb the neutrons:



The first of these processes produces ^{71}Ga , which is an electron acceptor, but ^{75}As is a donor and ^{77}Se , from Group VI, is a double donor. Since the reactor core²—our neutron source—is much larger than Ge wafers, the radiation flux is very uniform. The distribution of Ge isotopes within the wafers is also highly homogeneous. Finally, the mean free path of thermal neutrons in Ge is significantly greater than the wafer thickness³, and thus we get a very even volume distribution of all three dopants inside the semiconductor. It is found that with a nominal neutron exposure of $4 \cdot 10^{18} \text{ cm}^{-2}$, we get a net dopant concentration of $|N_A - N_D| \leq 2 \cdot 10^{10} \text{ cm}^{-3}$ [59, 60].

To instrument the bolometers of CUORE and its predecessor experiments, we create thermistors from wafers of NTD Ge. This is done by dicing $3 \times 3 \times 1 \text{ mm}^3$ chips and depositing gold pads on two sides of each one for gold wire bonding. All CUORE NTD development and fabrication was performed at LBNL Material Science Division by J. W. Beeman. An example of a CUORE NTD is shown in Fig. 4.2a. Prior to detector construction,

²CUORE made use of the research reactor at MIT.

³Wafers are 1 mm thick, while mean free path $l \simeq 20 \text{ cm}$ from $l = 1/(\sigma \cdot n)$, where $n \simeq 4.4 \cdot 10^{22} \text{ cm}^{-3}$ is Ge number density, and cross section σ is about 1 barn.




Figure 4.2: **(a)** Photograph of an NTD Ge thermistor (credit CUORE collaboration). **(b)** Example of NTD resistance v. temperature measurement and fit from UC Berkeley.

we characterized NTDs with different irradiation levels via four-wire measurements with a resistance bridge in the UC Berkeley R&D cryostat and at collaborating institutions. The temperature dependence of their resistance is exponential:

$$R(T) = R_0 e^{\sqrt{T_0/T}} , \quad (4.10)$$

where R_0 and T_0 are the key characteristic parameters of the thermistor. The former represents the chip's minimum electrical resistance—the value to which the R-T curve asymptotes at higher temperatures. The latter is the $1/e$ roll-off temperature, defining the exponent slope. Fitting resistance v. temperature data, we find typical numbers to be $R_0 \sim 1 \, \Omega$ and $T_0 \sim 4 \, \text{K}$. Note that T_0 is *not* the same quantity as the bolometer base temperature T_b in Section 4.2. An example of one NTD characterization is given in Fig. 4.2b.

The NTDs are coupled to the TeO_2 crystal absorbers by means of Araldite Rapid [61] epoxy robotically deposited in a pattern of nine spots. This epoxy has been found to contain $< 2.2 \cdot 10^{-10} \, \text{g/g} \, {}^{232}\text{Th}$ chain and $< 8.2 \cdot 10^{-10} \, {}^{238}\text{U}$ chain contaminants [57], and cures fully in about an hour at room temperature. We use $25 \, \mu\text{m}$ diameter gold wires anchored by gold ball bonds to make a two-wire electrical connection between the NTDs' pads and the Cu-PEN⁴ strip cables running down two sides of each CUORE tower. These strips, in turn, are connected to 36 AWG NbTi⁵ ribbon cables running to the top of the CUORE cryostat. The connections are passed out through vacuum bulkhead Fischer connectors to the CUORE

⁴Copper on polyethylene naphthalate substrate

⁵Niobium-Titanium is a superconducting alloy with a high critical temperature of $\sim 10 \, \text{K}$, popular in cryogenic wiring.




Figure 4.3: Photographs of fresh glue spots deposited on an NTD (left) and Si heater (right), from the CUORE assembly line. Image from [57].

electronics. NTDs at ~ 10 mK can have a resistance of $O(100 \text{ M}\Omega)$ or even more, and our electronics bias the NTDs through $O(10 \text{ G}\Omega)$ load resistors to perform a conventional voltage measurement. The signal is passed through an analog Bessel filter [62], amplified [63, 64, 65], and then digitized continuously at a rate of 1 kHz [66, 67, 68].




Figure 4.4: An example of a CUORE pulse. The red marker indicates the trigger. Typical rise time is ~ 100 ms (10% to 90% of amplitude), fall time— ~ 400 ms (90% to 30%).

In Section 4.2, we discuss an idealized version of a bolometric detector. In that model, we work under the assumption that the thermistor is ideally coupled to the crystal—that is, that their temperatures are always equal. In practice, the connection is via nine epoxy glue

spots, each of which has a thermal conductance of about 2×10^{-9} W/K [69]. Furthermore, the two gold wires connecting the NTD to the Cu-PEN cable are an alternative path for heat to escape the bolometer. These two factors contribute to the features we observe in physical CUORE pulses, such as the one in Fig. 4.4.

Structure

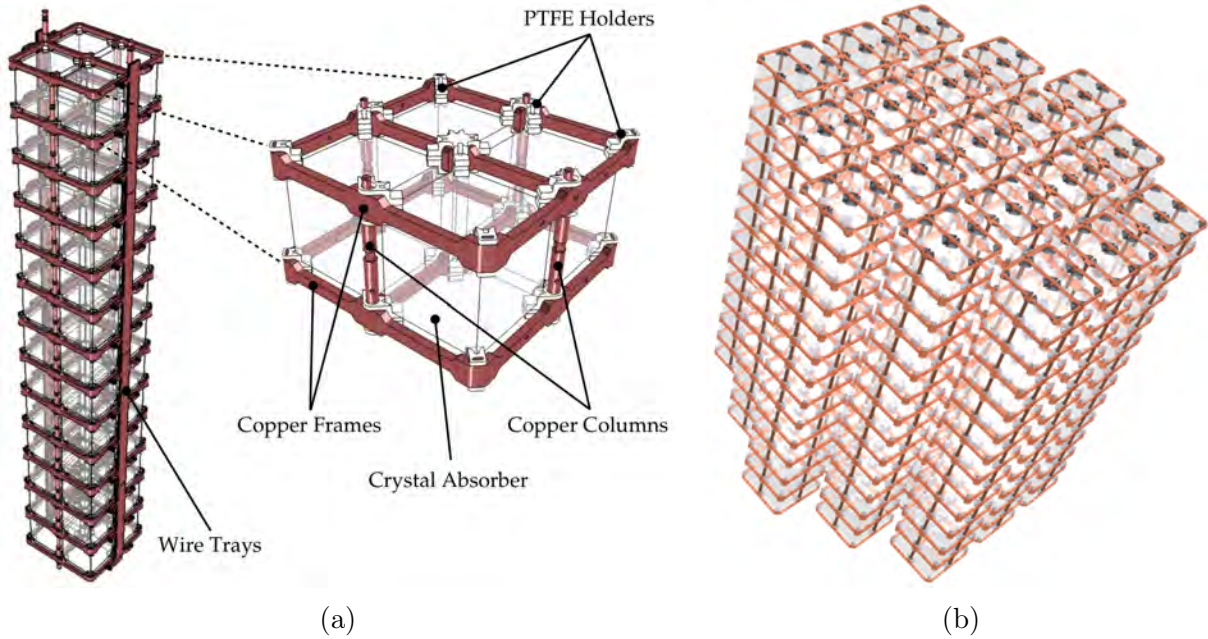


Figure 4.5: Renderings of the CUORE bolometers: **(a)** one 52-crystal tower, with zoom to one floor of 4 crystals held in PTFE supports in a copper frame [57]; **(b)** the full 19-tower, 988-channel array (credit CUORE collaboration).

Each of CUORE’s 988 bolometric channels is based on a TeO_2 crystal instrumented with an NTD Ge thermistor, as described in detail immediately above. In addition to this, there are the surrounding structural components that simultaneously perform the role of the weak link and heat sink from Section 4.2 and Fig. 4.1.

The basic CUORE bolometer assembly, shown in Fig. 4.5a, consists of a group of four crystals arranged in a square and mounted with PTFE (polytetrafluoroethylene) holders in a copper frame. Copper has very low thermal diffusivity⁶ thanks to its low heat capacity and high thermal conductivity, even at low temperatures because it is a non-superconducting metal. This makes it an ideal material for most cryogenic structural components and specifically thermalization hardware (see Chapter 5). In this case, the copper frames, which have a direct thermal connection to the mixing chamber stage of the CUORE dilution refrigerator, serve as the bolometer’s heat sink/cold bath. For one four-crystal group, the copper

⁶ Thermal conductivity divided by specific heat capacity.

structure consists of two square frames further subdivided into four square sections—one on top, and one on bottom—connected to each other via a set of four threaded vertical posts.

There are eight PTFE supports per crystal—one on each corner. The PTFE pieces have an angled shape that allows them to make contact with all three surfaces of a corner of a crystal. They are anchored to the copper using the same threaded posts that link the two square frames. During assembly at room temperature, the crystals are held somewhat loosely, but are intended to be gripped tighter when cooled thanks to differential thermal contraction of the components. The PTFE serves as the weak thermal link between the bolometric absorber crystal and the copper heat sink. This synthetic polymer material, more commonly known by its DuPont brand name Teflon, was selected for this role due to its low thermal conductivity (see measurement in Appendix A and [70]), good radiopurity, affordability, and ease of machining. That said, it is also the biggest source of non-ideal behavior in the CUORE bolometers. The ideal bolometer discussed in Section 4.2 defines the weak coupling as having a thermal resistance R , and no heat capacity of its own. This means the decay of pulses would be a simple exponential. In reality, it is the earlier-mentioned quantity of thermal *diffusivity* that characterizes the link, and the PTFE holders’ heat capacity is certainly non-negligible [71]. When an incident particle’s kinetic energy is converted into heat in a CUORE bolometer, it does not simply dissipate via the coupling into the copper frame. Rather, some of it is briefly stored in the PTFE, from where it then dissipates via all available paths, including back through the TeO_2 crystal, the NTD, and its gold wire electrical connections. This phenomenon is responsible for the long second time constant that can be seen with the naked eye in CUORE pulses (Fig. 4.4), which has been confirmed by detector thermal modeling [69].

In addition to the NTD, each bolometer hosts a naked silicon heater. This is the second, smaller chip mounted adjacent to the NTD in Fig. 4.8b. Thanks to its smaller size compared to the thermistor, the heater is adhered to the TeO_2 with just five dots of epoxy. Besides this, its attachment and electrical connection are identical to those of the NTD. We use the heaters to generate “fake” pulses for detector baseline stabilization, diagnostics, and various tests [72].

The four-bolometer assemblies (Fig. 4.5a) are stacked vertically into towers. The top frame of one “floor” serves as the base for the next one. Each CUORE tower has 13 levels, for a total of 52 crystals. Along the side of each tower run Cu-PEN strip cables—essentially long, thin printed circuit boards with copper traces on a polyethylene naphthalate substrate—used to make the electrical connections of the NTDs and heaters. The detector comprises an array of 19 such towers, seen in Fig. 4.5b. They are attached to the vibrationally-isolated “tower support plate,” or TSP, at the 10-mK stage of our large custom dilution refrigerator discussed in Chapter 5.

4.4 The CUORE Assembly Line

Overview

CUORE was preceded by the CUORE-0 detector and experiment. This demonstrator was a single CUORE-type tower⁷, like we have just described in Section 4.3, operated in the cryostat of the Cuoricino prototype [57, 73]. Compared to the older Cuoricino, CUORE and CUORE-0 leverage a major improvement at the construction stage in the form of a new cleanroom facility and detector assembly line. The special class-1000 environment occupies the entirety of the second story of the CUORE experiment “hut” underground at LNGS, and the entirety of the detector building process takes place within its confines. Furthermore, we perform all work inside gloveboxes over-pressurized with boil-off nitrogen gas. All completed and partially completed components are likewise stored in nitrogen-flushed containers. The implementation of this system, discussed in more detail below, has allowed us to improve detector hardware regularity and build quality over Cuoricino, which manifests directly in detector performance [33, 57, 73, 74, 75, 76, 77, 78]. No less significantly, the glovebox-based assembly line and cleanroom—together with improved materials selection and cleaning, purer crystals, and size reduction of copper parts—helped to facilitate an approximately tenfold reduction in CUORE-0’s α background with respect to Cuoricino, as we can see in the comparison of spectra in Fig. 7.1a [76, 77, 78, 79]. CUORE’s background budget, detailed in Fig. 4.6, aims for a further passive background reduction of a factor ~ 6 over CUORE-0 near the ^{130}Te $0\nu\beta\beta$ decay Q-value of ~ 2528 keV [74, 80]. We have been able to meet this goal with the help of CUORE’s custom cryostat (Chapter 5), built from radiopure materials and incorporating extensive cold lead shielding, and a special radon-free installation environment as discussed in Chapter 7 and [81].

NTD and Heater Attachment

As stated previously, the NTD Ge thermistors and Si heaters are adhered to the surface of CUORE’s TeO_2 crystal absorbers with fast-curing epoxy. A glue joint is fundamentally inferior in terms of bolometer performance to the perfect goal of a resistance-free thermal coupling. As such, it is important for us to maximize the effectiveness of this coupling. Additionally, with almost two thousand chips to glue, we must maximize channel-to-channel uniformity and assembly efficiency. As such, the CUORE gluing station (Fig. 4.7), housed entirely in one large nitrogen-flushed glove box, utilizes two computer-controlled robots to partially automate the process.

Front and center in the glove box are a pair of “Z-positioners.” These devices, composed of precision-machined stainless steel and titanium components, are used to properly locate the heater and thermistor chips prior to gluing. Each one features two groups of three metal pegs—one set for each chip—against which the thermistor and heater are placed upside down

⁷While structurally equivalent, the 19 CUORE towers are all new. The former CUORE-0 tower is currently in nitrogen storage.




Figure 4.6: Summary of the CUORE background budget. Figure from [74].

to define their X and Y coordinates. We can see an example of this in Fig. 4.3, which shows a heater and an NTD with glue already applied. More important is the vertical Z axis that gives these mechanisms their name. For this purpose, the chips rest on a pair of vertical pistons, which are coupled to micrometer calibrated screws. The operator manually adjusts these to position the chips' surface up against a pair of swing-out reference plates [57, 82, 83].

Next, we apply the adhesive to the chips. Before we can do this, the adhesive must be mixed—like all epoxies, Araldite [61] functions through a chemical reaction between a liquid resin and a hardener. A pneumatic piston squeezes equal amounts of the two components through a mixer tube into an open-topped syringe, tipped with a short hypodermic needle, which the operator then covers with a cap that incorporates another pneumatic line. The syringe is held vertically by a three-axis Cartesian robot—a computer-controlled apparatus that precisely positions the syringe in using three perpendicular rails. The Cartesian robot is pre-programmed with the X, Y, and Z coordinates of the heater and NTD on both of the Z-positioners. Once we are ready, the software automatically guides the syringe to these locations one at a time to deposit a pattern of nine ~ 0.6 mm glue spots on the NTD and five on the heater. This is achieved through a combination of pressurized air bursts from the pneumatic system and a vertical jerking motion of the Cartesian robot. As one can imagine, it is crucial to release the droplets at the proper height with respect to the chip surface, which is the main reason for using the Z-positioners. We found the spot pattern to be necessary, because a flat even layer of glue leads to crystal fracturing during the cooling of the detector due to excessive differential contraction. Finally, we inspect the glue spots. A camera is

mounted on a different portion of the Cartesian robot, and moves into position above the chips to make a high resolution photograph of the glue spots on both chips. Fig. 4.3 is one example of such a pair of photographs. Aided by an image recognition program placing circular target markers on top of the drops to gauge their size, the user makes sure that the quantity of epoxy deposited is neither too small nor too large, and that its distribution is sufficiently even.




Figure 4.7: Photograph of the interior of the CUORE gluing glovebox, used for attaching NTDs and Si heaters to TeO_2 crystals (credit CUORE Collaboration).

Once the glue spots are ready and the Cartesian robot withdraws, we can place the crystal on top of the chips. Crystals await their turn on a pair of PTFE shelves directly behind the Z-positioners (visible in Fig. 4.7). Again using pre-programmed coordinates, a robot arm manufactured by the ABB company [84] picks up a crystal and sets it down on the Z-positioner. This is that device's second function. The crystal's full weight is held up by a set of supports that located at precisely the right distance from the surface of the NTD and heater. This ensures that the epoxy droplets are sufficiently compressed to make a solid coupling, but are not squeezed into a single uniform layer [57, 82, 83].

Though Araldite Rapid is practically hard for most intents and purposes within just 5 or 10 minutes of application [61], it takes ~ 1 hour for it to fully cure. During this time, we let the crystal rest on the Z-positioner before placing it back on the PTFE shelf. To speed up the assembly, we take advantage of having two Z-positioner stations to stagger the gluing work. Once a full batch of bolometer crystals has been instrumented, the ABB robot arm moves them into an area in the back of the glovebox for placing into hard plastic vacuum

containers for transport and/or storage. On the way, the arm passes each crystal in front of a near-infrared camera that images the dried glue spots *through* the crystal for another quality check. At the time of the “boxing,” the operator visually inspects each of the future detectors for dirt, inclusions, cracks, or any other flaws [57, 82, 83].

Wire Bonding and Tower Assembly

The assembly of the CUORE towers, together with making the wire bonds for the electrical connections of the chips, is done at a separate workstation shown in Fig. 4.8. This is based on a work table integrated with a sealable nitrogen-flushed “garage” able to house one fully-assembled tower in a vertical orientation. This table can couple to different specialized gloveboxes. First, the tower is assembled from crystals (with NTDs and heaters already attached), copper frames, and Teflon supports. As shown in Fig. 4.8a, we do this one layer at a time. Once one “story” is completed, the tower is lowered several centimeters deeper into the garage for the mounting of the next set of four crystals. This is repeated 13 times for one full tower of 52 bolometer channels. Upon completion, the tower is fully retracted into the garage and sealed. At this point, we replace the compact assembly glovebox, which has convenient openings in a diagonal plane, with a tall one that allows us to attach the Cu-PEN strip cables to the sides of the tower [57, 83].




Figure 4.8: Photographs of **(a)** CUORE tower construction and **(b)** NTD and Si heater wire bonding (credit CUORE Collaboration).

Finally, the tower is stowed in the garage one more time and a Westbond 7700E bonding machine is installed inside the tall glove box (see Fig. 4.8b). This manual bonder has been specially modified to be operated in a vertical plane instead of the usual horizontal one, installed on a set of rails that afford it extra movement and reach. The operator is also provided with a digital video camera with high magnification to use while making connections. The bonding procedure for a single NTD or heater chip is as follows:

1. Attach one end of the gold wire (25 μm diameter) to the bonding pad on the chip with a ball bond.
2. Pull away from the bond to extract a length of wire of a few centimeters, and anchor the wire to the Cu-PEN cable’s pad with a wedge bond. The wire should not be stretched tight, but curve in a “U” shape.
3. Place an extra ball bond on the Cu-PEN end to reinforce the connection.
4. Repeat steps 1 – 3 four times: we utilize two-wire connections for both the NTD and the heater, but each one is made with two wires for redundancy and robustness [57, 83].

The bonding machine and its tip are both mobile to ensure that we can precisely place each wire in the correct location. Additionally, a second operator raises, lowers, and rotates the tower in the garage as necessary until 104 chips have been bonded. The final step in tower assembly is to bolt protective copper covers over the Cu-PEN strips, and place the tower in nitrogen-flushed storage so that the assembly table and garage may be freed up for the construction of the next one. We store the towers in cylindrical Plexiglas containers mounted on wheeled carts. During the detector installation, these are rolled directly into the CR6 radon-free cleanroom described in Chapter 7, which surrounds the cryostat [81].

4.5 Data Taking, Production, and Low-Level Analysis

Dataset Organization and Structure

CUORE data taking is subdivided into datasets and runs. The former, for practical reasons, are relatively short periods of continuous acquisition lasting a few hours to day or two. Besides general-purpose “test” runs used for optimization, diagnostics, and optimization, runs are divided into two categories: “calibration” and “background,” or “physics.” Our experimental setup features a novel cryogenic detector calibration system (DCS), which can lower 12 necklace-like source strings amidst the bolometer array without significantly raising the temperature. These consist of capsules with ^{232}Th sources inside, on Kevlar threads [33, 85], which we use to calibrate the energy scale. Calibration runs are taken while the DCS is deployed. Physics runs, used to hunt for rare processes, are obviously conducted without it. Each CUORE dataset comprises about one month of background measurements, with several days of calibration runs before and after [33].

Pre-Processing

Of CUORE’s 988 bolometers, 984 (99.6 %) are active (have working electrical connections). The raw data from each one is a continuous voltage waveform sampled at 1 kS/s. The first step in processing this data stream is to set and eliminate “bad intervals”—periods where

the detector is noisy; unstable due to interference from the earthquakes that are so common in the Gran Sasso mountains, human activity, or heavy equipment operations nearby; or compromised in some other fashion. Next, we convert the continuous data to individual thermal events. We do this using a derivative trigger, which trips on the steep rise of the pulses (Fig. 4.4). Response time performance varies channel-to-channel, and so this trigger corresponds to a different energy threshold in each one, between 20 and a few hundred keV [33, 86, 87]. We save a 10-s window for each event—three seconds before and seven seconds after the trigger, asymmetrically because the pulse decay time is significantly longer than the rise. The pulse amplitude is determined with respect to the “pulse baseline” taken from those three seconds of pre-trigger data. In the physics data, we observe an average event rate of ~ 6 mHz. During calibration runs with the DCS, this increases to ~ 50 mHz [33]. In addition to events from absorbed particles, we have the “fake” pulses generated with the Si heaters mentioned above in Section 4.3. These are generated using a constant voltage every few minutes. Finally, we employ random triggers to record baseline noise.

Optimal Filter

Our bolometric pulses, like the one illustrated in Fig. 4.4, exhibit a strong dependence of amplitude on event energy thanks to the low heat capacity of the TeO_2 crystals and the NTDs’ exponential R-T curve (Eqs. 4.6 and 4.10). We further improve our signal to noise ratio by means of “Optimal Filtering” (OF) [88]. This technique, which CUORE inherits from Cuoricino and CUORE-0, functions in the frequency domain. Working in this space, we can describe the pulse as

$$y(\omega) = as(\omega) + n(\omega) , \quad (4.11)$$

where a is the pulse amplitude, $s(\omega)$ is a function describing the detector response, and $n(\omega)$ is the usual additive noise. Practically, we can do this by taking a from the event itself, determining $s(\omega)$ by calculating an average pulse from the entire dataset or run, and evaluating $n(\omega)$ as the average noise power spectrum from our aforementioned random triggers. Having this information, we can write a filter that weighs each frequency bin to maximize the signal therein with respect to noise:

$$y_{\text{filtered}}(\omega) = \sum_i w_i(\omega) y_i(\omega) , \quad (4.12)$$

where i refers to the bin, and the weights $w_j(\omega)$ are defined as

$$w_i(\omega) \sim \frac{|s_i(\omega)|^2}{\langle |n_i(\omega)|^2 \rangle} . \quad (4.13)$$

We define the pulse amplitude as the maximum of this OF-filtered waveform in the time domain. An example is illustrated in Fig. 4.9.

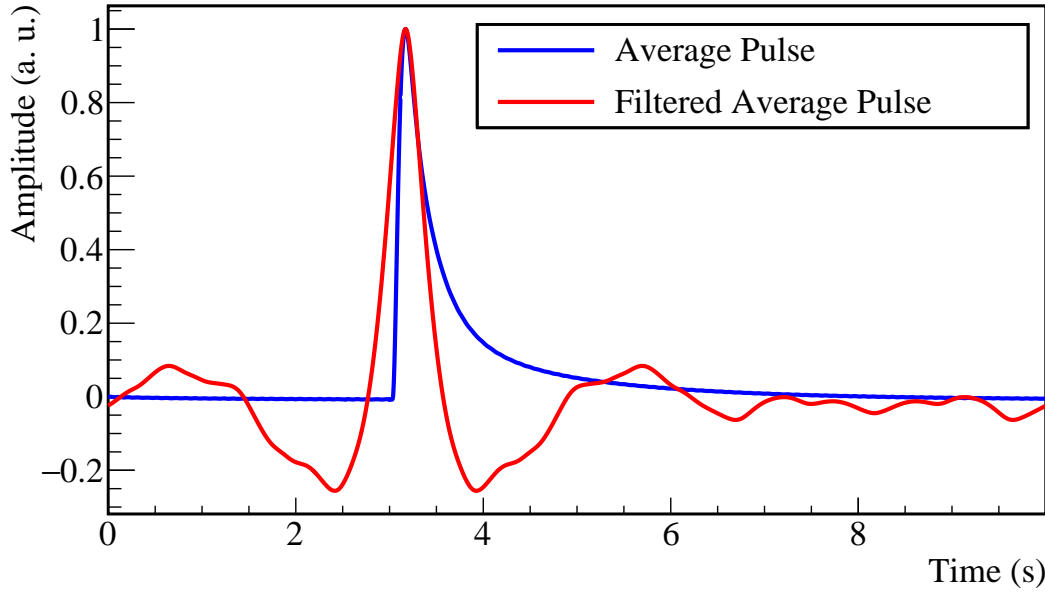


Figure 4.9: An example of a CUORE average pulse, before and after application of the Optimal Filter.

Stabilization

We make an effort to keep the detector baseline as steady as possible during data taking: the dilution refrigerator’s base temperature is stabilized with a PID feedback loop, and personnel access to the experiment’s underground hut is restricted to minimize mechanical vibrations. Nonetheless, the pulse baseline tends to shift event-to-event. Since detector sensitivity is strongly dependent on temperature, this means monoenergetic events reconstruct with varying amplitudes. This can lead to worsened resolution, “split” spectral peaks, and other artifacts. To address this issue, we utilize thermal gain stabilization (TGS). This is done by fitting a trend to the baseline shifting over time, using a set of events known to be monoenergetic and distributed throughout the dataset. By default (96.6% of channels), we do this with the pulser heater events, which are tagged and are always driven by the same voltage. In the remaining cases, we use γ events from the ^{208}Tl ~ 2615 -keV calibration line. While the unstabilized amplitude spectrum is indeed smeared, this peak from the ^{232}Th series is very prominent and can be easily identified [79]. There are two cases in which we choose to employ γ TGS: 1) the small number of channels where it is found to outperform pulser TGS; 2) $\sim 3\%$ of bolometers where the NTD electrical connections are intact, but the heater wires have been damaged [33].




Figure 4.10: The decay chain of ^{232}Th —the source used in the CUORE DCS. We use six γ lines from ^{228}Ac , ^{212}Pb , and ^{208}Tl (bold) to calibrate the CUORE spectrum (shown in Fig. 4.12).

Calibration

The calibration process consists of determining the relationship between the OF-filtered stabilized pulse amplitude and particle energy, and applying it to all detected events. We do this by making use of distinct, monoenergetic γ lines in the spectrum of the calibration sources from the DCS. The Thorium series, starting with ^{232}Th and ending with the stable ^{208}Pb , involves six α and four β decay processes of 11 different isotopes (see Fig. 4.10). A number of these decays release γ rays. Of these, we select six lines for calibrating our detector: the 239 keV from ^{212}Pb ; 338 keV, 911 keV, and 969 keV from ^{228}Ac ; and 2615 keV and 583 keV from ^{208}Tl [33]. These are reliably prominent and easy to identify in the amplitude spectrum both with the naked eye and with a computer peak finding algorithm. Additionally, they decently cover the entire energy range from 0 keV to above the ^{130}Te double-beta decay Q-value at (2527.515 ± 0.013) . For each channel-dataset pair, we locate these six peaks in the stabilized amplitude spectrum, determine their means and variances, and equate them to the known true energy values. Subsequently, these six points are fitted with a quadratic polynomial. The fit result gives the reconstructed energy as a function of stabilized OF-filtered amplitude, which we evaluate for every recorded event.

Selection

Besides discarding the “bad intervals,” which comprise $\sim 1\%$ of the total live time, we apply a number of other selection cuts to our physics data. Basic pulse quality criteria eliminate pile-up events—those with multiple maxima, inflections and kinks, and insufficiently stable

Table 4.1: Summary of CUORE efficiency for the first two datasets, taken from [33]. “Total cuts” is a combination of the previous three lines, and does not include $\beta\beta$ containment.

Selection efficiency (%)	Dataset 1	Dataset 2
Base	95.63 ± 0.01	96.69 ± 0.01
Pulse shape (Mahalanobis)	91.1 ± 3.6	98.2 ± 3.0
Anti-coincidence (random)	99.4 ± 0.5	100.0 ± 0.4
Total cuts	85.7 ± 3.4	94.0 ± 2.9
$\beta\beta$ containment	88.35 ± 0.09	88.35 ± 0.09

pre-trigger baselines. Using a simple amplitude-to-energy relationship, there is no way to disentangle the two or more particles that go into a pile-up event, thus making such a signature essentially non-physical. This cut, together with trigger and selection efficiencies, gives the “base” efficiency quoted in Table 4.1 [33].

A second-order pulse quality screening is done by comparing every pulse against a proper “signal-like” response, determined from events that can be reliably assumed to come from single γ -rays. Specifically, we take signals lying within 10 keV of the 1173-keV and 1332-keV ^{60}Co , and 1461-keV ^{40}K peaks. These lines are the most prominent in the background spectrum (Fig. 4.12), and lie near the center of the spectral domain below ~ 2800 keV—the portion of greatest interest and importance for all our calculations. The comparison to this template is done by means of a single parameter—the Mahalanobis distance [89]—calculated in six-dimensional space. The cut value is chosen such that the $^{130}\text{Te } 0\nu\beta\beta$ decay discovery sensitivity is maximized [33, 90].

Thus far, we’ve explicitly considered only the scenario where the full energy of a particle is deposited in a single CUORE bolometer. Indeed, the vast majority ($88.345 \pm 0.085\%$ as of this writing) of double-beta events are expected to be of this type: CUORE’s $5 \times 5 \times 5 \text{ cm}^3$ TeO_2 crystals are more than large enough to contain the electrons, which typically penetrate a few millimeters in a solid. As such, for the purposes of our main analysis, we cut multi-channel events, defined as those occurring within a 10-ms window. These come from various sources. High-energy cosmic muons are easily identified, depositing energy in crystals all along their tracks. The Compton length of γ rays in our crystals is ~ 5 cm, leading to a significant portion of multiple-scatter events. As we can see in Fig. 4.6, the biggest contributors to our α background are the surfaces of copper parts, and electrical hardware on the detector towers—Si heaters and Cu-PEN cables. Next, however, come surface contaminants on the TeO_2 crystals themselves. A significant portion of these events trigger in two channels: the nuclear recoil registers in the source-bolometer, while the α particle is absorbed in an adjacent crystal. Finally, there are random coincidences, which we also cut because there is no way to be sure that the events are unrelated. With a ~ 6 mHz/channel signal rate in physics data, this practice does not lead to a large loss in efficiency. In the relatively high-rate calibration data, random coincidences are much more common, so we refrain from applying this cut to boost statistics.

Blinding

To ensure our $0\nu\beta\beta$ decay analysis is unbiased and fair, it needs to be blind. One simple way to do this is to literally remove or obscure a portion of the data around $Q_{\beta\beta}$. This solution, however, is not very good: the gap in the data creates difficulties in testing and tuning fits to the region of interest (ROI). Instead, we perform “salting” of our data. That is, we take some real events from the strong 2615-keV γ line, and shift them down by 87 keV to build an artificial peak at the Q-value. The total number of shifted events is unknown: it is some random fraction between 20% and 40% of the total counts in a 40-keV-wide window around the ^{208}Tl line. As such, it is not possible for the analyst examining the physics spectrum in the ROI to determine if any portion of the salted peak is in fact real $0\nu\beta\beta$ decay candidate events. An additional benefit of this approach is that the fake peak has the distribution of a real CUORE spectral line. This permits us to ensure that our ROI fitting code could properly handle a discovery scenario. We unblind the data by reversing the salting process once the analysis is finalized [33, 79].


4.6 The First Data from CUORE

Table 4.2: Summary of CUORE detector performance for the first two datasets [33]. Combined background rate is exposure-weighted.

	Dataset 1	Dataset 2	Combined
Number of channels	876	935	1811 ch-ds pairs
TeO ₂ exposure (kg · y)	37.6	48.7	86.3
¹³⁰ Te exposure (kg · y)	10.5	13.5	24
Energy resolution at $Q_{\beta\beta}$ (FWHM, keV)	8.3 ± 0.4	7.4 ± 0.7	7.7 ± 0.5
Background rate in ROI (10^{-2} c/keV/kg/y)	$1.49^{+0.18}_{-0.17}$	$1.35^{+0.20}_{-0.18}$	1.4 ± 0.2

The CUORE cryostat, with the detector inside, was cooled to its ~ 10 mK base temperature in January of 2017. We spent the first few months of operation exploring the behavior of the bolometer array, lowering noise, and tuning its performance. A particularly important consideration is choosing the operating temperature. We do this by means of a temperature scan between 7 mK and 25 mK. Ideally, as we explain in Section 4.2, a bolometer should be run as cold as possible to minimize the absorber’s heat capacity and maximize sensitivity through thermal gain. In practice, we need to balance this with optimal signal bandwidth. With CUORE, we found that the signal-to-noise ratio is inferior at the lowest temperatures, in part due to the challenge of reading out the NTDs at their highest resistances—in the gigaohm range. As of this writing, CUORE’s chosen base temperature is ~ 15 mK [33], though we are interested in lowering this in the future through noise reduction.

Our first physics measurements were performed over the spring and summer of that year. In Fig. 4.11, we plot the TeO₂ exposure in kg · y versus the date. The two upward-

Figure 4.11: CUORE TeO₂ exposure versus date.

sloping portions correspond to two datasets, with the interruption in between being a ~ 1 -month period of detector optimization. Detector performance is summarized in Table 4.2. Improvement from dataset 1 to dataset 2 is readily apparent in terms of energy resolution: (8.3 ± 0.4) keV versus (7.4 ± 0.7) keV [33]. This parameter is of cardinal importance to maximizing $0\nu\beta\beta$ decay discovery sensitivity, since the small peak at the Q-value must be separated from the $2\nu\beta\beta$ continuum. Nonetheless, we are still short of the ~ 5 keV achieved by the CUORE-0 demonstrator experiment. We aim to approach this target in the future. In addition to the analysis cuts described in Section 4.5 and Table 4.1, we manually eliminate entirely the worst-performing channels. Thanks to the tuning, we manage to keep 59 more active bolometers in dataset 2 compared to dataset 1. This, in part, is the reason for the greater exposure of dataset 2. In total, we analyze data from 1811 channel-dataset pairs, which corresponds to 92% of live bolometers (1986 pairs) [33].

The results of datasets 1 and 2 are analyzed in Chapters 9 – 11 of this thesis. The full physics spectrum is plotted in Fig. 4.12, overlaid with the calibration spectrum normalized such that the amplitude of the ^{208}Tl 2615-keV line is equal to that in the background. Fig. 4.13 shows the blinded and un-blinded background spectrum in the vicinity of $Q_{\beta\beta}$ of ^{130}Te . The salting method described in Section 4.5 is illustrated well here. We give a detailed treatment of fitting the shaded ROI by various methods in Chapter 10.




Figure 4.12: CUORE background (physics) spectrum from datasets 1 and 2, overlaid with the calibration spectrum scaled to the amplitude of the 2615-keV line. Labeled are γ lines from: 1) ^{212}Pb , 2) ^{228}Ac , 3) e^+e^- annihilation, 4) ^{208}Tl , 5) ^{54}Mn , 6) ^{60}Co , 7) ^{40}K , and 8) ^{214}Bi .

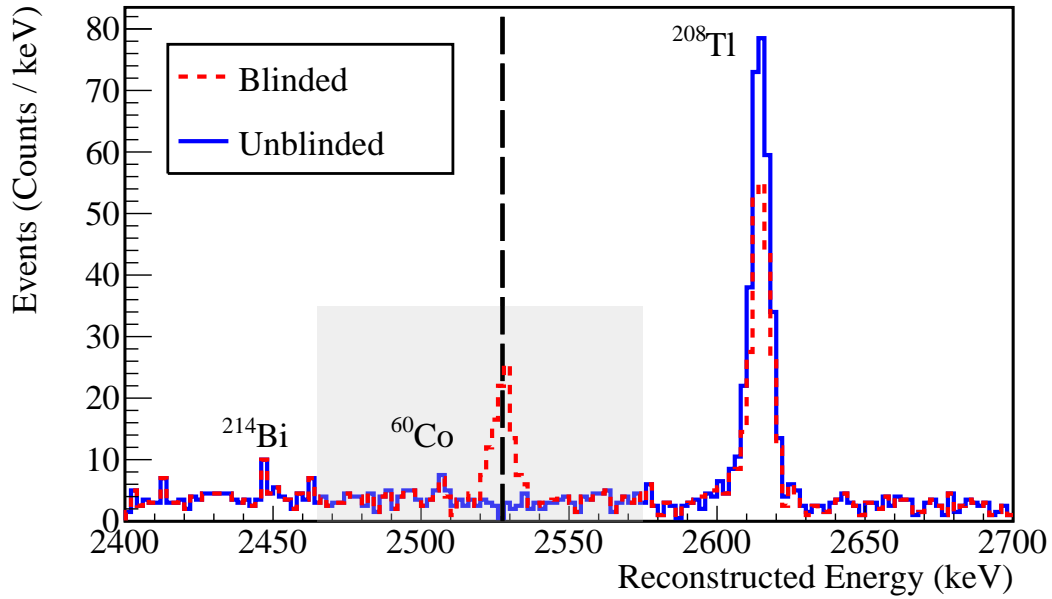


Figure 4.13: CUORE blinded and unblinded background (physics) spectra near the ^{130}Te $0\nu\beta\beta$ decay ROI. The black dashed line indicates $Q_{\beta\beta}$.

Chapter 5

The CUORE Cryostat

5.1 Introduction

The CUORE bolometers must be operated at the ultra-low temperature of ~ 10 mK. This serves several purposes. Most importantly, by cooling the detector we lower the heat capacity of the TeO_2 crystal absorbers, in accordance with the Debye model, to values sufficiently small for the effective detection of temperature changes resulting from the absorption of single particles. Furthermore, these temperatures are in the design operating regime of the neutron transmutation doped germanium thermistors (NTDs), where they have a strong dependence of resistance on temperature. We discuss the operation details of bolometers in significantly more detail in Chapter 4.

Such cold environments are typically created and maintained for long periods of time using dilution refrigerators. These devices can achieve sub-Kelvin temperatures by exploiting a phase transition in a liquid mixture of ^3He and ^4He , coupled to the “standard” cryogenic liquid-gas helium phase transition [91, 92]. Given CUORE’s status as the first ton-scale bolometer, taking up ~ 1 m³ of space, the cryostat for this experiment needs to feature a large cold volume and a high-power Dilution Unit (DU).

To realize this goal, we make use of a cryostat designed and built specially for CUORE. With five Cryomech pulse tube cryocoolers producing a total of 7.5 W of cooling power at 4 K [93], and a DU from Leiden cryogenics providing ~ 2 mW at ~ 100 mK and ~ 3 μW at ~ 10 mK [94, 95], this machine allows us to create what we believe to be the coldest contiguous cubic meter in the known universe [52]. Specifically, we cool the ~ 1 t detector to temperatures below 10 mK, along with 7.5 t of lead shielding to ≤ 4 K [94]. The ability to include extensive cold shielding, as close as possible to the detector, is an additional advantage of constructing a custom cryostat for CUORE. Along with careful selection of radiopure materials such as NOSV [96] copper, this leads to a background reduction of $\sim 6\times$ with respect to the CUORE-0 detector operated in a commercial Oxford Instruments cryostat [74, 80]. Finally, the CUORE cryostat features a Detector Calibration System (DCS) that allows the deployment of ^{232}Th calibration sources inside the bolometer array

during data acquisition, without warming up [85, 94].

5.2 Dilution Refrigerator Theory and Operating Principles

$^3\text{He}/^4\text{He}$ Mixture Low-Temperature Phase Separation

The two stable isotopes of helium— ^3He ($N = 1$) and ^4He ($N = 2$)—manifest significantly different behaviors at low temperatures. Due to the 0.98 amu difference in mass, the boiling point (3.19 K) and critical point (3.35 K) of ^3He are both lower than those of ^4He (4.2 and 5.2 K). Furthermore, since the ^4He nucleus is a boson, it becomes a superfluid at 2.7 K—three orders of magnitude warmer than the fermionic ^3He (2.5 mK)¹ [97, 98].




Figure 5.1: A phase diagram of a liquid $^3\text{He} - ^4\text{He}$ mixture, showing the separation between the ^3He -rich (concentrated) and -poor (dilute) phases below ~ 870 mK. Figure from [91, 92]

¹The ^3He isotope turns superfluid *at all* because at sufficiently low temperatures helium atoms form molecule-like diatomic structures, which are bosonic. This phenomenon breaks the dilution refrigerator phase separation behavior, and sets the fundamental 2.5 mK limit on the temperature these devices can reach.

Thanks to the different quantum physics of the two isotopes, the mixture of ^3He and ^4He used in a dilution refrigerator (DR) undergoes a phase separation below a temperature of ~ 870 mK. As illustrated in the phase diagram in Fig. 5.1, the concentration of ^3He in the “ ^3He -poor” phase (Fermi liquid ^3He in superfluid ^4He) drops quickly to 6.6%. The “ ^3He -rich” normal fluid phase, meanwhile, approaches nearly 100% ^3He concentration. These specific percentages are explained by optimization of enthalpy—decreased entropy at low temperature favors nearly complete splitting of the isotopes, but having some small quantity of fermionic nuclei in the bosonic condensed ^4He allows for closer packing. The forced, endothermic dilution of ^3He across this phase boundary is the foundation of the dilution refrigerator working principle, described below.

Precooling With Pulse Tube Refrigerators

To reach the temperature regime in which this phase separation becomes accessible, we must precool and liquefy the helium mixture. Traditionally, dilution refrigerators (DRs) have used a large bath of liquid helium, which needs to be regularly refilled, to cool the circulating mixture to ~ 4 K. This is followed by a smaller liquid helium bath that is actively pumped, colloquially known as the “1-K pot”, to cool to $\lesssim 1$ K, at which point the dilution unit (DU) takes over. The commercial Oxford Instruments machine used for the smaller Cuoricino and CUORE-0 prototype/demonstrator detectors predates the widespread use of mechanical cryocoolers, and is a “wet fridge.” The large liquid cryogen bath can potentially offer some advantages in terms of simplicity and vibrational noise reduction.




Figure 5.2: A schematic diagram of a single stage pulse tube refrigerator. E2, at the center, is the cold stage—real pulse tubes are curved such that this stage is at the bottom. The pulse tubes used for precooling dilution refrigerators are two-stage units. Figure from [91].

We ruled out the use of a “wet fridge” for CUORE given the experiment’s ton-scale cold mass and its long intended live time of ~ 5 years, which would result in tremendous expenditure of liquid helium and reduced live-to-upkeep time ratio. Instead, we selected a “dry,” or “cryogen-free,” design with a pulse tube precool scheme, which has become very common in the last two decades.

A pulse tube refrigerator (PTR or just PT) is a type of closed-cycle mechanical cryocooler, invented by Mikulin in 1984 [99], and illustrated schematically in Fig. 5.2. A compressor and reciprocating motor, located outside of the cold tube itself (CP in Fig. 5.2), periodically pump helium gas into and out of the tube via a regenerator (RG in Fig. 5.2) with a heat exchanger on either end. On one side of the regenerator, at the tube’s hot end, warm gas enters from an external buffer when pressure is low. Meanwhile, on the cold side, gas is being sucked out. Gas passes through the heat exchangers and regenerator while expanding adiabatically, dropping in temperature and cooling its surroundings. Heat extracted from the gas deposits in the regenerator. A fraction of a second later, in the high pressure phase, fresh gas fills the cold end, while heat that had been deposited in the regenerator is extracted at the warm end as gas is forced out of the tube [91, 100]. This is not unlike the process of cooling below 4.2 K by pumping on the liquid helium in the earlier-mentioned “1-K pot,” but done continuously, cyclically, and in the gas phase.

The above-described process is fairly similar across all types of cryocoolers. The PTR is unique with respect to other designs (Gifford-McMahon, Stirling, etc.) in that the cold tube itself—the only part of the device that must necessarily be anchored to the cryostat—is a passive conduit, and does not contain any moving parts. This noticeably improves reliability during prolonged operation, while the ability to house the piston, motor, and compressors externally reduces mechanical vibrations and electromagnetic interference with experimental equipment. Thus, the PTR is the cryocooler of choice for dilution DR precooling [91].

A single-stage PTR is currently unable to cool below ~ 10 K, but in DRs we use a two-stage device. Here, the first stage pre-cools the second, while both are driven by the same compressor and motorhead. The Cryomech PT415 cryocoolers we use for the CUORE cryostat, which are typical modern devices, each provide ~ 40 W of cooling power at ~ 45 K with the first stage (PT1), and ~ 1.5 W at ~ 4 K with the second (PT2) [93]. In fact, with this much surplus cooling power at the “liquid helium” level, temperatures of < 3 K, very close to helium’s λ point—the normal fluid to super fluid helium transition temperature—can be achieved. As a result, DRs cooled by PTRs can forego not only the liquid helium bath, but also the “1-K pot,” as illustrated in the schematic in Fig. 5.3, making do with a Joule-Thompson stage incorporated into the DU itself.

As shown in Fig. 5.3, a DR works in stages. We find a vertical stack of plates at different temperatures between the above-described pre-cool hardware at the top and the Mixing Chamber (where the phase-separated $^3\text{He}/^4\text{He}$ mixture dilution occurs) at the bottom. This is necessary, because cooling power drops with temperature, and each level pre-cools the next. We discuss this in detail in the next subsection, but one important feature of this design is that these plates need to stably operate with very large thermal gradients between them. Typical values are: 300 K (roof), 40 – 45 K (PT1 or liquid nitrogen), 4 K (PT2 or liquid helium), 600 – 700 mK (Still), 50 – 100 mK (“cold”), ~ 10 mK (“base”). Hence, structural connections between the plates are purposefully fabricated as thermal isolators, making use of materials such as G10 and stainless steel. This would prevent our pulse tube(s) pre-cooling the *entire* length of the DU, which is necessary for the helium mixture to condense inside.

Most cryogen-free dilution DR designs address this issue in one of two ways. The first

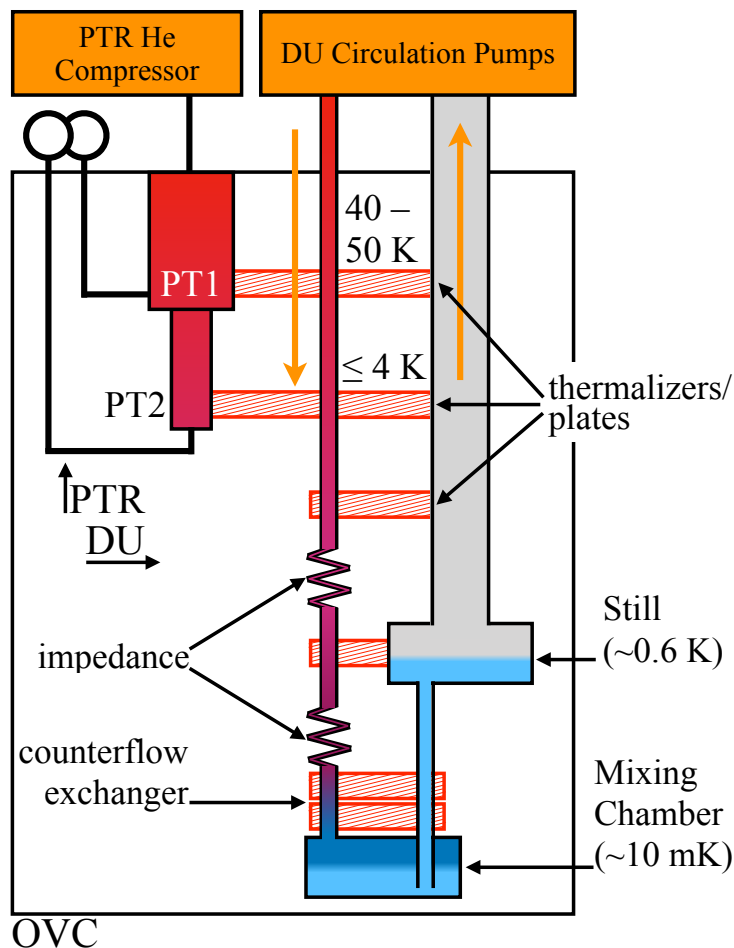


Figure 5.3: A schematic of a dilution refrigerator with “dry” pulse tube precool. Figure 5.4 shows a more detailed view of the Dilution Unit.

method involves the circulation of helium gas (often the mixture used for the DU) through a separate set of tubes. The precool circuit is thermalized on the two pulse tube stages, and on all the plates to be cooled below. It is separate from the DU, and is optimized for a high flow rate with a relatively small volume of gas. Because the precool circuit is a closed cycle, this system gives us the advantage of needing only a single vacuum vessel—the outer vacuum chamber (OVC) at 300 K. The entirety of the cryostat’s volume is maintained under vacuum for the entire cooldown cycle, and only non-hermetic thermal radiation shields are necessary at the intermediate stages. The Oxford Instruments Triton 400 R&D cryostat operated by the Kolomensky group at UC Berkeley (featured in Appendix A) exemplifies this approach.

The second design encloses the cryostat volume below the 4 K (PT2) stage in an additional inner vacuum chamber (IVC). During the precool phase, the PTR(s) operate in the vacuum of the OVC, while the IVC is pressurized with helium exchange gas to conduct

heat from the lower stages to the PTR cold head(s). Once precool is complete, we evacuate the IVC before beginning the condensing. This system is frequently found in refrigerators designed by Giorgio Frossati of Leiden Cryogenics, as well as the CUORE cryostat. With sufficient exchange gas, it provides very high effective cooling power, thus shortening precool time and making it an attractive option for a large machine with a massive payload.

Condensation and Dilution

Once the DU has been precooled, we can begin $^3\text{He}/^4\text{He}$ mixture circulation and condensing. This will cool the three temperature stages below the pulse tube—the 500 – 700-mK Still plate, ~ 50 – 100-mK cold plate, and ~ 10 -mK Mixing Chamber plate. In a sufficiently large cryostat, the circulation is typically driven by two or three pumps in series—a turbomolecular pump (highly effective for low pressures/molecular flow), backed by a Roots blower (high speed in a narrow pressure range), in turn backed by a membrane pump/compressor ensuring the effectiveness of the Roots. The mixture first passes through a 77 K cold trap, comprising a cassette of highly porous material submerged in liquid nitrogen, to freeze or condense out any impurities such as air or water (helium is one of very few substances to condense at temperatures colder than nitrogen). If they are not scrubbed out, the liquefied or solidified impurities can obstruct the thin tubing of the DU.

When we first turn the DU on, the entirety of the mixture—both ^3He and ^4He are actively circulating. In a stable operating regime, the ^4He is condensed, superfluid, and at rest, with only ^3He flowing. The gas enters the DU-proper at a pressure of several hundred millibar, and is cooled by the 40-K and 4-K stages of the PTR(s), to which the mixture lines are thermalized. Though at this point temperatures are at the “liquid helium” level, the lower-than-atmospheric pressure means the mixture remains gaseous. Condensation occurs shortly afterwards, at around 1.2 K—in a “wet fridge,” through the thermalization of the lines to the actively-pumped “1-K pot,” and in a “dry fridge” via a Joule-Thompson valve. This is a very simple device that cools gas by passing it through a nozzle from a higher to a lower pressure. The valve’s cooling power, of course, is much smaller than that of pumped boiling helium, but we are able to use it in a PTR-precooled setup because the nominally 4-K stage actually sits at temperatures closer to ~ 2.5 K [91]. At this point, our helium is ~ 1.2 K. Flowing ^3He condenses, continues through the first impedance² to the level of the Still (see Fig. 5.4).

On the way down, the mixture does not actually pass *through* the Still, but rather is thermalized to it so as to cool the helium to ~ 600 mK (thermalization connections are shown as dashed red blocks in the diagram in Fig. 5.3). Notice that this temperature is already below the ^3He -rich/ ^3He -poor phase separation discussed earlier in the subsection

² the impedance is a narrow capillary through which the mixture flows on its way down through the DU. The first component, often called the “main impedance,” is located right above the Still level, while the “secondary impedance” is immediately below it. The diameter of the impedance, given the mixture volume and the power of the pumps, determines the flow rate of the DU. Careful tuning of this value is important for maximizing cooling power.

“ $^3\text{He}/^4\text{He}$ Mixture Low-Temperature Phase Separation”. This is achieved by the functioning of the Still, which our mixture reaches on its way back *up*, out of the DU. Like the eponymous piece of liquor production equipment, the Still derives its name from “distillation,” and is a chamber where a liquid-gas phase transition is observed (see Fig. 5.4). Inside the Still, the volume is actively pumped to a relatively low pressure of ~ 0.1 mbar. At this pressure and temperature, ^4He is completely superfluid, and remains condensed and motionless at the bottom of the Still, while the ^3He is flowing through the ^4He , and evaporating thanks to its significantly higher vapor pressure. The resulting cooling effect is equivalent to that of the pumped boiling ^4He in the 1-K pot, but achieves colder temperatures due to ^3He ’s lower boiling point and heat of vaporization. The efficiency and power of the Still, and consequently of the whole DU, is predicated upon maintaining sufficient flow of ^3He through the Still. This often occurs at temperatures slightly higher than the minimum achievable at this stage, so the Still typically comes equipped with a resistive heater that allows us to warm it up as needed. Usually we drive this heater at some constant non-zero power, though it can theoretically be connected to a feedback controller (see Chapter 6), whereas some DUs are found to be optimally run with zero additional power on the Still.

The ^3He pumped out of the Still is circulated by the pumps back through the cold trap and into the DU for another cycle. Meanwhile, down-flowing ^3He cooled by the Still continues on to the Mixing Chamber at the bottom. With the superfluid ^4He condensed elsewhere, this ^3He is in the concentrated phase. As it goes down, we are cooling it further by passing it through a counter-flow heat exchanger (Fig. 5.3), where already-cold ^3He is being pumped up an to the Still. Approximately half-way down this heat exchanger, we have the “cold plate,” which despite the name, is the second-coldest and sits stably at $\sim 50 - 100$ mK. Unlike the other cryostat plates above and below, this one doesn’t correspond to a particular active cooling element like the Mixing Chamber, Still, or cold head. However, it provides a useful barrier protecting the ~ 10 mK Mixing Chamber plate below from heat radiation impinging on it from the ~ 600 mK Still plate, as well as an important thermalization point. As we have already mentioned, cooling power decreases significantly with temperature. By definition, the refrigerator has no surplus cooling power at the Mixing Chamber level if it has reached its lowest achievable temperature—otherwise, it would keep cooling down. If a small amount of heat is applied there, raising the temperature by a couple of millikelvin, a few microwatts of cooling power are available at most. At 100 mK, on the other hand, the DU provides us with $O(100) - O(1000)$ μW of cooling power. For experimental payloads, surplus power is necessary only while actively cooling down, and once at base, temperature must simply be maintained. The 100 mK plate, however, becomes essential for thermalizing hardware that passes between the plates and carries heat from warmer stages above, such as electrical wiring for thermometers and experimental sensors.

Lastly, the down-flowing mixture reaches the Mixing Chamber—the final, coldest stage of the DR, where the experimental payload is attached (Fig. 5.4). This is where we find dilution process that gives this device its name. Inside the Mixing Chamber, the concentrated (^3He -rich) and dilute (^3He -poor) phases are completely separate, and in an equilibrium state. Given that temperature of ^3He entering the Mixing Chamber is $\lesssim 30$ mK, the concentrated




Figure 5.4: A schematic diagram of the Dilution Unit, showing the ^3He concentrated-dilute phase transition in the Mixing Chamber, the liquid/gas boundary in the Still, and the flow of ^3He .

phase is nearly 100% ^3He , while the dilute phase is about 6.6% ^3He and 93.4% ^4He (see Fig. 5.1) [101]. Spatially, the dilute phase is on the bottom of the Chamber, while the concentrated phase is on top, with a horizontal boundary. The outlet of the Mixing Chamber, leading up through the counterflow heat exchanger to the Still, originates *below* the phase boundary, within the dilute phase liquid helium. Pumping up through this tube, augmented by osmotic pressure from the evaporation in the Still above, forces the ^3He to flow down from the concentrated phase and into the dilute. This is the opposite of what the helium “wants” to do in these low-temperature, low-entropy conditions, and is a highly endothermic process, removing heat from the surrounding metal. This, in the end, is the useful cooling power of the dilution refrigerator, and is approximately equal to

$$\dot{Q}[\text{W}] = (\dot{n}_{^3\text{He}} [\text{mol/s}]) (95(T_{\text{MC}} [\text{K}])^2 - 11(T_{^3\text{He}} [\text{K}])^2) , \quad (5.1)$$

wherein $\dot{n}_{^3\text{He}}$ is the flow rate of the ^3He , $T_{^3\text{He}}$ is its temperature as it enters the Mixing Chamber, and T_{MC} is the temperature of the Mixing Chamber itself [92]. Thanks to the

endothermicity of the dilution, T_{MC} is actually noticeably colder than the flowing helium itself:

$$T_{\text{MC}} = \frac{T_{3\text{He}}}{2.8} \quad (5.2)$$

in equilibrium [92]. Efficient design of the earlier-mentioned counterflow heat exchanger (often a tightly-wound spiral within a larger spiral with silver sinter couplings to maximize the contact surface area) is essential for minimizing $T_{3\text{He}}$. A significant reason why a majority of DRs do not reach base temperatures *significantly* below ~ 10 mK is that the Kapitza resistance—thermal boundary resistance between liquid helium and the metal—below $T_{3\text{He}} \sim 30$ mK is very significant, and requires an impractically large heat exchanger surface area.

Thermal Radiation Shielding

A final crucial component of any dilution refrigerator’s construction is thermal shielding. Black body radiation can provide a significant heat load, and it is important to absorb and dissipate as much of this load as possible, at the warmest stage possible. Specifically, from the Stefan-Boltzmann law, the radiative thermal transfer between an object of surface area A and its surroundings is

$$\dot{Q} = A\sigma(T_1^4 - T_2^4) , \quad (5.3)$$

where σ is the Stefan-Boltzmann constant and T_1 and T_2 refer to the surroundings’ and the object’s temperatures. Thus, we see that the quartic temperature dependence is very significant, but it is the *relative* temperature value that is key.

If we imagine a 10 mK payload with a ~ 1 m² surface area in a DR, exposed directly to the 300-K OVC, we find it would be bombarded by nearly 500 W of infrared radiation! Such an arrangement would not be possible to achieve, let alone maintain, with even the most powerful DU available. If, on the other hand, this hypothetical payload were to be surrounded by a 600 mK shield in the form of cylindrical copper can attached to the Still plate, the radiation heat load would go down to the order of 10 nW, which is more than manageable. Of course, with the Still stage having a very large surface area thanks to the presence of the shield, now *it* is being heated by excessive amounts of room temperature radiation. Thus, we must also have 40-K and 4-K shields at the two pulse tube stages.

Like Russian nesting dolls, concentric cylindrical shields are attached to the cryostat’s vertically- stacked plates. Most commercial designs, like the Oxford Triton 400 mentioned in Appendix A come equipped with shields as far as the Still level. Given the relatively small surface area of payloads that can fit in such a refrigerator, colder cans are not necessary. The CUORE setup, being significantly larger, features thermal shields at all levels down to base, which can be seen in the cutaway rendering in Fig. 5.6.

5.3 Renderings of the CUORE Cryostat




Figure 5.5: A three-dimensional rendering of the CUORE cryostat setup within the CUORE experiment “hut.” The top of the main support plate (MSP) is ~ 2 ft above the level of the third floor, where we find the Faraday room and computers. The Y-beam, which rests on minus-K springs and holds up the detector suspension while bypassing the rest of the cryostat hardware, can be clearly seen. The interior of the cryostat is accessible from within the CUORE cleanroom facility (7), which occupies the second floor. The concrete walls are on the level of the first floor, which houses the Cryomech helium compressors for the pulse tube coolers, pumping systems for the Dilution Unit and the fast cooling system, and other large and noisy peripherals. The external lead and polyethylene shield is shown in the storage/open configuration on the first floor. During normal experimental operations, the shield is lifted into position surrounding the cryostat. Image credit CUORE Collaboration.




Figure 5.6: A cutaway three-dimensional rendering of the CUORE cryostat, with key components and temperature stages/plates labeled. For clarity, the pulse tube coolers have been omitted. They are found around the perimeter of the cryostat roof, and extend down to the 40-K and 4-K plates (see Fig. 5.8). Image credit CUORE Collaboration.

5.4 Precooling the CUORE Cryostat

The CUORE cryostat features a unique example of a pulse tube (PT) based precooling system. While “dry fridges” have seen widespread use in scientific research for more than a decade, and are now beginning to find commercial applications as quantum computing moves into the private sector, CUORE is an unprecedented sub-Kelvin project in terms of size, cold mass, and time spent stably at base temperature.

While a typical cryostat comes equipped with one, or occasionally two, pulse tubes, the CUORE machine boasts five Cryomech PT415 two-stage devices. We find these cryocoolers provide ~ 1.5 W of cooling power each at the 4-K stage which is in accordance with the manufacturer specifications [93, 94]. In fact, just four of the PTs are sufficient for maintaining base temperature in the CUORE cryostat [103]. The fifth unit improves performance during the cooling down, and is a redundancy in the design that would allow us to continue operating without warming up if one cryocooler were to fail. We chose the “RM,” or “remote motor,” version of the PT415, illustrated in Fig. 5.7. In this configuration, the motorhead is housed separately from the top of the pulse tube itself, coupled to it by a 2 ft flexible stainless steel line curved into a broad upside-down U-shape. An electrical isolator is incorporated into the line. The twin ballast tanks are attached in a similar fashion. Whereas the pulse tubes are mounted directly on the roof of the cryostat (the 300-K plate) on firm isolating foam rings, we suspended the motorheads and tanks from the roof of the Faraday cage surrounding the experiment with flexible steel cables and lengths of bungee cord. This arrangement serves to isolate the experiment from the ~ 1.4 -Hz vibrations of the motors and ~ 140 -Hz vibrations of the Cryomech helium compressors. Additional passive vibration isolation of the PTs is found within the cryostat: instead of rigidly bolting each PT’s 40-K and 4-K cold elements to their respective plates, we make use of flexible thermalizers comprising a pair of gold-plated copper flanges connected by multiple gold-plated copper braids. Two of the five PTs also incorporate mixture line thermalization in the form of a thin tube spiraling around the body of the pulse tube between the 40-K and 4-K levels. The thermalization features can be seen in Fig. 5.8.

We find another layer of passive vibration reduction in the CUORE pulse tube arrangement on the other side of the motorhead. Because they take up a significant amount of space and are very noisy, the five Cryomech helium compressors are housed on the ground floor of the CUORE “hut”—together with the pumps and other large hardware, on the same level as the external shield storage and concrete support structure shown in Fig. 5.6. With the cleanroom facility occupying the second floor (Chapter 7) and access to the PT cold heads on the third, we must route the ten compressor lines up through a long vertical conduit. To minimize the amount of ~ 140 Hz vibration from the compressor traveling up together with these pipes, we first run them through a ~ 1 m³ box of sand (Fig. 5.9), which is a highly effective dissipator. Once the lines exit the conduit on the third floor, they enter the Faraday room. As they do so, they pass through a set of ceramic electrical insulators. Each pair of lines is routed to its respective pulse tube along the ceiling of the Faraday room, where they are suspended by flexible bungee cords.






Figure 5.8: One of CUORE’s Cryomech PT415 cold heads. *Right*: external to the cryostat and displaying the DU mixture line thermalization in the form of a thin gold-plated tube winding around the pulse tube body between the 40-K and 4-K stages. *Left*: inside the cryostat—the 40-K and 4-K flexible thermalizers, used for vibration isolation. The silver material seen in these photos is super-insulation (multi-layer aluminized MylarTM), which we use in the CUORE cryostat outside the volume of the IVC.

Even with vibration damping devices incorporated into the construction, pulse tube refrigerators are a source of vibrational noise at ~ 1.4 and ~ 140 Hz, which degrade the performance of sensitive detectors like CUORE. The former is especially problematic for CUORE’s slow pulses (chapter 4). Furthermore, CUORE is vulnerable to interference and beat frequencies from the five separate cryocoolers—though we made attempts to slightly adjust the frequencies *manually* so that none are exactly equal, phase was nearly impossible to control in this way with sufficient precision for effective noise reduction. We were able to achieve some improvement in reducing the *amplitude* of noise from this source by making use of linear drives for PT motor control. These devices, manufactured by Precision Motion Controls [104], cause the reciprocating motor to move smoothly, like a triangle wave, as opposed to toggling sharply back and forth. However, the linear drives also have the ability



Figure 5.9: Photograph of the 10 helium lines for the CUORE cryostat’s five PTRs passing through the sandbox, for vibration isolation, on the ground floor of the CUORE hut. The “naked” ends of the lines lead to the compressors, which can be seen in the background. The ends sheathed in black foam go up to the top of the cryostat on the third floor.

of precise computer control of position, and hence phase. CUORE implemented a novel solution of dynamically setting the PT phases using a LabView program with a feedback algorithm connected to pressure gauges on the PT compressor lines. This system offers us the potential to reduce the PT-related noise by up to an order of magnitude at low frequency [103]. The feedback mechanism was first operated, though not fully optimized, during the acquisition of CUORE’s second physics dataset (“dataset 2” in this thesis), in the summer of 2017. It will be tuned properly for the next phase of data taking commencing in the spring/summer of 2018.

To precool all of the hardware below the 4-K plate, the CUORE cryostat uses a “Leiden-style” arrangement with an IVC pumped full of helium exchange gas. Even with so much available cooling power, however, precooling CUORE’s ~ 1 ton detector and 7.5 t of cold lead shielding to 4 K would take months if we were to use this system in its standard passive mode. To address this issue, we employ a unique fast cooling system, or FCS. A second cryostat, featuring a large copper heat exchanger and three Gifford-McMahon (GM) coolers—another

type of mechanical cryocooler—rapidly precools helium gas to ~ 20 K, at which point we use a large Roots blower to circulate it through the IVC. Below ~ 20 K, the helium exchange gas is static. Thanks to the FCS, CUORE is able to precool to ~ 4 K in approximately three weeks.

5.5 The Dilution Unit

The Dilution Unit (DU) of the CUORE cryostat is a custom device produced for us by Giorgio Frossati of Leiden Cryogenics BV [105]. This outfit is known for their high-power, innovative dilution refrigerator designs, and the CUORE DU is a supreme example of this. Firstly, this unit is simply very large. Fig. 5.10b, a photo of the DU in the test cryostat before it was installed in the CUORE cryostat, illustrates this to good effect. The test setup is essentially a standard Leiden cryostat, and its plates are $\sim 30 - 35$ cm in diameter—still slightly larger than the majority of laboratory-size dilution refrigerators. The scale of the Still, heat exchangers, and Mixing Chamber can be readily seen. The mixture contains 120 L of ^3He , and ~ 4 times that amount of ^4He . We circulate this large volume of gas using a pair of Adixen ATH-3200 turbomolecular pumps, each of which has a pumping speed of 2500 l/s and throughput of >84.5 mbar l/s for helium [106]. These are backed by two Roots pumps and a compressor in series. Like in the pulse tube system described in Section 5.4, we house these pumps on the ground floor of the CUORE experiment hut, with the two turbos mounted near the ceiling (approximately level with the top of the concrete structure in Fig. 5.5). When not circulating, the helium mixture is also stored here (^3He and ^4He in separate volumes).

The CUORE DU also features a novel design of its main impedance. There are, in fact, two primary capillaries in parallel, both equipped with valves. Generally, only one is used at any given time, but having the ability to switch from one to the other in the event of a line clog is very beneficial to CUORE, which demands slow, lengthy cooldowns and even more prolonged stable base temperature operation. Additionally, both impedances incorporate tiny spring-loaded needle valves that, if enabled, automatically vary the flow rate of the mixture depending on temperature. While the refrigerator is cooling down, the impedance is in the “high-flow” regime of about 8 mmol/s, maximizing the cooling power. As the temperature approaches base, the pen spring inside the valve naturally stiffens, constricting the flow to approximately 1 mmol/s. This reduces the quantity of warmer ^3He impinging on the Mixing Chamber, and thus the heat load, allowing us to achieve a lower base temperature if necessary. As of this writing, we prefer to operate CUORE with a wide-open impedance at all times, because our operating temperature is limited by NTD readout optimization.

The manufacturer specifies a cooling power of ~ 2 mW at 100 mK—the CUORE DU is a special uprated version of Leiden’s standard 1.2 mW device. During testing, we have observed performance consistent with or surpassing this—3 mW at 123 mK at the Mixing Chamber [94, 95]. We are able to reach a base temperature of nearly 5 mK, and achieve surplus cooling power of ~ 5 μW with the Mixing Chamber heated to 10 mK by a resistive

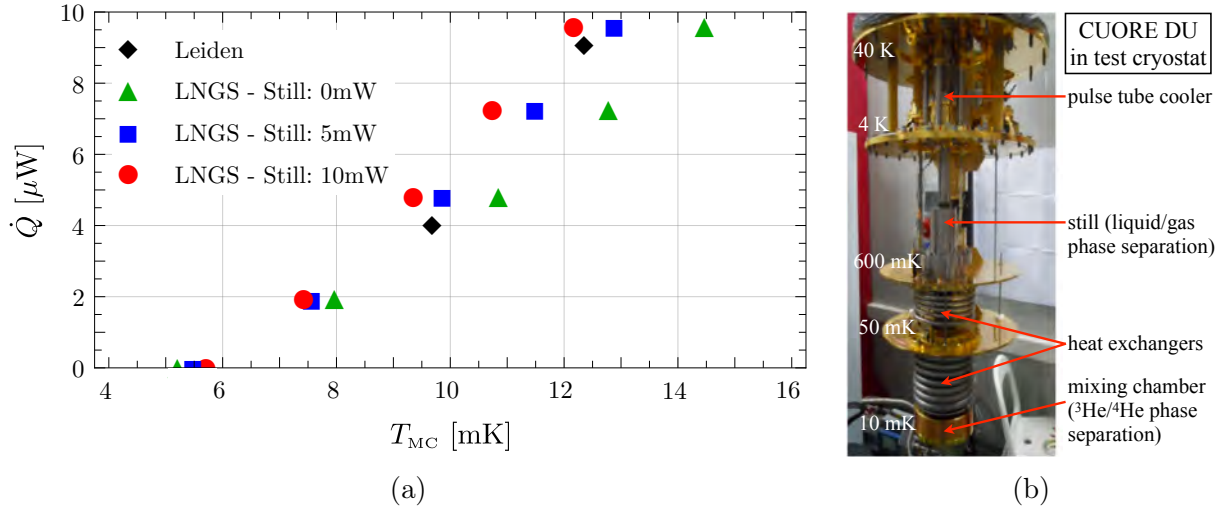


Figure 5.10: **(a)** The cooling power of the CUORE DU, measured at the Mixing Chamber plate, as a function of temperature, from four different tests conducted at the Leiden facility and at LNGS. The three LNGS tests correspond to three different powers—0, 5, and 10 mW—applied by the Still heater. Higher temperature inside the Still increases the flow rate of ^3He through the DU, thus increasing the refrigerator’s cooling power, but also the heat load on the Mixing Chamber plate. Thus, we are able to achieve the lowest base temperature—just over 5 mK—with no heater power on the Still, but have more surplus cooling power available at higher temperatures when heater power is applied [107]. **(b)** A photograph of the CUORE DU with components and temperature stages labeled. For clarity, we show it in the test cryostat, before installation in the CUORE cryostat. The plates of this cryostat are still oversized compared to most typical laboratory-scale units, giving us an idea of the immense size of this DU. Image credit CUORE collaboration.

heater [95]. This performance is plotted in Fig. 5.10a. As is the standard practice, our DU is equipped with a Still heater, allowing us to adjust the flow rate of the ^3He by raising and lowering the temperature at the liquid-gas boundary (see Section 5.2). As Fig. 5.10a clearly demonstrates, a higher flow rate (higher power) allows us to access similar cooling powers at lower temperatures. The greater quantity of relatively warm ^3He flowing in, however, also places a heat load on the Mixing Chamber stage. As such, the minimum base temperature is achieved when the power on the Still heater is zero and flow is minimized—the same rationale as the afore-mentioned variable impedance. Thanks to the CUORE cryostat’s great mass and the large volume of mixture in the DU, the base temperature is very stable. Fig. 5.11 shows the base temperature of the tower support plate (TSP) over a 12-hour period during a test run before the CUORE detector was installed, with no heater power applied on the Still level and no feedback temperature control on the Mixing Chamber level.




Figure 5.11: CUORE cryostat base temperature over a 12-hour period during a test run. All hardware besides the detector itself is in place. No power is applied on the Still heater, and PID feedback temperature control on the mixing Chamber level is switched off.

5.6 Unique Structural Features of the CUORE Cryostat

Suspension and Vibration Isolation

The structure of the CUORE cryostat is well illustrated by the renderings in Figs. 5.5 and 5.6. Overall, this unique machine appears as simply a scaled-up version of a more normal, laboratory-size dilution refrigerator. However, it features a number of unique structural features aimed at improving experimental performance through vibration isolation and background reduction.

Vibrational sources of noise are particularly problematic for bolometers. At sub-Kelvin temperatures, heat capacities and thermal conductivities of materials are minimized, and surplus cooling power from the dilution refrigerator is low, so even the smallest amount of mechanical movement results in the deposition of sufficient heat energy to disrupt the detector. We have already discussed the major and obvious issue of noise from the pulse tube refrigerators, and mitigating features incorporated into the precooling system itself, in Section 5.4. Beyond this, the CUORE cryostat has a novel suspension that seeks to isolate the detector from vibrations, regardless of the source.

Whereas most dilution refrigerators make use of rigid connections between the plates in the form of vertical or diagonal stainless steel or G10 struts, in the CUORE cryostat each plate is attached to the cryostat roof/300-K plate independently. This arrangement is demonstrated in Fig. 5.12. The roof itself is suspended from the main support plate (MSP), which is the red-colored four-legged structure seen in Fig. 5.5. Higher/warmer plates have openings with ample clearance to allow the stainless steel suspension rods of the lower plates to pass through. Furthermore, we make use of Cardan joints (aka universal joints, essentially double gimbals, not unlike those on the drive shaft of a rear-wheel-drive car) for the attachment of said suspension rods to the plates, with thermalization achieved using copper wires or braids. Thus, the plates are free to move with respect to each other, dissipating vibration. The plates are still not completely decoupled: functional components, including the DU, pulse tubes, and 40 mm wiring tubes connect them to each other. These attachments,

however, are not load bearing (low tension), and much less rigid than structural supports would be—the wiring tubes incorporate bellows, the pulse tubes are thermalized with copper braids, and the DU is inherently flexible.




Figure 5.12: Side-view cutaway rendering showing the CUORE cryostat’s plate suspension and cold lead shields. Only the bottom third of the Roman lead side shield (4 K) is shown so that the detector, TSP, top lead shield (50 mK), Mixing Chamber plate (10 mK), and cold plate (100 mK) can be seen. Image credit CUORE Collaboration.

The ~ 2.5 -ton top lead shield and the ~ 1 ton CUORE detector itself, via the tower support plate (TSP), are coupled to their respective suspension rods with loops of Kevlar ropes. At each anchor, a ring of rope is passed around a pulley-like wheel on either end. With three attachment points on both the lead and the TSP, this arrangement gives us another level of flexibility beneficial not only for noise reduction, but also seismic safety and reducing mechanical strain.

The most advanced component of the CUORE detector suspension is located outside of the OVC, on top of the MSP. Unlike the rest of the internal cryostat plates, the TSP, to which the detector towers are rigidly mounted, is not suspended from the 300-K plate. Rather, the

three detector suspension rods penetrate the cryostat roof through vacuum-tight steel tubes, and are joined to the “Y-beam” just (gray Y-shape support just above the center of the MSP in Fig. 5.5). The three ends of the Y-beam each rest on a Minus-KTM spring—a patented vibration isolator based on negative-stiffness mechanism (NSM) [94, 108]. In these devices, a stiff conventional spring supports the large weight of the detector, while the NSM—which behaves like it has a negative spring constant—offsets this, thus resulting in a very smooth suspension with a nearly-zero net stiffness without compromising weight bearing qualities. Thus, the heavy detector is essentially floating. A horizontal motion isolator is connected in series with this vertical one, followed by a third isolator for tilt motion [109].

The Minus-KTM suspension is one of the most effective vibration reduction options currently available for large-scale scientific equipment. It is particularly effective for low frequency motions, which is especially important for CUORE. With 1-2 second-long pulses (see Chapter 4), we are particularly vulnerable to fluctuations on the order of a Hertz, since they cannot be effectively filtered out. Slower motions could significantly interfere with the NTDs’ baselines, degrading detector performance to the point that we are unable to take care of it with stabilization during analysis, or even take the thermistors out of their effective working regime. We are concerned about this given the unique design of the cryostat structure with rods and Cardan joints. While advantageous for the reasons described above, that setup is vulnerable to pendulum-like low frequency swinging. Thanks to the Minus-KTM, we can keep the detector stable even if such motions do develop. We have been able to confirm this empirically, with working personnel walking on top of the MSP, causing the entire cryostat to move visibly, without raising the temperature as measured by a noise thermometer at the Mixing Chamber.

Cold Lead Shielding

Virtually all passive rare event searches require extensive shielding from external background radiation, typically consisting of large liquid tanks or lead screens. For cryogenic experiments like CUORE, this question presents a unique challenge. The roof of a dilution refrigerator hosts a range of bulky and sensitive hardware, and must remain accessible. Furthermore, the interior of a dilution refrigerator (Fig. 5.6) is a large volume filled with a lot of hardware besides the detector itself, all of it a source of background. In the case of CUORE, the primary background sources are the structural components [74, 80], including the massive cryostat plates and “cans.” The optimal shield would be located as close to the detector as possible, screening it from γ rays emanating from the cryostat as well as the outside world, from all sides. Necessarily, this shield would have to be cold.

Standard commercial cryostats provide minimal spacing between thermal shields in place, strongly limiting the amount of cold lead that could be installed. Furthermore, lead is dense and massive and superconducting below 7.2 K, making cooling it down a challenge. The CUORE-0 cryostat, for instance, includes only a couple centimeters of cold lead shield on the bottom and sides, relying almost exclusively on the warm external lead that does nothing to protect the detector from the less-than-radiopure refrigerator. Our R&D cryostat in Berkeley




Figure 5.13: (*Left*) Side-view technical drawing and (*right*) cutaway three-dimensional rendering of the CUORE cold lead shielding. The ~ 6 -cm thick 4-K Roman lead side and bottom shield, as well as the 60-cm thick 50-mK modern lead top shield are shown in gray, while the brown components represent the cryostat plates and thermal radiation shields. The modern lead is a greater source of background radiation from ^{210}Pb and its daughters compared to the Roman variety, but we stop the α s and most of the β s with 5 cm of clean NOSV copper on the bottom of the top lead shield. Image credit CUORE Collaboration.

lacks cold shields entirely, with only an external one available. This situation is perfectly acceptable for proofs of concept and prototypes, but the ton-scale CUORE experiment would fail to reach its sensitivity goals and $\sim 6\times$ background reduction [74, 80] with this limited level of protection.

Thankfully, the ample surplus cooling power provided by the CUORE cryostat’s five PTs and world-leading DU, together with the custom-designed plate and shield geometry, give us the opportunity to include significantly more cold shielding *ab initio*. We have a ~ 2.5 -t, 60-cm “tablet-shaped” top shield directly above the detector, and a ~ 6 -cm cylindrical side and bottom shield located between the 4-K and 600-mK thermal shields, weighing about 7.5 t [94, 95]. These are shown in Fig. 5.13. The top shield needs to be significantly thicker than the sides and bottom, because the former are also screened by the massive external (warm) shield shown in Fig. 5.5, which does not cover the top. The top lead is located below

the Mixing Chamber (10 mK) plate of the cryostat, but is thermalized to the 50 mK stage with flexible copper strips. It is not a monolithic casting, but rather consists of five ~ 6 -cm slabs with thin copper sheets sandwiched in between. These copper sheets, connected to the thermalization strips, assist in the cooling. We do this because the DU provides some two orders of magnitude more cooling power at that temperature, while the lead doesn't actually have to be as cold as the detector below—just cold enough to not heat up the 10-mK components with its black body radiation. The detector itself is thermalized to the Mixing Chamber plate in a similar way, with flexible connections going up around the lead. Below the five slabs of lead, the tablet is covered with a ~ 5 -cm layer of clean copper. This serves to protect the detector from background radiation stemming from the ^{210}Pb and its daughters found in the lead.




Figure 5.14: A comparison of the radioactivity of (1) modern common lead, (2) modern low-activity lead, and (3) ancient Roman lead. Figure from [110].

The side lead shield is comprised of crescent-shaped bricks, also interspaced with thin copper thermalization aids. Unlike the top shield, this one is made from Ancient Roman lead—among the most radiopure available. The original ingots were naval ballast weights recovered from a two thousand year-old shipwreck in the Mediterranean Sea. They were recovered with the specific purpose of the CUORE shield. Each ballast brick contained a stamp, which was sawed off and preserved for its archeological value. The remainder of the bricks were melted down and recast at a specialized foundry in Dresden, Germany. Obviously, Roman lead dates to before the year 1945 and the era of nuclear testing, which is a source of uranium-chain contamination in modern lead and many other materials. Furthermore, the Romans' smelting process was designed to remove silver. Unintentionally, this also separates out natural uranium found in the ore, and is the original source of uranium-chain contaminants therein. ^{210}Pb has a 22-year half life, so over the course of two thousand years on the sea floor the vast majority present in the ingots has decayed away. Since the shipwreck was resting above ground, no significant amount of fresh ^{210}Pb from radon gas was introduced. The cleanliness of the Ancient Roman lead—not only with respect to common modern lead, but also specially-engineered modern clean lead—is clearly demonstrated by the comparison of spectra in Fig. 5.14, taken from a study conducted for the CUORE shielding [110].

Chapter 6

A Cryogenic PID Temperature Controller

6.1 Introduction

Successfully operating a bolometric detector array like CUORE demands the careful regulation of the device's temperature. As explained in Chapter 4, a bolometer's energy threshold, sensitivity, and intrinsic energy resolution all stand to benefit from minimizing the heat capacity of the TeO_2 crystal absorbers [111]. According to the Debye Model, this value is proportional to T^3 , suggesting that the minimal achievable temperature is best. The resistance of the NTD thermistors that sense the bolometer pulses, however, increases strongly as the temperature falls:

$$R_{NTD} = R_0 e^{\sqrt{T_0/T}} . \quad (6.1)$$

The CUORE electronics make use of a load resistor, the resistance of which must be significantly higher than that of the NTD to achieve the highest quality readout with a good signal-to-noise ratio. Our choice of operating temperature must take this into account as well. In fact, thus far we have found the optimum point to be ~ 11 mK. As such, we need a temperature control system that allows us to:

1. Scan the base temperature through a series of stable steps between ~ 5 and ~ 25 mK to determine the value for maximum detector performance.
2. Set and maintain this temperature over the course of a prolonged run.
3. Even though we stabilize the bolometers' baselines during data production/analysis, this temperature needs to be kept steady, minimizing or eliminating fluctuations resulting from vibrations and other events, especially the deployment of the Detector Calibration System (DCS). After all, there is little point in calibrating the detector if its temperature and state during the calibration run significantly differs from that during background measurements.

A good solution to this problem is to implement feedback control of the cryostat temperature. CUORE-0 had a clever but basic system in place—a heater and a temperature-dependent resistor were mounted on opposite ends of the tower, and current was driven through them in a loop. As temperature changed, so did the impedance of the resistor, and in turn the current going through the heater [57]. Probably the most popular feedback control algorithm currently in use is the PID (**p**roportional-**i**ntegral-**d**erivative) loop. Taking into account the current deviation of the temperature from the desired value, the integral of the past history of that deviation, and a differential prediction of future behavior, it is able to accurately approach and maintain the setpoint. The design and creation of a PID controller for CUORE was prioritized by the collaboration, and several systems were delivered in parallel. In this chapter, we discuss a software-based setup written in National Instruments’ LabView graphical programming environment, and making use of a combination of existing CUORE electronics and off-the-shelf commercial hardware. This system is not currently the main working controller for the experiment, but it does serve as a backup, saw limited use during testing, and has sufficient flexibility for future applications in cryogenic detector R&D and other areas.

6.2 PID Control Principles

A PID mechanism is a widespread variety of a feedback control algorithm—a system which uses the readout data from a sensor measuring an environmental parameter or, process variable $PV(t)$, to adjust the value of a control parameter, or manipulated variable MV , to manipulate $PV(t)$ such that it reaches and stabilizes at some set point value SP . In our case, the SP is the desired temperature in our cryostat, so $PV(t)$ is the readout of a thermometer measuring the true temperature over time, and MV is the current/power setting on a resistive heater. The acronym PID stands for **p**roportional-**i**ntegral-**d**erivative, describing the three terms of the algorithm’s feedback calculation, given in Eq. 6.2:

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau) d\tau + K_d \frac{de(t)}{dt} , \quad (6.2)$$

where $u(t)$ is the controller output (MV), and $e(t)$ is the difference between the intended set point and true value of the process variable at time t — $e(t) = SP - PV(t)$. The constants K_p (proportional), K_i (integral), and K_d (differential) are adjustable parameters tuned by the user. Having all three of these terms allows us to adjust our MV (heater) taking into account not only the current value of the error (temperature difference between the measurement and SP), but also the past behavior of the system and a prediction of the future based on the current trend [112, 113].

The proportional term,

$$P_{\text{out}} = K_p e(t) , \quad (6.3)$$

simply sets the MV/heater power to the current value of the PV/temperature error multiplied by a constant. While any PID controller is typically empirically tuned in situ, this gain tends to be dominant, while the integral and derivative terms provide corrections. In fact, the proportional algorithm is able to function as a proper feedback controller on its own, though with limitations. Namely, if the gain is set to a high value for a fast response, we find that the controller tends to overshoot its setpoint, respond quickly to this again-large error in the opposite direction, overcompensate a second time, and thus proceed to oscillate before finally converging. If the gain is set too high, the convergence time could be effectively infinite and the system could be unstable. We can minimize oscillations by setting the gain to a low value, but then the correction is slow, which can be bad if environmental changes are sufficiently frequent and severe. Furthermore, the oscillations are never completely gone—since P_{out} in Eq. 6.3 is zero when error is zero, even a well-tuned and “converged” proportional-only controller always has a steady state error [112].

Logically, we can afford to increase our proportional gain for a faster response, as well as removing the steady state error, if we can take into account not only the current value of the PV (temperature) error, but also its history. This is done by integrating over time in the “I” term:

$$I_{\text{out}} = K_i \int_0^t e(\tau) d\tau . \quad (6.4)$$

The value of the integral in Eq. 6.4 depends on the duration of a deviation, not just its magnitude. Multiplying this integral by a second gain K_i and adding it to the proportional term output compensates for accumulated past error that we failed to correct immediately. Since the integrated error is non-zero even when the PV = SP, we eliminate the proportional term’s steady state error. In a well-tuned controller, K_i tends to be quite a bit smaller than K_p , because accumulated error from the integral can *increase* the controller’s overshoot and settling time [112].

While the integral term helps us consider the condition of our system in the past, with the “D” term we can use the time derivative of the error to predict its future behavior given current settings:

$$D_{\text{out}} = K_d \frac{de(t)}{dt} . \quad (6.5)$$

If the error is increasing or decreasing more rapidly, this term contributes a greater correction and helps to decrease the time it takes for the system to stabilize. In an ideal system, increasing the derivative gain decreases the overshoot and settling time, though there is no effect on steady state error, and only a minor effect on rise time and final stability [112, 113]. Since physical sensors providing the PV(t) feedback information, such as our thermometers, are non-ideal and have noise, real-world PID controllers usually require some kind of low-pass filtering to make the derivative term work correctly. Even so, only a minority of PID controllers in use are configured with non-zero values of K_d because the derivative action often causes challenges in non-ideal systems. Unless shortening the settling

time and overshoots is absolutely critical, the benefits of having a fully functioning derivative contribution are not usually worth the effort. As of this writing, the PID loop in use with CUORE uses only the proportional and integral parts of this algorithm.

6.3 Heaters and Thermometry for Temperature Control in CUORE

In the CUORE experimental setup, we utilize dedicated cold hardware for feedback temperature control at the 10-mK stage. As discussed in Chapter 5, the CUORE cryostat has two plates at base temperature: the actively-cooled mixing chamber (MC) plate, which is the final stage of any dilution refrigerator, and the tower support plate (TSP) below the ~ 60 -cm lead shield, thermalized to the MC and holding up the ton-scale detector. For cryostat monitoring, all of the 50-mK and base temperature components are equipped with RuO_2 thermometers read out in a four-wire arrangement by resistance bridges. Since at base temperature, always below 20 mK, these resistive sensors no longer have reproducible calibration curves, we equip the MC plate with a noise thermometer. Manufactured by the Magnicon company of Germany, this device uses a SQUID (superconducting **q**uantum **i**nterference **d**evice) array to read out the temperature-dependent Johnson noise on a small, well-shielded piece of copper [114]. The noise thermometer is very sensitive in the ultra-low temperature regime, but unfortunately it requires about a minute of integration time between readings, because sufficient statistics must be accumulated to construct a sufficiently precise Fourier spectrum of the noise for fitting. This low-speed readout is perfectly adequate for temperature monitoring, but is likely to be problematic for a PID controller. Since no commercial resistance thermometer is currently available that is sensitive and reproducible in the 5 – 20 mK range¹, we decided to fabricate our own thermometers. The standard CUORE NTDs (**n**eutron **t**ransmutation **d**oped germanium thermistors, see Chapter 4) are already optimized for sensitivity range. When the NTDs for CUORE-0 and CUORE were fabricated and characterized, several batches were prepared with varying levels of neutron irradiation. One of these lots, carrying the internal designation NTD40A, received a larger dose than the others and consequently has a lower electrical resistance. Rejected for this reason for bolometer instrumentation and pulse readout, these thermistors are a good choice for the role of thermometers.

The custom thermometers are deployed in the immediate vicinity of the detector, and so radiopurity was a major concern during their design. The NTD40A chips are hand-glued with Araldite epoxy on top of a thin piece of plastic film inside cylindrical cases cleanly machined from NOSV [96] copper. The finish of the copper matches the mounting plate: the MC thermometers are gold-plated, while those on the TSP are plasma-cleaned

¹Historically, carbon resistance thermometers functional in this range were available, but have become expensive and uncommon since commercial resistor manufacturers switched from carbon to thin-films. One capable resistor is available on the market, but is challenging to adapt for thermometry [115].

the same way as the detector frames. Solder beneath the lead shield is strictly forbidden in CUORE, so electrical connections were made by crimping the wires—25 μm gold bonding wires from the chips and Constantan alloy cryogenic cables outside—inside a thin copper tube. These holders were designed by Nicholas Chott of the University of South Carolina, and are illustrated and described in Fig. 6.1. Using naked silicon chips of $O(100\text{ k}\Omega)$ in place of the NTDs, heaters are fabricated in an identical way.




Figure 6.1: **(a)** A photograph of two PID heaters and one NTD thermometer, in their special copper holders, on the tower support plate (TSP). **(b)** A schematic (top view, cover off) of a CUORE thermometer or heater holder. The chip (dark gray) is adhered with Araldite epoxy on top of thin insulating plastic sheet (light gray) inside the clean NOSV can. Wire connections are made with a pair of copper crimping tubes, epoxied in small holes in the walls of the can, to avoid having radioactive solder near the detector.

Three NTD thermometers and six heaters dedicated to PID temperature control are mounted on each of the two base temperature plates, for flexibility and redundancy. The arrangement of the devices on the MC plate and TSP are shown in Fig. 6.2a and Fig. 6.2b, respectively. Two heaters are “paired” with each thermometer spatially, but this does not necessarily dictate which thermometer-heater pair is used for feedback temperature control, and multiple devices could be used at once.

The readout of the NTD thermometers is performed by the standard CUORE electronics as used for the bolometer channels. This arrangement has some drawbacks compared to conventional thermometry for temperature monitoring—the two-wire electrical connection is problematic for accurate measurement of resistance and absolute temperature values com-

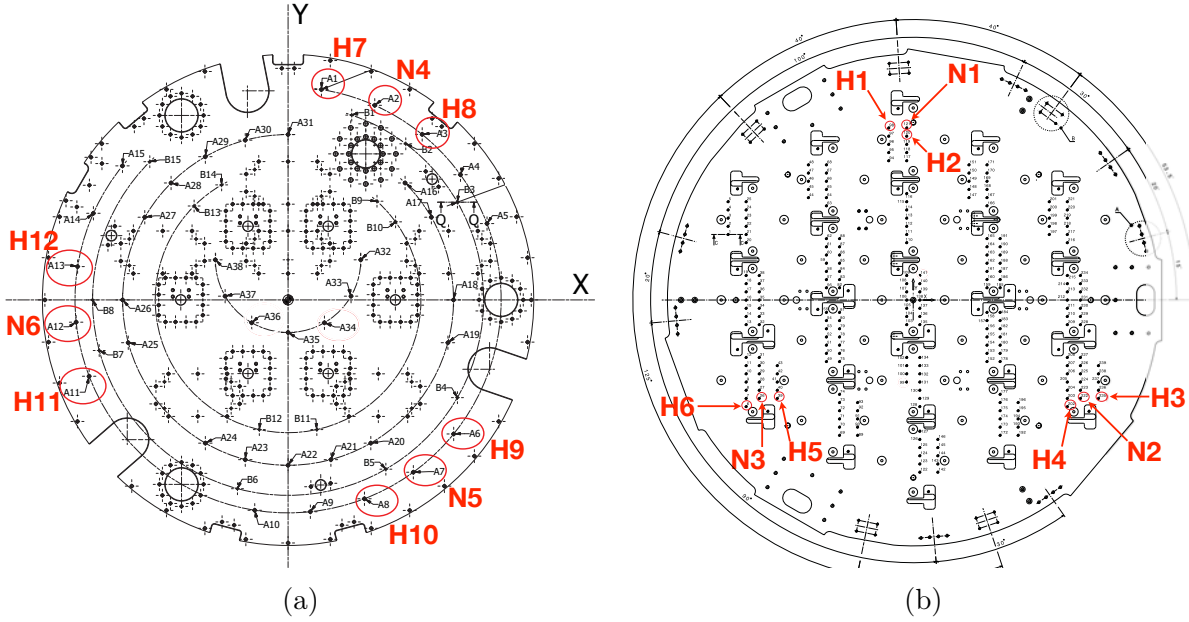


Figure 6.2: **(a)** Mixing chamber and **(b)** tower support plate (TSP) drawings with NTD thermometer (N1 – N6) and heater (H1 – H12) locations. Both of the ~ 10 mK plates have three thermometers and six heaters each, for redundancy. The LabView-based PID controller discussed in detail in this chapter is capable of operating up to two heaters and reading out up to eight thermometers at one time.

pared to four-wire bridge readings. For a PID loop, it is quite good, however. As long as we are able to track the true value of the temperature with a properly calibrated device such as the noise thermometer, the actual feedback for stabilization could be done with any reproducible temperature-dependent parameter—in this case, the voltage across the NTD. The CUORE electronics' readout of the NTD signal is continuous and can be digitized at a high rate—averaging groups of 1000 samples acquired at a 15 kS/s rate is typical in the PID application—thus not limiting the speed of the controller.

The planned primary PID system for use with the experiment will be operated on a modified CUORE pulser board (a microcontroller). As of this writing, the pulser firmware has yet to be adapted, so a temporary system consisting of a DAQ² and PID algorithm is being run on a PC connected to the CUORE electronics. The LabView- and Keithley-based setup discussed in this chapter was developed for the cryostat testing phase, and currently serves as a backup. In any of these cases, the same cold hardware and thermometer readout is used, though the methods of digitizing the NTD signal and biasing the heaters differs.

²Data Acquisition system

6.4 Room Temperature Hardware for a LabView-based PID Setup

The wiring of the CUORE PID NTD thermometers and heaters is unified with that of the bolometers. Ribbon cables comprising 13 twisted pairs of 36-AWG NbTi³ wires lead out of the cryostat inner vacuum chamber (IVC), through 40-mm bellows tubes traversing the outer vacuum chamber (OVC), to breakout boxes on the cryostat roof with vacuum-tight 27-pin Fischer connectors (one per ribbon cable). Heaters and thermometers are always on separate cables/connectors. In the case of the standard setup, shielded cables connect the NTDs to the front end boards (FEB) of the CUORE electronics via 50-pin Centronix cables. Immediately after the FEB comes an analog Bessel filter, followed by the digitizer and CUORE DAQ with the “Apollo” software. The heater chips mounted on bolometers are used to generate controlled pulses for baseline stabilization, and connect directly to pulser boards [65].

The thermometer NTDs are connected to the FEBs and Bessel boards in the usual fashion. For the LabView-based CUORE backup PID controller, we make use of multi-purpose DB25 connectors located between the Bessel boards and the CUORE DAQ to break the signal out and take it to a dedicated digitizer. This configuration is illustrated in the block diagram of the PID system in Fig. 6.3. The digitizer is a relatively inexpensive National Instruments NI 6321 16-bit PCIe card [116]. This allows us to read in up to eight differential analog channels, with a maximum aggregate sampling frequency of 250 kHz and an input voltage range of ± 10 V. This card interfaces directly with the LabView control software.

The PID heaters are DC-biased using Keithley 6220 precision current sources [117]. The standard CUORE heater electronics are optimized for rapid pulses—hence the need to modify their firmware for the planned fully-integrated, microcontroller-based CUORE PID loop. Meanwhile, our LabView/Keithley apparatus does not make use of any standard CUORE electronics. Its arrangement is shown in Fig. 6.3. Each Keithley’s triaxial BNC cable output is broken out into a pair of standard coaxial BNCs, which are then coupled to the heater Fischer connector via a BNC-to-Fischer breakout box. Triaxial or “triax” BNC cables have a similar appearance to the standard coaxial variety, but have three layers of conductors. This allows a single cable to be operated differentially while still shielded by a grounded jacket, making these connectors a popular choice in precision current applications and standard on Keithley devices. Two parallel coaxial cables are necessary to achieve equivalent performance, which we do to save funds—standard coaxial BNC cables are much less expensive—and utilize existing infrastructure—the BNC breakout box existed prior to this system, previously used for tests and diagnostics. The Keithley’s are controlled via GPIB protocol, by the same PC that hosts the digitizer card for the NTD thermometers. At the time of the system’s construction and testing, as well as currently, two Keithley 6220 current sources, driving one heater each, are available for CUORE PID control use. Technically, the system is able to

³a superconducting alloy of niobium and titanium, NbTi has a relatively high critical temperature of ~ 10 K and is an excellent material for cryogenic wiring.

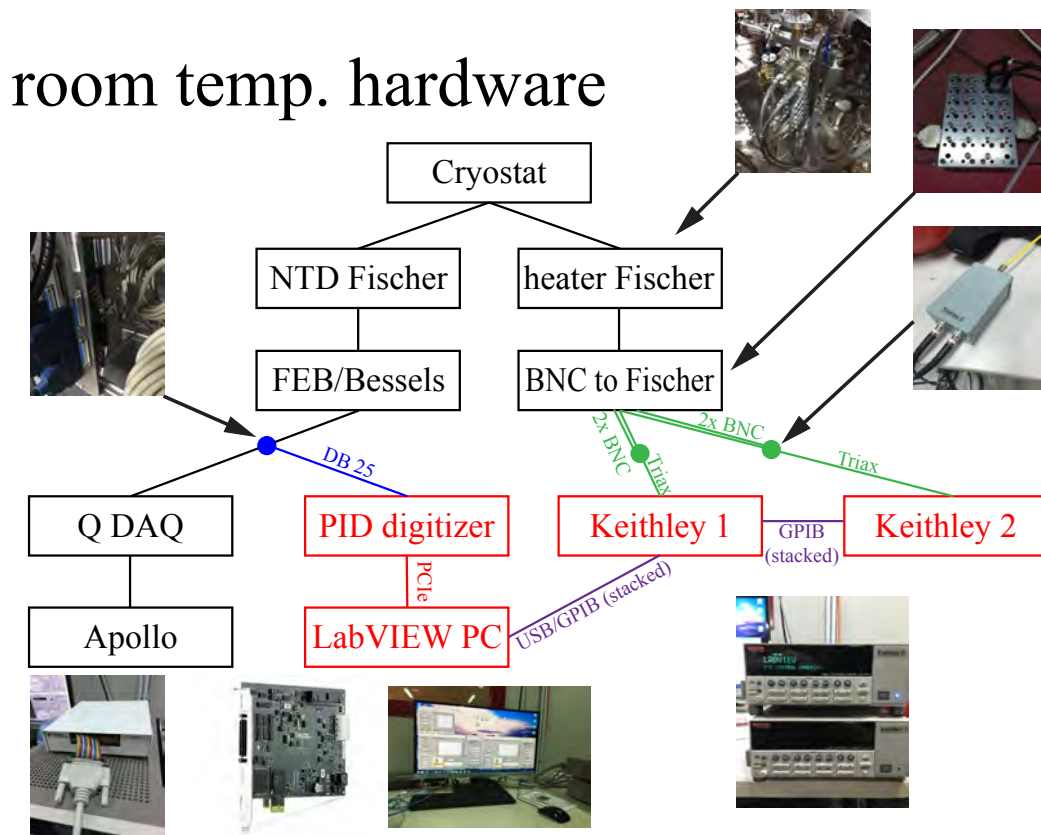


Figure 6.3: A block diagram of the room temperature hardware of the CUORE LabView-based backup PID system, with photographs of key components. Black lines and boxes refer to the existing CUORE electronics and DAQ infrastructure, while the colored boxes and lines represent the added dedicated PID hardware.

accommodate any number of Keithleys.

6.5 LabView PID Controller Software

The National Instruments (NI) LabView visual programming language is highly suitable for the rapid writing and deployment of DAQ and/or instrument control software with a graphical user interface (GUI). The NI Virtual Instrument Software Architecture (VISA) standard includes many templates and pre-written tools for interfacing with GPIB⁴ devices (including our Keithley current sources) [118]. LabView interacts very cleanly with NI's own DAQ hardware, in this case our NI PCIe 6321 card. Consequently, this environment is

⁴General Purpose Interface Bus, also known as IEEE-488 and HP-IB

a natural choice for creating CUORE’s first PID controller during the cryostat testing and commissioning phase.

The functions of the PID control software can be divided into three basic tasks:

1. Interface with the DAQ and read in the digitized thermometer waveform(s). With multiplexing, the NI PCIe 6321 card can read up to eight devices. The CUORE cryostat hosts six dedicated NTD thermometers for temperature control (see Section 6.3).
2. Using the thermometer measurement and user-specified parameter settings, actually run the PID algorithm and evaluate the correction (see Section 6.2). This process needs to be done for as many thermometer channels as are being used for control, in parallel.
3. Interface with the Keithley(s) via GPIB, configure the current source(s) and set the appropriate output current(s), as determined by step 2, to drive the heater(s).

We notice two features of these three processes that directly influence our program design. Firstly, besides needing to pass the thermometer readout from step 1 to step 2, and subsequently the PID setting from step 2 to step 3—a total of just two variables—the three tasks are computationally separate (and different). Secondly, we have the question of instances. Namely, the DAQ card can only be safely accessed by one process at a time, and even if we use multiple thermometer channels, we would do this by multiplexing the single digitizer. The tasks of calculating the PID correction and applying it to the current source, conversely, are completely independent channel-to-channel because the electronics are separate, and can be run in parallel. Finally, in this particular case, we have the hardware capacity for six NTD thermometers (PV values), but only two heater channels (MV values), which could be paired with any of the thermometers. It could be potentially advantageous to run more thermometers than heaters at a given time, and have the ability to switch which NTD measurement is used by the PID algorithm.

The PID software consists of three LabView “virtual instruments,” or VIs, compiled as stand-alone Windows executables. The scheme is illustrated by the flow chart in Fig. 6.4. The three servers correspond to the three basic tasks outlined above and run independently, passing sensor readings and PID settings to each other via two global variables.

The Digitizer Server (PIDDigitizerServer.exe)

A single instance of the Digitizer Server reads in all eight differential multiplexed channels of the NI PCIe 6321 board and writes them to network variables that may then be read as needed by PID control server instances or other programs. The GUI allows the user to select the digitizer device (as of this writing, there is only one real option), set the sampling frequency, and set the number of samples averaged to make a single PV (NTD voltage) reading for use by the PID. There is also a waveform graph displaying one channel at a time. An image of the interface can be seen in Figs 6.4 and 6.5a.

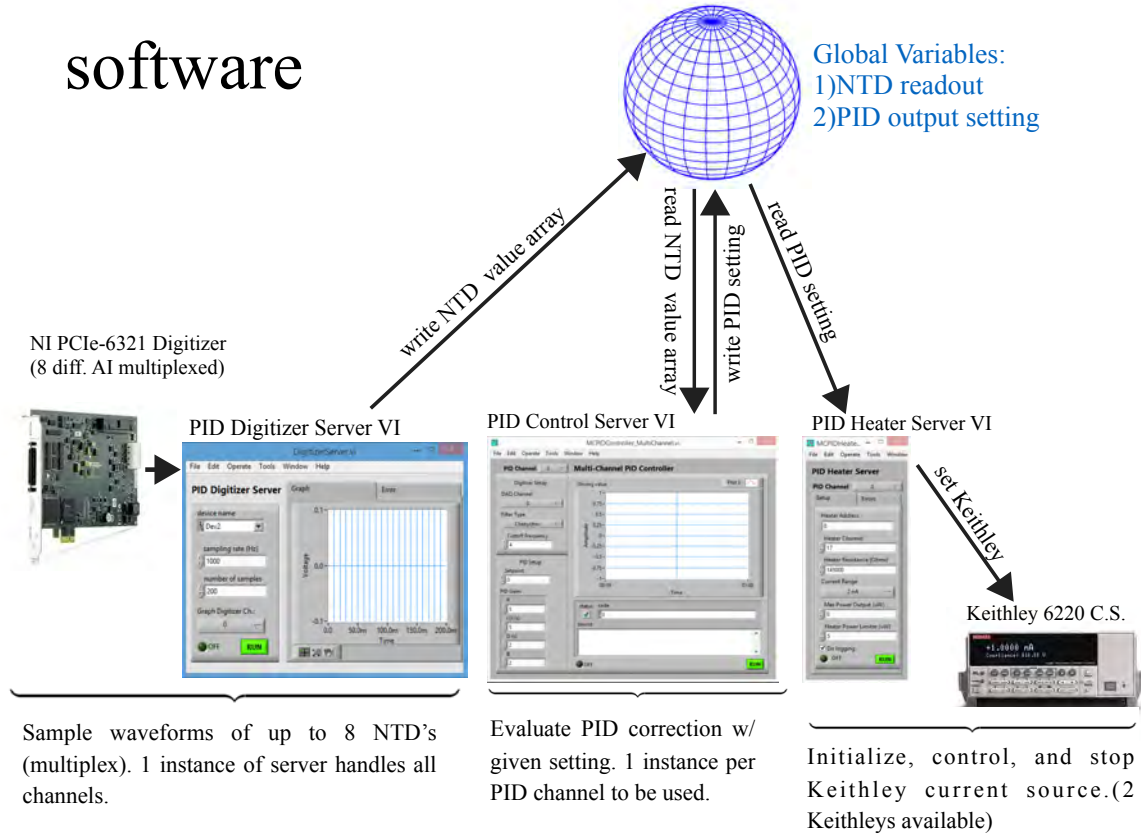


Figure 6.4: A flowchart schematic of the PID control software.

The PID Control Server (PIDControlMulti.exe)

The PID Control Server reads in network variable of one digitizer channel, and carries out PID feedback calculation to determine what the heater driving value should be. The algorithm we use differs slightly from the “ideal” PID described in Section 6.2 and Eq. 6.2. Of course, since we are doing a numerical calculation in a while loop, our implemented version is discretized. The correction is defined as

$$U_j = P_j + I_j + D_j, \quad (6.6)$$

where j refers to the number of the iteration of the loop, and we are evaluating the proportional and derivative terms as one would expect:

$$P_j = K_p e_j, \text{ and} \quad (6.7)$$

$$D_j = K_d \frac{e_j - e_{j-1}}{t_j - t_{j-1}} . \quad (6.8)$$

The integral is a modification of the trapezoidal rule:

$$I_j = K_j \frac{e_j + e_{j-1} - K_B \times \text{Excess}_{j-1}}{2} (t_j - t_{j-1}) + I_{j-1} , \quad (6.9)$$

where we have added a fourth gain parameter K_B and the variable *Excess*.

We insert this calculation to alleviate integral windup, a scenario where the integral term in a feedback controller “runs away” and accumulates a large and potentially problematic error, particularly after a significant change to the setpoint [119, 120]. For example, in the CUORE case, this could happen if the experiment is undergoing a scan to determine the optimal working temperature for the NTDs. The PID Control Server source code contains variables for the minimum and maximum correction limits⁵. If the calculated correction goes outside these constraints, it is set equal to the limit value. *Excess* is defined as

$$\text{Excess}_{j-1} = (P_{j-1} + I_{j-1} + D_{j-1}) - U_{j-1} , \quad (6.10)$$

which is trivially zero *unless* the calculated correction passes outside the limits. Depending on the value of the gain K_B , this term combats the winding-up of the integral to varying degrees.

All four gains— K_p , K_i , K_d , and K_B —are set by the user together with the setpoint in the server GUI, which is shown in Figs. 6.4 and 6.5a. As of this writing, we have not had the opportunity to properly tune this controller for stabilizing the base temperature of the CUORE cryostat. As such, there are no meaningful defaults for the gain parameters in the software.

The program writes the PID output to a global variable. As many instances of the PID Control Server as there are PID channels can be run. In addition to the four feedback loop gains and setpoint, in the GUI window the user selects which DAQ channel is to be used for the PID calculation, as well as assigning the “PID Channel” for driving the heater—thermometer/heater pairs are not locked in. Finally, one chooses a type of low-pass filter and its cutoff frequency in Hz. This functions on top of the filtering effect of averaging samples by the Digitizer Server. A Chebyshev filter with a 4-Hz cutoff is default, though a Butterworth filter, an FIR filter, or none at all can also be selected. Low-pass filtering is particularly important if we configure our algorithm with a non-zero derivative term, since higher-frequency noise fluctuations in the thermometer reading can confuse it. As with the Digitizer Server, the PID Control Server comes with a waveform graph, this time plotting the driving value as it is being written to the global variable and read by the Heater Server.

⁵These limits are intentionally not modifiable from the GUI, and their values should to be determined empirically

The Heater Server (`PIDHeater.exe`)

A single instance of the Heater Server uses GPIB protocol to operate one Keithley 6220 DC current source. When started, it sets the output current on the Keithley to zero and disables its direct manual control. Typically, there is one Heater Server for each PID Control Server running. As of this writing, we are limited to two channels because this is how many Keithleys we have available for this use. With the GUI (see Figs. 6.4 and 6.5), the user selects the PID channel to be used. The value of this parameter matches that designated in one of the running PID Control Servers. The program reads in the corresponding control value network variable.

IEEE-488 cables are “stackable,” and so all Keithleys are plugged into a single USB port in the PC. As such, while the Heater Server controls only one Keithley at the time, it has the ability to access *any* of them. For this reason, we have assigned an address and channel to each current source and must also select the correct values of these parameters in the GUI. When the software is configured and compiled, these variables are initialized. Before running, we also provide the information for the physical circuit: the resistance of the heater (145 k Ω for CUORE), the current range for the Keithley, the “heater power limiter” to set the Keithley’s compliance (just in case), and the value of the “maximum power” that is scaled by the gain from the PID calculation to determine the actual current, as long as it is within range and compliance.

Finally, the Heater Server is also responsible for logging time, channels/addresses, heater resistance, and applied current to JSON logfiles compatible with the CUORE Online Run Control website (CORC), which allows remote monitoring (though not control) of most of the experimental systems. Multiple instances of the program write to the same logfile.

Additional Software and Running

The three servers described earlier in this section form the core of our LabView-based PID system, and on their own are sufficient for successful operation. For convenience, we provide two additional pieces of software. Because a control/monitoring PC can be quite heavily loaded under some situations, we have also prepared a “light” single-channel program—`PIDControlSingle.exe`. This folds together the digitizer serving and PID control functionality into a single executable that reads in one sensor and calculates the appropriate correction. The Heater Server remains independent, and is the standard version. Under the hood, the functionality is identical to that of the full multi-channel version described above. Examples of a “full” two-channel configuration and the “light” one-channel setup can be seen in Figs. 6.5a and 6.5b, respectively.

Since the arrangement with multiple instances of multiple programs can be confusing to an unfamiliar user, we also provide the Launcher executable (`PIDLauncher.exe`). With this simple GUI, the user selects the number of channels he or she would like to operate and clicks “Launch.” At this point, the Launcher boots up the appropriate number of instances of each server and quits itself. If one channel is selected, the “light” single-channel suite is

started by default. While the use of the Launcher is completely optional, we do compile the Digitizer Server in a foolproof way that prevents running more than one instance: doing so is never necessary, while having two pieces of software attempting to communicate with the same DAQ card can cause an error or even a crash.

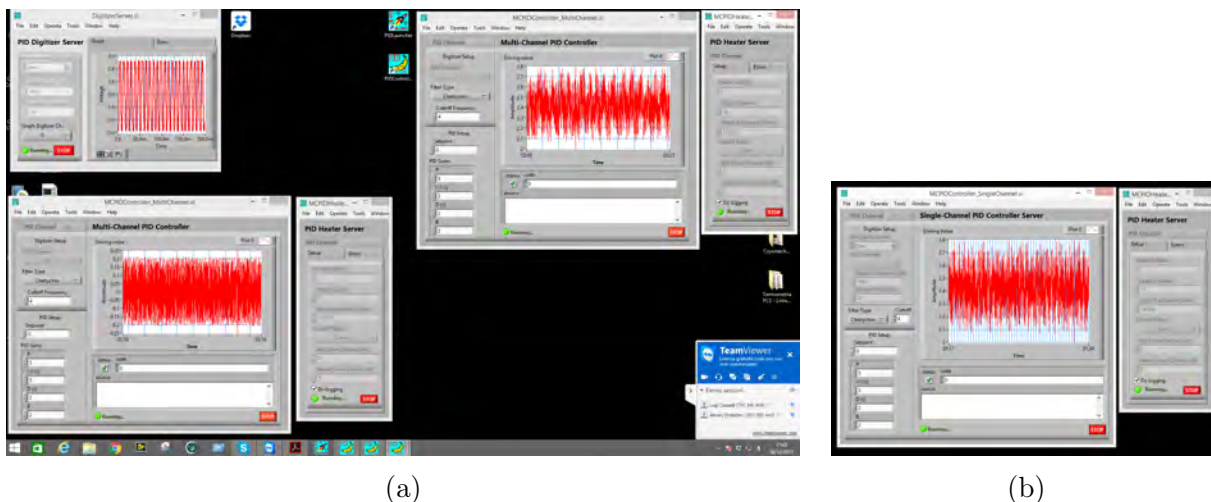


Figure 6.5: (a) A screenshot of a two-channel PID setup—one Digitizer Server, two PID Control Servers, and two Heater Servers—running on a PC, with fake signals. (b) An equivalent screenshot of the “light” single-channel setup, combining digitizer serving and PID control into one program.

6.6 Testing and Usage

The practical usage of the CUORE LabView-based PID temperature controller has been very limited thus far. Over concerns over the long-term stability of a LabView program operating on a Windows PC, it was decided by the collaboration that the permanent solution would be based on a microcontroller. The system discussed in this chapter has been assigned the role of a back-up. Additionally, it was dedicated to serve as the temperature controller during the CUORE cryostat’s third test cooldown—the first run with all cold hardware in place besides the detector itself, including the cold lead shields. Unfortunately, this cooldown was prematurely aborted without completing all of the planned tasks, as a result of an obstruction in the dilution unit (DU) mixture line.

Among the tests that we were unable to carry out was the control and stabilization of the base temperature of the mixing chamber and tower support plates with a PID loop. As the cryostat was already undergoing forced warmup, hurriedly adapted our software to interface with the hardware on the *Still* level and performed a limited test.

As we describe in Chapter 5, this component of the DU consists of an actively-pumped 500 – 700 mK chamber wherein superfluid ^4He remains condensed at the bottom, while the

^3He is actively passing through it and evaporating. The Still comes equipped with a heater used to adjust the temperature of this boundary, which affects the flow rate of the ^3He , and in turn the dilution unit's cooling power at the mixing chamber level (see Fig. 5.10a). This is an important parameter, and applying some optimum amount of heat to the Still during dilution refrigerator operation is standard practice. That said, this power is almost always constant, and not regulated by a feedback mechanism. While small changes to the Still temperature can significantly impact the cooling power of the DU *at base*, the cooling power of the $\sim 600\text{-mK}$ stage is orders of magnitude higher and its temperature remains very steady. In spite of standard practice, we had some interest in a Still PID—possibly a third role for the LabView system—to mitigate the thermal disturbances from the deployment of the source strings of the cryogenic DCS [85]. Experience has shown this to not be necessary. Furthermore, it is the opinion of the author that a temperature PID for the Still would not be the best solution even if it were. Rather, the addition of a feedback loop to manipulate the speed of the lowering of the DCS strings, combined with the action of the always-active MC-stage cryostat PID, would be better at taking care of this concern.

In any case, under the unique circumstances of Cooldown 3, stabilizing the Still temperature was our only opportunity to test the controller. If anything, given the severely reduced cooling power of the obstructed DU, the Still was a much better testing environment than it would be during standard operation. The Still level of the CUORE cryostat is equipped with two RuO_2 thermometers, read out by a PICOWATT AVS-47B resistance bridge [121]. One of these is mounted inside the Still itself, and due to its immersion in helium, gives visibly noisy readings. The second is bolted to the plate, and was chosen as the stabilization channel. The Still heater is driven by a Leiden Cryogenics Triple Current Source (TCS). The device is not optimal for PID use—meant to control the temperature of the *helium*, it is mounted in the Still like the first thermometer, and not on the plate like our MC and TSP design (Section 6.3). We modified the software described in Section 6.5 to be compatible with these electronics: the Digitizer Server was eliminated, and instead the RuO_2 reading was written to a global variable through a modification of the standard temperature monitoring software; the Heater Server was re-written to send the TCS RS232 commands in place of the Keithley's GPIB protocol. We were able to run, with no tuning and only a proportional action (no integral or derivative) for a total of about nine hours. The temperatures of the two Still thermometers during this trial are plotted in Fig. 6.6.

Over the course of the run, the cryostat was gradually warming up and two DCS strings were deployed, one at a time, in a concurrent test. It can be immediately seen from the two prominent spikes that the P loop was unable to fully compensate the heat load of the strings, though we cannot say to what extent this is a result of the feedback mechanism's non-existent tuning, or simply that of an excessively cold setpoint combined with low cooling power—the controller is able to supply varying amounts of heat, but it relies on the DU to make downward corrections passively. Under normal CUORE operations, 12 strings are routinely deployed with no prominent spikes. The amplitudes of the spikes from strings lowered with the P loop *are* lower than that of the first spike (circa 18:30 in Fig. 6.6), during which the temperature was not stabilized. We can also observe that temperature to which the system




Figure 6.6: Temperature of the Still v. time during the brief impromptu PID test on the Still level. Note that the loop is not tuned and only the proportional gain is enabled. A forced warmup of the cryostat due to a mixture line clog cut this test short and hampered further tuning. The temperature spikes are caused by the deployment of the calibration sources. Even the un-tuned PID decreases the magnitude of the spikes. Steady state error from the lone P term is observed in the inconsistently-stabilized baseline.

returns upon stabilizing is not consistent. This can possibly be explained by steady state error of the proportional correction, since integration is disabled (see Section 6.2).

While the test on the Still level was impromptu and incomplete, we were able to demonstrate that our system is able to manipulate the temperature inside the CUORE cryostat. If and when the LabView-based PID is operated in its intended configuration, proper tuning of the gain parameters (and possibly algorithm modifications) would have to be performed; however, knowing the settings from the currently-running PC+microcontroller hybrid system would likely speed up that process. Furthermore, we were able to successfully transfer our software to a new computer and a completely different set of hardware with different behavior and communications protocols in a matter of hours. This fact demonstrates the flexibility and adaptability of our software suite, and hence a good potential for its future use. There is an existing possibility that this PID controller will be adapted for use for bolometer R&D research currently ongoing at UC Berkeley in the Kolomensky group.

Chapter 7

A Radon-Free Environment for CUORE Installation

This chapter is derived from [81].

7.1 Introduction

To maintain the radiopurity of the CUORE detector, all assembly work was carried out in nitrogen-atmosphere overpressure glove boxes, making use of robots. Completed components were stored in nitrogen-flushed acrylic containers. All of the gloveboxes and storage are located inside the same class-1000 cleanroom facility.

We first used the CUORE assembly line for the construction of the CUORE-0 demonstrator experiment, which consists of a single CUORE-type tower [57, 73]. CUORE-0 achieved a factor of ~ 10 reduction in α background [76, 79] over the Cuoricino prototype, though they were operated in the same cryostat (Fig. 7.1a) [77, 78].

Though the CUORE detector's design and fabrication are identical to a scaled-up CUORE-0, its target background in the region of interest around the ^{130}Te Q-value (2.53 MeV) is lower by a factor of ~ 6 (Fig. 7.1b) [74, 80]. This is possible in a new custom-built cryostat, made with clean copper and carefully selected materials [122], as well as improved shielding (see Chapter 5). We need to ensure, however, that no significant recontamination occurs during the installation of the instrument in the cryostat. During this operation, the towers are removed from their nitrogen-flushed storage, and are exposed to the surrounding atmosphere until the cryostat's inner vacuum chamber (IVC) is evacuated.

Given that the towers are very fragile, ~ 1 m tall, and weigh a total of nearly 1 t, we opt to perform this operation in an air, rather than nitrogen, atmosphere. A temporary higher-specification cleanroom enclosure (CR6) of ~ 32 m³ volume with a radon-free air supply is erected around the CUORE cryostat, within the confines of our standard class-1000 cleanroom (CR5).




Figure 7.1: (a) Comparison of CUORE-0 and Cuoricino spectra, demonstrating background reduction. (b) Simulated CUORE energy spectrum from background budget study [74]. The peak at ~ 3.3 MeV in the experimental spectra comes from a ^{190}Pt crystal bulk contamination, which does not contribute to the background at $Q_{\beta\beta}$ [80], and is not included in the CUORE background budget.

Our goal, to avoid recontamination, is to keep the radon activity of the air below a 1 Bq/m^3 limit for the duration of the installation [123]. The clean air is supplied by the Radon Abatement System, which scrubs ambient air through carbon filters. Inside CR6, the air is circulated and cleaned of airborne dust by two ULPA (Ultra-Low Particulate Air) [124] filter cabinets. The radon level is constantly monitored by an electrostatic radon monitor (RM).

7.2 Description of the Setup

Radon Abatement System

The Radon Abatement System (RAS), manufactured by the Ateko company of the Czech Republic, functions by cleaning air from the surrounding environment. Upon intake, air is pressurized by a compressor to $\sim 9 \text{ atm}$, passed through an oil vapor separator and three microfilters, and fed into a dryer. There, the air is desiccated such that the dew point is below -70°C , and filtered again of any eventual liquid or dust. This, in turn, permits us to cool the air down to $\sim -55^\circ\text{C}$, prior to flushing it through two large activated carbon filters (in series) to trap the radon. The cooling is necessary, because activated carbon's efficiency for radon absorption decreases exponentially with temperature. These filters are a potential source of dust, particularly problematic for cleanroom use, so we pass the air through a set of coarse particle filters followed by HEPA (High Efficiency Particulate Air) filters after heating it back to room temperature (between 14 and 25°C). Fig. 7.2a shows a diagram of the RAS. It is deployed on an external platform/balcony of the underground hut

housing the CUORE experiment, adjacent to and one story above the cleanroom facility, allowing us to conveniently feed clean air through a port near the ceiling of the cleanroom.

At the output of the RAS, the air has a ^{222}Rn activity reduced to $< 5 \text{ mBq/m}^3$ (from $\sim 30 \text{ Bq/m}^3$ ambient), and is produced at a rate of $\sim 120 \text{ m}^3/\text{h}$.




Figure 7.2: **(a)** Diagram of the RAS, labeling key components, with ducting paths altered and simplified for clarity. The physical arrangement shown corresponds to the CUORE setup. The system is located adjacent to the cleanroom facility [81]. **(b)** Schematic of the MPIK RM, showing the gas inlet (1) of the $\sim 700 \text{ L}$ vessel (2) on the bottom, and the outlet at the top (4). The vessel sits at $\sim 10 \text{ kV}$ on insulated feet (3). The Si PIN diode (5) is grounded and mounted on an insulated cylindrical finger near the vessel outlet [81].

Radon Monitor

CUORE is vulnerable to contamination from the surrounding atmosphere over the course of the detector installation process, as well as the subsequent closing of thermal and radiation shields and other hardware work. This corresponds to a period of about four months, during which it is crucial for us to track the level of radon in the cleanroom, in case of failure or degradation of containment and/or the RAS. Given the generally low level of radon, the detector must be very sensitive. Additionally, because sudden failures are possible, and we need to be able to respond to those quickly, it must be fast. This fact prevents us from using conventional radon meters, such as the DurrIDGE RAD7, which can reach a sensitivity below 1 Bq/m^3 only by integrating over a long period of time.

Thus, we make use of a highly sensitive electrostatic radon monitor (RM) belonging to the Max Planck Institute for Nuclear Physics (MPIK) of Heidelberg, Germany. This is located at the facility of the GERDA experiment, also in Hall A of LNGS, about 80 m away

from CUORE. The RM is comprised of a 700 L vessel, through which the monitored gas is flushed. A silicon PIN diode is positioned at the outlet to maximize efficiency, at a 10 kV potential offset from the vessel (Fig. 7.2b). The RM has a static sensitivity of 0.5 mBq/m^3 , which improves to $50 \text{ } \mu\text{Bq/m}^3$ with an air flow rate of 7 L/min .




Figure 7.3: **(a)** The decay of ^{222}Rn , from the ^{238}U decay chain, which is the primary source of the radon contamination we are mitigating. Solid arrows indicate dominant decay channels, while dashed arrows are secondary processes. The polonium isotopes we detect directly are highlighted in bold [81]. **(b)** RM energy spectrum: ^{218}Po (red) and ^{214}Po (black) α peaks at 6.1 MeV and 7.8 MeV, respectively. Smaller peaks from ^{210}Po at 5.4 MeV and ^{212}Po at 9 MeV are also visible [81].

When a ^{222}Rn nucleus α -decays inside the RM, the ^{218}Po daughter is positively ionized, and drifts along the voltage gradient towards the grounded diode. If this nucleus does not decay or recombine to become electrically neutral on the way, it adheres to the diode surface and decays there to ^{214}Pb ($T_{1/2}^{218\text{Po}} \simeq 3 \text{ min.}$). The resulting α particle is detected with a $\sim 50\%$ geometric efficiency. Moving along the uranium chain (Fig. 7.3a), the ^{214}Pb and its ^{214}Bi daughter undergo relatively fast β^- decay ($T_{1/2}^{214\text{Pb}} \simeq 27 \text{ min.}$, $T_{1/2}^{214\text{Bi}} \simeq 20 \text{ min.}$) to produce another polonium isotope— ^{214}Po . This α decay, to ^{210}Pb ($T_{1/2}^{214\text{Po}} \simeq 164 \text{ } \mu\text{s}$), can also be detected by the diode. In fact, the probability of observing the ^{214}Po α is greater than that from ^{218}Po : in the event that ^{218}Po decays in the gas, all of the daughters mentioned above tend to keep moving in the direction of the diode, increasing the chance that the ^{214}Po is resting on the diode surface at the moment of decay.

Thus, the experimental spectrum, shown in Fig. 7.3b, is characterized by the ^{218}Po α peak at 6.1 MeV and a somewhat taller ^{214}Po α peak at 7.8 MeV. The value of the ratio between the observed peak amplitudes depends on the probability of recombination, and ultimately

on the properties of the monitored gas. These data are consistent with a total efficiency of $\sim 30\text{--}40\%$ for ^{218}Po and ^{214}Po . As we can see in Fig. 7.3a, the ^{238}U chain also contains ^{210}Po . Producing a 5.4 MeV α , this isotope is the source of the small bump in the tail of the ^{218}Po peak. Its amplitude is much smaller due to the longer half-lives involved, particularly for ^{210}Pb (about 22 y). The peak includes residual ^{210}Pb from all past measurements made by the meter. The contribution from the ^{232}Th chain is much smaller. The 6.9 MeV α from ^{216}Po is not visible in the spectrum beneath the ^{214}Po tail, while ^{212}Po 9 MeV α is responsible for the residual peak near 415 a.u. in Fig. 7.3b.

Cleanroom

Our cleanroom facility is located on the second (middle) floor of the CUORE experiment's underground hut. As originally constructed, it comprises five separate rooms (Fig. 7.4a: CR1 for entry, CR2 and CR3 for detector assembly, CR4 for storage of completed towers, and CR5 housing the cryostat).

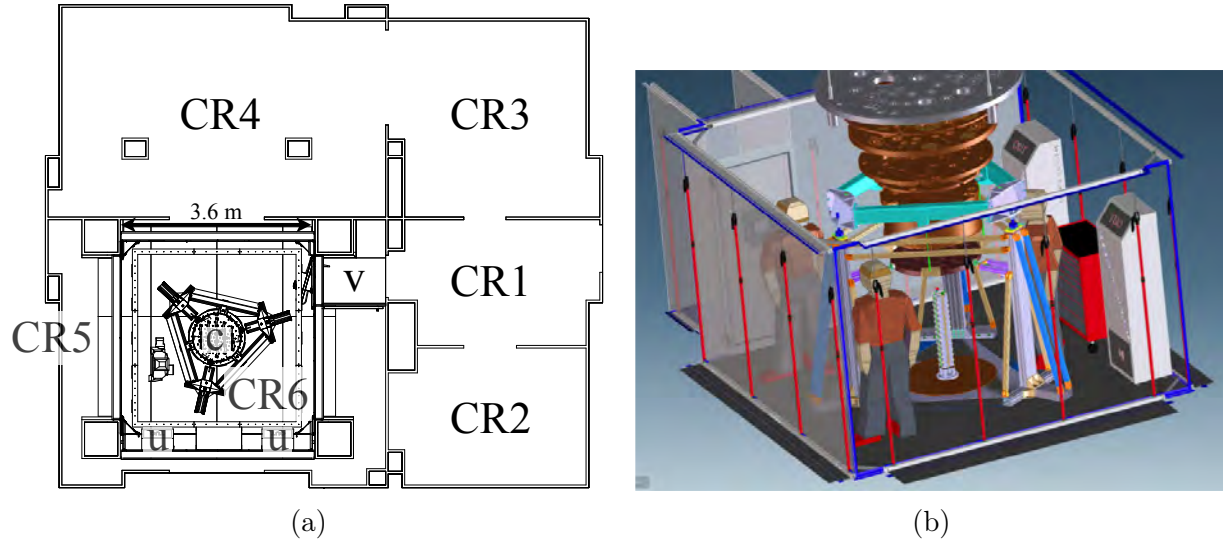


Figure 7.4: **(a)** A line drawing of the CUORE cleanroom facility showing the separate spaces: entry area (CR1), detector assembly (CR2, CR3), tower storage (CR4), and the cryostat (c) room (CR5). The temporary cleanroom (CR6) for detector installation is found within CR5, connecting to CR1 via the vestibule (v). Air from the RAS feeds into CR6 from a port in the ceiling adjacent to CR4, and is recirculated through two ULPA filter cabinets (u) [81]. **(b)** Rendering of CR6, showing the soft walls, vestibule, seals in blue, two ULPA filter cabinets, as well as the cryostat and detector installation hardware inside [81].

The assembly and storage areas are certified as class 1000, with a handler supplying ambient air (containing $\sim 30\text{ Bq/m}^3$ of radon) through HEPA filters. This is sufficiently clean for our purposes, because the detector assembly gloveboxes and tower storage containers are

all nitrogen-flushed. CR5, housing the cryostat, is equivalent to the others, but additionally features a large trap door in the center for raising and lowering the thermal and radiation shields, and a double door to the hall for forklift access. This floor also saw a significant amount of heavy equipment operations over the course of ~ 4 years of cryostat construction. Thus, we deploy a temporary softwall cleanroom (CR6) with a higher cleanliness level to serve as the radon-reduced work environment for detector installation.

Placed within the confines of CR5, CR6 consists of a sealed structure of thick plastic sheeting on an aluminum frame (Fig. 7.4b). The dimensions of the enclosure are $\sim 3.6 \times 3.6 \text{ m}^2$, with a ceiling height of $\sim 2.5 \text{ m}$ (equal to that of CR5) and an effective volume of $\sim 32 \text{ m}^3$. To lower dust levels, we laid a new stainless steel floor backed with plastic sheeting on top of the original tile floor, extending beyond the perimeter of CR6 itself. The Plexiglas and aluminum door features a floor sweep and an interlocking hand-through port. This door is aligned with the permanent door between CR5 and CR1. We use the same softwall construction as CR6 to set up an enclosed vestibule between the two doors, creating a “radon lock” in which assembly workers wait for air to recirculate before entering the clean area (Fig. 7.4). We find that there is no change in the radon levels in CR6 when personnel remain in the vestibule for one hour or more before entering the cleanroom, though the increase is acceptably small for a wait time of about 10 minutes. A second point of access is a heavy duty zipper in the plastic sheeting, aligned with the large double doors leading to the adjacent storage cleanroom (CR4). This can be used for the transport of bulkier equipment. All necessary electrical cables, hoses, and nitrogen lines are brought in through sealed feedthroughs mounted in metal panels attached to the plastic walls. We take advantage of the transparency of the enclosure, mounting the LCD displays of various monitoring devices outside of CR6, on the solid walls of CR5.

The air supply for CR6 is provided by the RAS described previously, in Section 7.2, through a port near the ceiling. Two Enviroco IsoClean HEPA filter cabinets [125] retrofitted with cleaner Teflon-based ULPA [124] filters perform the circulation and filtration of air within CR6. Besides being even more effective than regular HEPA filters, the ULPA units are radiologically superior—the HEPA filter medium is made up of randomly arranged fibers spun from borosilicate glass, which is a potential radon contributor. We disable CR5’s standard air handler, and seal off the HEPA filter outlets in the ceiling with metal panels to prevent the entry of radon-containing air. One of the two cabinets is fitted with a heat exchanger and connected to an external chiller, helping to keep ambient temperatures lower during work.

We make an effort to render CR6 nearly hermetic to CR5, using cleanroom-standard foam and tape on all edges, as well as 3M VHB (Very High Bond) double-sided tape on softwall seams. Combined with the disabling of the CR5 air handlers, this sealing ensures that the RAS air flow is sufficient to over-pressurize CR6 with respect to its surroundings. We can observe the softwalls “inflate,” and that the majority of air current out of CR6 passes through the entry vestibule, allowing it to function as designed. Design and installation of CR6 was led by Joseph Wallig of the LBNL Engineering Division.

Beyond the technical capabilities of the RAS and CR6, we can improve cleanliness by

adopting certain work practices. Two particularly important ones during CUORE installation are a daily washing of the interior surfaces of CR6 as well as tools, and the usage of a nitrogen-flushed “radon bag” (RB). This is a cylindrical plastic barrier hung around the bottom of the cryostat, covering those towers that are already mounted. It is continuously flushed with nitrogen, and is deployed at all times during which installation is not being performed. While playing a role in protecting the detector from recontamination under normal circumstances, the RB is especially important in the event of a failure of the RAS.

Once detector installation is complete and the innermost (10 mK) cryostat vessel is closed, we must open the hatch in the floor of CR5 to allow the raising of the remaining shields. This is impossible to do with CR6 in place. To minimize exposure during this phase, instead of returning CR5 to its standard configuration, we make use of an “intermediate” cleanroom setup. We keep the standard HEPA air supply disabled and blocked off, continuing to rely on the RAS and one of the two ULPA cabinets. To prevent radon-contaminated air entering from the hall and adjacent cleanrooms, we seal the doors and other gaps with cleanroom-rated foam and tape. The vestibule remains in place, permitting personnel to enter and exit without flooding the room with external air.

The conditions inside CR6 are actively monitored. We use a KNF diaphragm pump to send cleanroom air to the RM described in Section 7.2. Besides the radon level, we also keep track of the airborne particle count, humidity, temperature, O_2 and CO_2 levels. These sensors are equipped with alarms, and at least one shifter is watching the readouts at all times.

The softwall modular design of CR6 allows us to deploy or dismount it in a period of less than one week. This means that we can create a radon-reduced environment around CUORE on short notice if this becomes necessary in the future.

7.3 System Performance

The achievable sensitivity of the RM measurement depends on the integration time. During the installation of CUORE, we operated the RM with continuous flushing at a rate of 7 L/min, which gave us a response time of ~ 10 min for ^{218}Po . This delay primarily stems not from integration, but from the time needed to pump the air from CR6 and fill the vessel of the RM, as well as the recirculation time for CR6 itself.

In Fig. 7.5, we show the activities of ^{218}Po and ^{214}Po as a function of time, measured immediately after the completion of CR6 construction. The first 11 hours represent a readout of desiccated ambient air from the hall, with an average level of ~ 30 Bq/m³. The rapid decrease of ^{218}Po activity down to ~ 100 mBq/m³ starting around 11:00 corresponds to switching the measurement to CR6 air. The subsequent increase in count rate around 17:00 occurred during a test where the CR6 door was left open for several minutes. Something we notice here is the delay of observed ^{214}Po levels with respect to ^{218}Po , which occurs due to the 20–30 minute half-lives of the intermediate ^{214}Pb and ^{214}Bi nuclei (Fig. 7.3a).




Figure 7.5: ^{218}Po (red) and ^{214}Po (black) count rate v. time, immediately after the completion of CR6 construction. The first 11 hours represent radon-containing air, while the subsequent measurements are of CR6 air, showing time response to opening and closing CR6 [81].

Over the course of the tower installation period, we are able to maintain the radon in CR6 well below the target of 1 Bq/m^3 at all times (Fig. 7.6). When the team is working inside, the activity is stable at the $\sim 100 \text{ mBq/m}^3$ level. If no activities are being performed, the activity decreases to $\sim 10 \text{ mBq/m}^3$, not counting the nitrogen-flushing of the RB immediately surrounding the detector. When personnel enter the cleanroom under normal working conditions, opening the door does not cause increases in count rate as large as the one seen in Fig. 7.5 thanks to the use of the vestibule. Reducing the radon level in the empty CR6 from the ambient $\sim 30 \text{ Bq/m}^3$ to the minimum $\sim 10 \text{ mBq/m}^3$ takes approximately 7 hours. We find that the airborne particle count in CR6 remains below $100 \text{ particles/ft}^3$, which would correspond to an improvement over the standard CUORE cleanroom by one class.

In the “intermediate” configuration of CR5—with radon-free air from the RAS, a vestibule, and sealed gaps, but without the CR6 enclosure—we are able to reach a baseline radon level of $\sim 50 \text{ mBq/m}^3$ when not performing work (Fig. 7.6), not counting the nitrogen flushing of the 10 mK cryostat vessel. While this is five times higher than with CR6 deployed, it is well within our target levels. A quantitative summary of the RAS and cleanroom performance is given in Table 7.1.




Figure 7.6: ^{218}Po (red) and ^{214}Po (black) count rate v. time over the course of the entire ~ 4 month installation and cryostat closing procedure, with different operation phases labeled. With the “intermediate” configuration, 50 mBq/m^3 could be stably achieved. The large excursions above that level occurring in October and November correspond to the frequent openings of the $\sim 2 \times 2 \text{ m}^2$ hatch in the floor of CR5 [81].

Table 7.1: CUORE RAS and cleanroom performance summary. Particle count, relative humidity, and ambient temperature were actively monitored only in CR6. CR5 with the floor hatch closed is rated class 1000. Relative humidity in CR6 while unoccupied is stable at 0.5%, and can reach as high as 25% during active manual work. Ambient temperature remains at $\sim 19^\circ\text{C}$ most of the time [81].

	CR6	Intermediate, sealed	Intermediate, open floor hatch
Rn: ^{214}Po and ^{218}Po (mBq/m ³)	~ 10	50–100	1000–2000
Particle count (part./ft ³)	0–100	<i>rated class 1000</i>	<i>not measured</i>
Relative humidity (%)	0.5–25	<i>not measured</i>	<i>not measured</i>
Room temperature (°C)	18–23	<i>not measured</i>	<i>not measured</i>

Chapter 8

A Room-Temperature Calibration System for CUORE

8.1 Introduction

The CUORE detector, an array of bolometers, has the potential for excellent energy resolution. In the latest dataset released as of this writing, we have achieved an average resolution of 7.4 ± 0.7 keV peak full width at half maximum (FWHM) at the ^{130}Te Q-value of 2527.5 keV [33], and believe the ~ 5 keV FWHM performance observed in CUORE-0 [76, 79] is achievable with further hardware and analysis optimization. For a neutrinoless double-beta ($0\nu\beta\beta$) decay search, good energy resolution is beneficial because the correctly determined energy of the would-be signal at $Q_{\beta\beta}$ is the primary identifier of this rare nuclear process. Furthermore, a $0\nu\beta\beta$ peak would have to be separated out from the two-neutrino ($2\nu\beta\beta$) spectrum, and possibly other backgrounds. The latter is especially crucial for a device like CUORE, which only measures particles’ energies, without actively discriminating between β s and α s—our primary background in the region of interest [33, 74, 80].

For this precision to be meaningful however, our spectrometry must also be *accurate*, and calibration is essential. During standard operation, CUORE data acquisition consists of “datasets,” broken up into day-long “runs.” Besides including about one month of uninterrupted “background” or “physics” measurements, each dataset begins and ends with a ^{232}Th calibration run lasting several days. The CUORE Detector Calibration System (DCS) is a piece of very advanced low-temperature technology. It is capable of deploying 12 ^{232}Th source strings down through thin tubes to positions *within* the ~ 1 m³ array of bolometric detectors that comprise CUORE, and do so without disturbing the cryostat vacuum or the sub-Kelvin operating temperatures of the detector [85]. CUORE-0, Cuoricino, and other cryogenic detectors of the past and present do not feature equivalent hardware, typically introducing room-temperature calibration sources inside the external lead shield and outside the cryostat itself. The CUORE collaboration, primarily the University of Wisconsin/Yale University group, undertook the difficult task of designing and implementing the DCS be-

cause our detector is much larger than predecessors, being about as wide as it is tall. Getting the uniform illumination of all 988 channels by the calibration sources is not possible unless they are deployed throughout the bolometer array. Because CUORE's behavior is not necessarily reproducible if it is partially warmed up and then re-cooled, the lowering of the source strings must happen without heating the detector.

The features necessary to achieve this, however, also leave the DCS difficult to repair without a lengthy interruption to data-taking in the event of a technical failure. Namely, if the DCS were to fail to deploy a sufficient number of its strings for a final calibration, due to icing in its tubes or another reason, the experiment would risk losing the previous dataset. While, thankfully, this scenario has not transpired and all CUORE datasets taken thus-far were successfully calibrated with the cryogenic DCS, the collaboration has decided to design and deploy an External Detector Calibration System (External DCS) as a back-up device. A more advanced adaptation of the CUORE-0 calibration approach (string-like sources deployed inside the external lead shield, but outside the cryostat Outer Vacuum Chamber (OVC)), this system is unable to provide the uniform exposure of the full bolometer array. However, it is technically simple and robust, accessible from outside the cryostat, easy to fix, and can be operated in either automatic or manual regimes.

The External DCS also holds advantages for non-standard circumstances. Not needing to thermalize strings down to sub-Kelvin temperatures, it deploys and retracts them in seconds to minutes, versus hours or even days for the cryogenic system. As long as uniform illumination is not essential, this speed renders the External DCS a superior machine for quick tests and diagnostics. Furthermore, its sources are not stored on tiny mechanized spools inside vacuum boxes, but hang freely in the Faraday cage when not in use, letting us replace them easily. We performed ^{60}Co and ^{56}Co calibrations to better understand the detector physics of CUORE-0, and similar campaigns are planned for CUORE. Given that replacing the usual thorium sources with cobalt and then back again is very difficult with the cryogenic DCS, the External DCS is an excellent choice for these measurements.

Finally, we have observed that when cryogenic DCS has a malfunction during a normal calibration, it is not usually a total failure. Rather, only some strings fail to be lowered properly. The incomplete deployment is usually sufficient to calibrate most if not all channels, but quality can suffer. Being able to *augment* the cryogenic DCS with external sources in such a scenario could be very beneficial. As such, the two systems are complementary to each other.

Design and implementation of the External DCS was led by Joseph Wallig of the LBNL Engineering Division. Successful installation and technical testing of the complete system was carried out in January and February of 2018. In this chapter, we discuss its design and functionality. Given some of the intricacies and unique decisions involved in retrofitting such a system to an already-built experimental apparatus, we also cover the installation process.

8.2 External DCS Description

The CUORE External DCS is a larger-scale and mechanized evolution of the CUORE-0/Cuoricino calibration approach: string-like sources are deployed within the thick external lead and polyethylene shielding, but at room temperature and outside the cryostat outer vacuum chamber (OVC). A schematic of the system is given in Fig. 8.1.




Figure 8.1: A schematic diagram of the CUORE External DCS. While only three winch/pipe/source combinations are drawn for simplicity, eight are installed. The three sources are shown in the “home” (1), “intermediate” (2), and “calibration” (3) positions. Both of the alternative winch mounting schemes (strut and toggle bolt) and PVC pipe routes (through an opening in the 300-K flange and around its edge) are depicted.

Source Strings

Though we design the External DCS to be flexible in this regard, by default it is equipped with ^{232}Th calibration sources—the same isotope used by the main cryogenic DCS (see Section 4.5 and Fig. 4.10). The “sources,” as currently installed, are fabricated from standard

thoriated tungsten welding rods manufactured by Diamond Group Products. The thorium content as specified by the company, and assumed during initial planning, is 4% by mass. Unfortunately, when counted it was found to be lower by a factor of about four. The measured activities are summarized in Table 8.1.

Table 8.1: Summary of manufacturer-specified and measured thorium content of welding rods used for fabricating the calibration sources for the External DCS. Measurements were carried out on two *different* pieces of rod from the same order, at Lawrence Berkeley National Laboratory (LBNL) and Laboratori Nazionali del Gran Sasso (LNGS), using the Merlin and GePaolo Germanium detectors, respectively. “Early” ^{232}Th , via its daughter ^{228}Ra , is determined from ^{228}Ac . “Late” ^{228}Th numbers are from ^{212}Pb , ^{212}Bi , and ^{208}Tl . The thorium chain was found to be in equilibrium.

	Nominal (% mass)	LBNL, Merlin (% mass)	LNGS, GePaolo (% mass)
Th-232	4		
Ra-228 (“early Th”)		1.03 ± 0.01	1.09 ± 0.06
Th-228 (“late Th”)		1.14 ± 0.01	0.95 ± 0.05

The welding rods are cut up into $\sim 1''$ ($\sim 2.5\text{-cm}$) pieces, and strung together with a flexible mesh nylon sheath, visually similar to sausage links. There are eight strings in total. Each one is $\sim 1\text{ m}$ long, because they are deployed vertically around the $\sim 80\text{-cm}$ -tall CUORE detector, and we would like to allow for some overshoot on either side. Spacing of rod sections or “capsules” is determined from the measured activity of the tungsten, and the desired total activity as determined from Monte Carlo simulation studies performed by Oliviero Cremonesi and Giovanni Benato.

Due to the material being $\sim 4\times$ less active than anticipated, the spacing of capsules in the strings currently installed is effectively determined by the maximum achievable density that would still preserve necessary string flexibility. This corresponds to a total activity of $\sim 70\text{ kBq}$ from all eight strings, which is slightly “cold” compared to the suggested 102 kBq from the simulation. To that end, we have also fabricated eight add-on strings with a total activity of $\sim 30\text{ kBq}$, which can be easily stitched on to the existing ones and deployed together.

Besides the total activity, the simulation shows us that the efficiency of the different areas of the detector differ significantly (Fig. 8.2). Along the z axis (vertical), efficiency is highest in the middle of the detector and falls off quadratically toward the ends. We can attempt to compensate for this behavior by fabricating source strings with non-uniform capsule density. The main strings currently installed are “close-packed,” and density cannot be adjusted. We can partially alter the distribution with the add-on strings, but not entirely—the ends of the complete sources would have to be $\sim 2.5\times$ hotter than the center.




Figure 8.2: Simulated CUORE detection efficiency to the vertical sources from the External DCS. The z -dependence of the efficiency can be compensated for by a non-even activity density of the sources. The $\sim 10\times$ lower efficiency of the inner layer of the detector (7 of 19 towers) compared to the outer layer (12 towers) is an unavoidable weakness of using external sources.

After a discussion, it has been decided that the add-on strings will not be added unless it is determined to be useful after testing. If the current uniform 70 kBq setup is found to be insufficient, there is a chance that we would fabricate entirely new, non-uniform sources from fresh rods with higher thorium content, to fully account for the z -dependence of the detection efficiency. The much weaker efficiency of the inside channels of CUORE is the main downside of the External DCS, and the main motivation for the development of the cryogenic DCS.

Deployment and Hardware

When deployed, the source strings are guided into their correct positions along vertical PVC pipes surrounding the cryostat. This arrangement is illustrated, in simplified form, in Fig. 8.1. Seven of these pipes penetrate the roof of the cryostat cleanroom (CR5) through holes in the plastic sheet covering the gap between the 300-K plate (cryostat roof) and the edges of the opening in the CR5 hard ceiling. Natural tension of the plastic keeps these holes leak tight. The tops of the pipes are positioned at locations around the perimeter of the 300-K plate, to which they are held by rigid clamps. A single, eighth pipe passes through a pre-existing opening in the 300-K flange near its edge. They penetrate ~ 10 cm above the cryostat roof, and are equipped with thermoformed funnels for guiding the string.

Surrounding the OVC, the PVC pipes' locations are designed to be as symmetric as possible around the vessel and preferably aligned with the corners of the octagonal external shield for better clearance. On the roof, however, the pipes outlets/funnels must have a clear, unobstructed line of sight to the ceiling of the Faraday Room. As such, each pipe has a unique thermoformed curved shape (Fig. 8.3a) within the ~ 30 cm of vertical space below the cryostat roof and above the top of the external shield.

The strings are deployed by means of computer-controlled winches—the Winch 5 and the Winch 10—produced by the Wahlberg company of Denmark [126, 127]. These devices are originally intended for theater/stage applications, and as such are optimal for smoothly and precisely lifting and lowering light objects, such as our sources, to and from pre-programmed position. The winches are anchored to the ceiling of the Faraday Room, and aligned with the pipe openings with a laser level. Three lead fishing weights on the end of each string help ensure successful entry into the mouth of the pipe. The strings are identical and interchangeable, since the unique positioning of each string is programmed on the winch level, with each winch deploying the correct amount of steel cable. The strings are clipped to the winch cables with quick-release carabiners that require no tools, and can easily be replaced and modified.

Wahlberg Winches—Control and Software

The Wahlberg Winch 5 and Winch 10 model wire winches are originally intended for the rapid and precise movement of props and other stage objects in a modern theater setting [126, 127]. With this application in mind, they have a low weight capacity¹, but are very accurate in their deployment speed, positioning, and timing, with remotely supplied instructions from the lighting desk or a programmed script. These same qualities make these devices an excellent choice for deploying our calibration sources, which are light weight and need to be placed repeatably and accurately with no physical operator contact.

The winches are controlled with DMX512 (**D**igital **M**ultiplex) digital communications standard, which is ubiquitous in the stage lighting and effects industry. This uses cylindrical 5-pin XLR5 connectors (similar to microphones, visible in Fig. 8.4b) and a serial protocol—not dissimilar to the RS-series of standards more common in scientific instrumentation. The protocol is used over a two-wire bus, and allows multiple devices to be controlled from one port as long as they are properly addressed [128, 129]. We make use of this in the External DCS, connecting our winches in series—each one hosts both input and output XLR5 ports. The DMX address of each one is set by a set of three screwdriver-operated dials, which are the row of yellow dots visible in Fig. 8.4b. Because seven DMX channels are used for the control of each winch [126, 127], we set the addresses in increments of seven: 001, 008, 015, etc. Only one of the eight devices is connected directly to the control computer, with a DMX-to-USB converter.

¹5 kg and 10 kg for the Winch 5 [126] and Winch 10 [127] models, respectively.

The fourth yellow dial on each winch housing sets its function mode, for which there are seven options, of which two are for “positioning.” In this context, this term refers to the behavior of the winch moving its cable to some position specified by a command, at a preset speed, and remaining there until the next command is received. The location is specified in bits, as a number between 0 (home) and 255 (fully down). With a 10-m cable, this corresponds to ~ 4 -cm or ~ 1.5 ” increments. This number stems from the motor controller, which has 16-bit coarse and 16-bit fine channels, corresponding to DMX channels 1 and 2. The two positioning modes are Mode 1, with automatic top reset, and Mode 2, with manual top reset. We program our “top” or “home” position, in which our sources are fully retracted, hanging above the MSP, with no physical contact with the cryostat. Thus, we choose Mode 1, providing us automated reliable extraction of sources. Now, “positioning,” in this context, refers to the behavior of automatically moving the winch to a position

Any DMX512-compatible control interface can be used to operate the Wahlberg winches. We utilize the ChamSys MagicQ [130] commercial “virtual lighting desk” software, which is available to download without cost and is well-suited for operating winches out of the box. Like the winches themselves, this program is intended for stage effects automation and incorporates many features that we do not need. However, it auto-detects all of our devices when they are plugged in, and allows us to move them both manually and automatically. Specifying the position in bits in the software, we can tune our desired source locations and then save them. Once that is done, we can send our sources to any of these pre-recorded positions by clicking the appropriate button in the graphical user interface.

8.3 Installation Work

All parts for the External DCS were procured by and delivered from LBNL in the United States, and were then subsequently assembled in situ in the CUORE hut. Because simplicity and speed of use and installation—the opposite of the advanced cryogenic DCS—were a priority for this backup system, all components are off-the-shelf commercial products adapted to this application.

Source strings were prepared in January 2018 at LNGS, having used activity measurements from LBNL and LNGS, Monte Carlo simulation data, and practical spacing limitations to determine the density of tungsten capsules and their spacing. Vertically compressing the nylon mesh casing expands it, at which point we inserted the capsules. Capsules can be positioned within the casing in the same way. When the string is hanging, it is sufficiently tensioned by its own weight that the capsules remain in place and do not slide down. The finished string is melted shut with a flame on the ends, and clipped to the winches using fishing snap swivels and keychain carabiners, which allows free movement and quick release, and was an installation-time decision.

Given the sensitivity of many components of the CUORE infrastructure, we assembled and tested a preliminary demonstrator winch with a source string on February 1 2018, on the balcony outside of the CUORE hut underground, before beginning actual construction.



Figure 8.3: **(a)** View of the OVC with PVC pipes installed. The curving path can be seen at the top. **(b)** Close up of one source string deployed in the “calibration” position in its pipe.

The PVC pipes were installed first, on February 2. Each of the eight pipes consists of three parts: the main span of the pipe, containing the curve, a funnel on top near the cryostat roof, and a capped extender piece on the bottom to give each pipe the correct length. The connections between the three sections and the bottom cap are made with sleeves, without glue or threading, and can be easily disengaged if necessary.

The curve in the main span is unique to each pipe, determined by the desired location of the source near the detector at the bottom, and the position of the outlet on the cryostat roof. While the placement of the sources is known in advance, we had only an approximate idea of the funnel positions at the design stage. The cryostat infrastructure, cryogenic DCS hardware, and detector electronics crowd the interior of the Faraday room. The winches need to be attached to the ceiling, exactly above the tubes with an unobstructed line of sight between them. To achieve this in the safest and most optimal fashion, we decided to determine the locations over the course of the installation.

As such, we did not have the ability to pre-fabricate curved tubes. We purchased straight PVC pipe stock, and had to develop a technique to define and create the correct shape for each one. Besides correctly locating the pipe ends, the curved section had to be sufficiently steep for the sources to be able to drop through, and hug the OVC closely enough to not interfere with the top of the external lead and polyethylene shield. So, we used a length of flexible copper refrigeration tubing to make a template of the each PVC pipe’s path by hand, after which the PVC was thermoformed with a heat gun and water inside CR5, but away from the cryostat.

The tops of the pipes, right below the funnel, are held to the cryostat roof with durable

plastic clamps. The main lengths of the pipes are held tight to the OVC with a loop of copper wire and small pieces of cleanroom-compatible tape.

A stabilized laser level, pointing up, was placed on the mouth of each pipe in turn to precisely mark the location of each winch. The eight winches are installed with a single bolt. Three of the winches—right-most in Fig. 8.4a, photographed in Fig. 8.4b)—are attached by way of struts. Specifically, one is mounted to an existing strut used to support pulse tube lines (see Chapter 5). The other two are on a dedicated strut spanning the length between two of the pulse tube ones. On the other sides of the cryostat, an equivalent arrangement was impossible, so we drilled a total of five holes (one for each winch) in the ceiling of the Faraday Room. We anchored the winches in these holes with toggle bolts, making sure these did not create electrical shorts to the ceiling of the CUORE hut. Prior to drilling, calculations were done to ensure that neither the electromagnetic shielding qualities, nor the structural integrity of the room were compromised.




Figure 8.4: **(a)** 3-D rendering of winch locations above the cryostat and MSP. The three right-most winches are mounted on struts. The remaining five winches are attached to the Faraday Room ceiling with toggle bolts. **(b)** Photograph of three winches in situ. The red light from the laser level demonstrating alignment can be seen on the nearest winch, right where the cable exits the housing. Surrounding pulse tube lines and the CUORE electronics mounted over the cryostat are also visible.

Once installation was complete, we calibrated the winches to deploy the strings to correct positions. The PVC pipes are clear, allowing us to see the position of the string inside while the external shield was lowered. Each winch is programmed with 3 positions (depicted in Fig. 8.1):

1. “home”—hanging above the main support plate (MSP), completely outside of the pipe, not touching anything.
2. “intermediate”—at the top of the funnel, just barely penetrating the tube, useful for checking alignment.

3. “calibration”—fully deployed inside the pipe, almost to the bottom, as shown in Fig. 8.3b.

The MagicQ [130] software was downloaded and installed on the CUORE underground laptop computer and tested there.

8.4 System Use and Testing

We tested the completed External DCS in situ multiple times, ensuring that the strings enter the pipes and deploy to their correct positions in a robust and repeatable fashion. Many times we ran through both single manipulations and complete “mock calibrations,” consisting of:

1. Plugging in and powering the winches without stepping on the MSP and disturbing the cryostat; starting up the software.
2. Lowering all the strings from Position 1 (home) to Position 2 (intermediate).
3. Lowering all the strings from Position 2 (intermediate) to Position 3 (calibration).
4. Powering down/unplugging the winches and quitting the software with the strings deployed.
5. Starting the software again, and setting the winch positions to Position 1 (home) *before* powering them again.
6. Plugging in and powering the winches, causing them to immediately extract the source strings automatically to Position 1 (home).

We rehearsed this routine both slowly, observing each string carefully as it went down, and rapidly, with a blind operator, preventing user interference and testing the reliability of the system. We found that being careful and not rushing, a team consisting of one operator and one observer could bring the External DCS online, deploy the sources, and then shut down the system in under ten minutes.

We observed that the winches have the impressive ability to return to the same position within 1 mm if they remain powered on. If powered down for the duration of calibration and when not in use—the more realistic scenario—precision is within 2 – 2.5 cm, which we consider to be acceptable given that the source strings have 10 cm extra length on either end. Having a string miss its funnel and require manual intervention was found to be unlikely, but does happen occasionally. We noticed that lowering the string to Position 2 (intermediate) at the top of the funnel before commencing full deployment down the pipe seriously improves reliability, and also makes manual re-alignment much easier in the improbable event of failure. As such, we consider this to be a part of standard operating procedure for deployment.

The CUORE external shield was raised and the nitrogen “bag” installed approximately a month after External DCS installation, and the as-built system was found to not interfere

with this hardware. Mock deployments were run again after the closure, and were successful. The alignment of the tubes and winches remained intact.

Tests of the system in its actual calibration role will be conducted once CUORE is cold and stably operating, in the late Spring or Summer of 2018.

Chapter 9

CUORE Analysis: Spectral Line Shape Modeling

Не отдавайте в физики
детей...

В.С. Высоцкий

9.1 Introduction

The CUORE experiment is sensitive to rare nuclear processes, most significantly the neutrinoless double-beta ($0\nu\beta\beta$) decay of ^{130}Te [33], but also two-neutrino double-beta ($2\nu\beta\beta$) decay of ^{130}Te [80], beta-plus-electron-capture (β^+EC) decay of ^{120}Te [131], and others. Our detector is essentially a highly precise spectrometer measuring the energy of absorbed particles of all types. The “generic” rare decay analysis approach with CUORE is based on fitting the region of interest (ROI) in the energy spectrum with some model PDF (probability density function) that allows for either the absence or the existence of the process. For $0\nu\beta\beta$ decay, for instance, the process manifests as a monoenergetic peak at the $\beta\beta$ Q-value. The fit PDF, thus, needs to contain a “spectral line” component at that energy to handle any events that might occur there. Additionally, the PDF includes the background found in the ROI, including any γ lines. For example, we find the ^{60}Co two-photon ($\gamma\gamma$) line at ~ 2506 keV—very near to the ^{130}Te Q-value of ~ 2528 keV [33, 76].

For this type of analysis to be successful, we need create a model PDF for the spectral line shape in the CUORE detector. The peaks in this real-world, non-ideal device have a non-Gaussian shape. To characterize it, we use the RooFit Toolkit for Data Modeling [132, 133] to perform unbinned extended maximum likelihood (UEML) fits on calibration data. We fit the ~ 2615 -keV ^{208}Tl γ line, which is the most prominent in the ^{232}Th calibration spectrum, also significant in background data, and is similar in energy to $Q_{\beta\beta}$ of ^{130}Te . Our fit model covers both the photopeak and surrounding features in a 190-keV window around

it, and comprises both theoretical/analytical and empirical/effective components. We carry out 19 independent simultaneous fits—one for each of CUORE’s towers (see Chapter 4)—and produce a unique photopeak PDF for every channel-dataset pair. In this chapter, we explore the elements of CUORE spectrum in the vicinity of 2615 keV, our fitting technique, and the final results.

9.2 A First Look at the 2615-keV Line in CUORE Calibration Data

Thallium-208 is a radioactive isotope present in the decay chain of ^{232}Th —CUORE’s calibration source. Its ~ 2615 -keV peak is the most energetic and prominent γ line in the calibration spectrum (see Fig. 4.12), and also the nearest in energy to the ^{130}Te $\beta\beta$ Q-value. The event rate of ~ 2615 -keV photons is high, and the background around the peak is relatively low: all other thorium-series γ s are lower in energy, so we are above the Compton continuum as well as the ^{130}Te $2\nu\beta\beta$ distribution.




Figure 9.1: The CUORE calibration spectrum between 2530 and 2720 keV—the line shape fit domain. Data from the entire detector from datasets 1 and 2 with all analysis cuts applied is plotted (1811 channel-dataset pairs). Labeled are the features we include in the construction of our fit PDF.

Nonetheless, when we examine the spectral region between 2530 and 2720 keV, we find a number of features besides the photopeak itself, as shown in Fig. 9.1. The flat baseline

background is nearly an order of magnitude higher below the ^{208}Tl line than above it. This is a result of multi-Compton events: we are above the Compton edge, which is the maximum energy for a single Compton scattering event, and occurs when the scattering angle is equal to 180° [134]. However, between the edge and the photopeak we find the “Compton valley,” a uniform distribution of events stemming from multiple Compton scatters followed by photoabsorption [135].

Additionally, there are two small peaks—one near 2585 keV, and the other close to 2690 keV. The former is caused by the escape of one of six different ~ 30 -keV Te characteristic X-rays when a 2615-keV γ is absorbed [136]. The latter is a fairly rare coincidence event. When ^{208}Tl decays, 2615-keV photons are always emitted simultaneously with 583-keV ones. These two are rarely detected by the same bolometer crystal, but occasionally they do coincide. Our tiny peak results from the even more rare scenario when such a two-photon event further coincides with a single 511-keV photon escape. Since the energy of the 2615 keV γ -ray is greater than 1022 keV—twice the mass of an electron or positron (511 keV)—this γ can undergo pair production (create a e^+/e^- pair). When we measure such an event, the ^{208}Tl γ has an energy of $2615 - 1022 = 1593$ keV. The e^+ and e^- , in turn, annihilate to create two 511-keV photons. If both of these are absorbed in the same crystal as the original ^{208}Tl γ , we register a full 2615-keV event. However, one or both of the 511-keV γ s has a chance of exiting the detector, leading to single-escape and double-escape spectral lines. Thus, the true energy of our small two-photon + single-escape line is $2615 - 1022 + 583 + 511 = 2687$ keV.

Finally, the main photopeak itself appears to be non-Gaussian, having one or two shoulders. This is a known behavior in CUORE-type bolometers, and was previously seen in the CUORE-0 [76, 79] and Cuoricino [58, 137] experiments. Possible explanations for this phenomenon are the imperfect stabilization of pulse baselines, dependence on the position of the absorption in the crystal, or event topology dependence, though as of this writing we do not yet have a complete understanding.

In our physics analyses, we are interested in modeling the shape of just the photopeak. However, for our PDF to describe it properly, we must fit it in its entirety. This means that we must also fit the entire region shown in Fig. 9.1 and include all spectral features found therein.

9.3 Fit Function

The fit function we write to model our detector response comprises a sum of five components, one for each of the physical spectral contributions we observe between 2530 and 2720 keV. These are the 2615-keV photopeak, the ~ 30 -keV X-ray escape peak, the 2687-keV coincidence peak, the multi-Compton plateau, and the nearly-flat remaining background consisting mostly of random coincidences and some degraded α events. The complete function has the following form:

$$\begin{aligned}
f_{\text{total}}(E) = & N_{\text{sig}} \cdot (\text{PDF}_{\text{photopeak}} \\
& + R_{\text{Xray}} \cdot \text{PDF}_{\text{Xray}} \\
& + R_{\text{coin}} \cdot \text{PDF}_{\text{coincidence}}) \\
& + N_{\text{comp}} \cdot \text{PDF}_{\text{Compton}} \\
& + N_{\text{bkg}} \cdot \text{PDF}_{\text{background}} .
\end{aligned} \tag{9.1}$$

The X-ray escape and coincidence peaks are both composed of ^{208}Tl decay events, like the main photopeak, only coinciding with the detection or absence of additional photons of lower energy. Given that all three lines represent the same physical process—“the signal”—in our fit we account for them with a single total multiplicity N_{sig} . The amplitudes of the two smaller peaks are defined in terms of the ratios R_{Xray} and R_{coin} with respect to the main 2615-keV line. The quantities of multi-Compton, random coincidences, and degraded α events depend on independent processes, so the magnitudes N_{comp} and N_{bkg} are defined separately.

The functions we select to describe each of the five sub-PDFs are selected based on the initial observations described in Section 9.2 and previous experience with modeling the detector response of the similar CUORE-0 and Cuoricino bolometers [58, 137, 138].

The Photopeak

The 2615-keV γ -rays from the decay of ^{208}Tl are mono-energetic, and the theoretical ideal spectral peak would have the form of a δ function. Of course, no real world detector has zero resolution, and observed peaks have some shape and width. Assuming detector-derived imperfections in the energy measurement are random, or at least nearly so, the line shape would be well described by a Gaussian curve. As we have already mentioned, this is not true in CUORE-type bolometers. Fitting low-statistics peaks with normal distributions gives good agreement, but deviations become apparent with more counts. Specifically, these manifest as excess events in the tails of the peak, or “shoulders.” We saw the same behavior in precursor detectors, and CUORE-0 empirically described its line shape with a double-Gaussian function, with a small “subpeak” or “sub-Gaussian” filling out the lower-energy tail. The bulk of the photopeak is still described by a regular Gaussian, with amplitude and energy ratios defining the subpeak with respect to the main one [76, 79, 138]. Improved line shape models that included a third Gaussian were tested in CUORE-0, also on the left side of the photopeak. Some channels in CUORE-0 exhibit excess events on the right side of the peak, but the effect was small and not addressed in the line shape fit [138].

In CUORE, we observe meaningful excesses on both sides of the 2615-keV line. Drawing on CUORE-0 experience, we also select a multi-Gaussian fit approach. However, we choose a new arrangement with two subpeaks, one for each of the tails. The fit equation for the main photopeak PDF is given by Eq. 9.2:

$$\text{PDF}_{\text{photopeak}}(x) = \frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{(x-\mu)^2}{2\sigma^2}} + \frac{R_{\text{lo}}}{\sigma\sqrt{2\pi}}e^{-\frac{(x-\rho_{\text{lo}}\cdot\mu)^2}{2\sigma^2}} + \frac{R_{\text{hi}}}{\sigma\sqrt{2\pi}}e^{-\frac{(x-\rho_{\text{hi}}\cdot\mu)^2}{2\sigma^2}}. \quad (9.2)$$

The central Gaussian contains the majority of events. Its mean μ , near 2615 keV, and standard deviation σ of a few keV, are key fit parameters. Both sub-Gaussians are defined with respect to the main peak, with their centers given by energy ratios to μ , ρ_{lo} and ρ_{hi} . When fitting, these two parameters are constrained such that one is always above and the other below μ . Because energy resolution is a fundamental property of the detector, the sub-Gaussians' widths are set by the same σ . The quantities of events in the subpeaks are also set with ratios:

$$\begin{aligned} R_{\text{lo}} &= N_{\text{low subgauss}}/N_{\text{main gauss}} \\ R_{\text{hi}} &= N_{\text{high subgauss}}/N_{\text{main gauss}}. \end{aligned} \quad (9.3)$$

The total number of signal events is given by a seventh fit parameter, N_{sig} , which multiplies the photopeak PDF and those of the X-ray escape and the coincidence peaks, as shown in Eq. 9.1.

The triple-Gaussian line shape model works well for CUORE, but like the double-Gaussian of CUORE-0 it is an effective fit. Based on empirical observations and testing, it does not represent a theoretically understood detector behavior leading to the non-normal form. It is the author's suspicion that imperfect stabilization of all events passing the final analysis cuts leads to a minority of slightly "misplaced" events, but this remains to be explored as of this writing. The subpeaks are much more distinct in channels with good energy resolution, where they can be easily seen with the naked eye. In "fuzzier" channels, the sub-Gaussians are swallowed by the main line.

The X-ray Escape Peak

When a sufficiently energetic electron, ion, or photon is incident on an atom, it knocks out an electron from the inner shell. As an outer electron drops into the favorable lower shell, a photon corresponding to said electron's transition energy is emitted. Such radiation is termed "characteristic X-rays," because they are emitted at exact energies particular to the element.

2615-keV γ -rays incident on Te atoms in CUORE's TeO_2 crystal absorbers cause this type of fluorescence, with six significant X-rays between 27 and 32 keV given in Table 9.1 [136]. The majority of these photons are re-absorbed in the same bolometer, and we still register the full event energy. In some cases, however, they escape the crystal, resulting in a measured energy of ~ 2585 keV and creating the X-ray escape peak in our spectrum.

We know that the X-ray escape events happen with six different but similar energies, in predictable ratios. Thus, our fit function for this peak is a constrained sum of six Gaussians:

Table 9.1: The relative amplitudes and energies of six Te X-rays that create our X-ray escape peak. Data from the LBNL X-ray Data Booklet [136].

Relative intensity	Energy (keV)
25.1	27.202
46.2	27.472
4.26	30.944
8.21	30.995
2.37	31.704
0.147	31.812

$$\text{PDF}_{\text{Xray}}(x) = \frac{1}{\sqrt{2\pi}\sigma W} \sum_{i=1}^6 R_i e^{-\frac{(x-\mu+E_i)^2}{2\sigma^2}}, \quad (9.4)$$

where the six X-ray energies E_i are each subtracted from μ , the mean of the main Gaussian of the photopeak. The ratios R_i refer to the corresponding relative frequencies of these six energies, and weight W is their sum:

$$W = \sum_{i=1}^6 R_i. \quad (9.5)$$

Thus, in Eq. 9.4 we have a function that correctly describes the complicated six-photon line, but does so with the same number of degrees of freedom as a single normal distribution. In fact, the Gaussians' means are defined via that of the photopeak— μ —because the γ -ray's energy is physically the same value. The standard deviation σ is likewise common, because energy resolution is a fundamental property of a given bolometer channel. Hence, the only truly independent free parameter in the X-ray escape peak sub-PDF is the fraction of such events with respect to full-energy ones, R_{Xray} as shown in Eq. 9.1.

Table 9.2: The ratios and energies of six Te X-rays that create our X-ray escape peak, as reconstructed by fitting the detector response in the CUORE-0 Monte Carlo model in [138].

i	R_i	E_i (keV)
1	$3.20 \cdot 10^{-4}$	27.20
2	$6.08 \cdot 10^{-4}$	27.47
3	$2.61 \cdot 10^{-4}$	30.96
4	$5.90 \cdot 10^{-6}$	31.22
5	$6.01 \cdot 10^{-5}$	31.65
6	$5.02 \cdot 10^{-6}$	31.80

The values of the X-ray energies and their branching ratios are given in Table 9.2. An important detail is that these numbers are *not* the same as those quoted by Kortright and

Thompson in the LBNL X-ray Data Booklet and reproduced in Table 9.1 [136]. Rather, because we are attempting to best model the detector response, we take the X-ray values from fits done on the output of the CUORE-0 Monte Carlo simulation in [138]. This simulation uses the [136] numbers as true values, but our bolometers are not able to reconstruct them perfectly, especially the three more rare, higher energy emissions. Using the actual CUORE simulation would be an even more appropriate representation, but this data is not yet available at the time of this writing and the structurally-identical CUORE-0 demonstrator is considered an appropriate stand-in.

The (2615 + 583)-keV Coincidence Peak with a Single Escape

During the ^{208}Tl decay, 2615-keV and 583-keV photons are always emitted simultaneously [139]. These can travel in different directions from the source, and do not usually end up in the same bolometer crystal. Even more rarely, this two-photon absorption further coincides with a 511-keV single escape. This interesting two-photon + escape coincidence leads to a peak at ~ 2687 keV in the CUORE calibration spectrum. Like the main 2615-keV photopeak, this spectral line is monoenergetic. Furthermore, its amplitude is tiny, so it is not necessary to use the proper triple-Gaussian line shape for its description. Instead, we use a single normal PDF:

$$\text{PDF}_{\text{coincidence}}(x) = \frac{R_{\text{coin}}}{\sqrt{2\pi}\sigma} e^{-\frac{(x - \rho_{\text{coin}}\mu)^2}{2\sigma^2}}. \quad (9.6)$$

As with the X-ray escape peak and the two shoulders of the photopeak, we define the coincidence peak's mean as a ratio to that of the photopeak, μ . Due to very low statistics, though we do know the theoretically correct value of this ratio, we allow it to float slightly as the heavily constrained free parameter ρ_{coin} . The second free parameter is the fraction of coincidence events with respect to 2615-keV ones, R_{coin} . This peak is the third contribution to the total signal magnitude N_{sig} (Eq. 9.1).

The Multi-Compton Plateau

The Compton valley is a spectral feature below the photopeak that consists of multiple Compton scatter events. The multiple-scatters are distributed randomly between the Compton edge (the single-scatter maximum energy) and the photopeak. Above 2615 keV, they are naturally completely absent [134, 135]. Within our line shape fit domain, this continuum is locally flat, because at 2530 keV and higher we are avoiding the tail of the Compton edge. Thus, we use the complementary error function, $\text{Erfc}(x) = \frac{2}{\sqrt{x}} \int_x^\infty e^{-t^2} dt$, to effectively model the multi-Compton plateau as a “smeared step”:

$$\text{PDF}_{\text{Compton}}(x) = \frac{1}{2} \text{Erfc} \left(\frac{x - \mu}{\sqrt{2}\sigma} \right). \quad (9.7)$$

The $(x - \mu)$ in the function's argument indicates that the step cuts off at the photopeak's mean energy μ . While for physical event energy this is a sharp limit, in reality it is blurred by the detector resolution. As with the various Gaussian contributions, this is set by the common parameter σ intrinsic to each bolometer channel. The total multiplicity of Compton events is given by the independent free parameter N_{comp} in Eq. 9.1.

The Linear Background

The remaining background in this region of the CUORE spectrum is comprised primarily of random coincidences, as well as some degraded α events. It is essentially flat, and was described with a simple constant in the CUORE-0 and Cuoricino analyses [58, 76, 79, 137, 138]. In CUORE, its behavior is basically the same, though lower in overall magnitude. Empirically, we find that it tends to slope downwards very slightly in many of the channels. While a constant fit is adequate, the linear model given by Eq. 9.8 was selected due to superior performance.

$$\text{PDF}_{\text{background}}(x) = m_{\text{bkg}}x + 1 \quad (9.8)$$

The parameter m_{bkg} is the background's slope, which tends to take on very small negative values. The background rate is completely independent from that of the 2615-keV γ -rays, so the background multiplicity N_{bkg} is an independent fit parameter in the full PDF (Eq. 9.1).

9.4 Fitting Technique

To model the CUORE line shape, we perform a fit to the calibration data with the model described in Section 9.3. Despite all our efforts to achieve maximum channel-to-channel uniformity during detector construction through the use of a highly automated assembly line, the performance of individual bolometers varies significantly. The energy resolution of any pair of channels can differ by up to an order of magnitude, while many of them saw a significant improvement thanks to detector optimization done between datasets 1 and 2 (see Fig. 9.6). Hence, we need to have a unique set of photopeak parameters for each channel-dataset pair.

Simultaneous Fits by Tower

Fitting each bolometer completely independently, however, is problematic because of the background secondary peaks. Looking at the data from an individual channel-dataset pair, we find that these spectral features are significant enough to affect the photopeak fit if we were to ignore them. Statistics, however, are much too low to model them properly. As such, we decide to perform a simultaneous fit: the seven 2615-keV photopeak parameters (N_{sig} , μ , σ , ρ_{lo} , ρ_{hi} , R_{lo} , and R_{hi} in Eq. 9.1) are defined on the channel-dataset pair level,

while the remaining six describing the surrounding spectrum are “global.” This categorizing is summarized in Table 9.3, together with parameter ranges.

Using a single set of global parameters for the entire CUORE detector, with its 984 functioning channels, was rejected—such a large fit is computationally intensive and hard to tune, while offering few advantages. The statistics necessary for a high quality background fit can be achieved with a much smaller grouping of crystals. Furthermore, background levels are not uniform throughout the detector due to the 19 towers’ different positioning with respect to the 10-mK cryostat shield (see Chapter 5). Thus, we decide to break up the line shape modeling by tower, carrying out 19 simultaneous fits.

It is important to note that, besides the Compton and background rates and the background slope, our global parameters are mathematically defined as *ratios* with respect to parameters of the non-global main photopeak. This means that, while there is only one fit parameter value for any of these, there are in fact up to 104 *physical* values—two datasets and 52 channels in a CUORE tower—for quantities like the roll-off point of the multi-Compton plateau.

Unbinned Maximum Likelihood Fitting

We perform our simultaneous fits using the unbinned extended maximum likelihood (UEML) method. Our software of choice is RooFit, a CERN ROOT extension originally developed for the BaBar experiment and well suited for analyses like this one [132, 133]. The likelihood function $\mathcal{L}(\text{Data}|\boldsymbol{\theta})$ quantifies the plausibility of the fit parameter values $\boldsymbol{\theta}$ given the data:

$$\mathcal{L}(\text{Data}|\boldsymbol{\theta}) = \prod_{d,c} \left(\frac{e^{-\lambda_{d,c}} \lambda_{d,c}^{N_{d,c}}}{N_{d,c}!} \prod_{i=1}^{N_{d,c}} \frac{f(\text{Data}_{d,c,i}; \boldsymbol{\theta}_{d,c})}{\lambda_{d,c}} \right). \quad (9.9)$$

In Eq. 9.9, N is the total number of events, while λ is the quantity predicted from the fit—ideally, very close to N . The function f is our fit PDF, as defined in Eq. 9.1, with the vector $\boldsymbol{\theta}$ representing all the free parameters therein. The subscript indices d and c refer to the CUORE dataset and channel, into which the data and parameters are grouped.

Best agreement between the function and the data occurs when the likelihood is maximized. Thus, when we are fitting, we are essentially finding the set of parameters that solve the differential equation

$$\frac{\partial \mathcal{L}}{\partial \theta_i} = 0. \quad (9.10)$$

In practice, RooFit uses the ubiquitous MINUIT algorithm to iteratively perform the mathematically-equivalent calculation of minimizing the “negative log-likelihood” (NLL)—the negative natural logarithm of Eq. 9.9:

$$\text{NLL} = -\ln \mathcal{L} = -\sum_{i=1}^N \ln \left(\frac{f(\text{Data}_i|\boldsymbol{\theta})}{\lambda} \right) - N \ln \lambda + \lambda + \ln N!, \quad (9.11)$$

written for simplicity for a single channel-dataset pair.

UEML is among the best-performing fitting approaches, especially for scenarios with relatively low statistics in certain components of the PDF. Other techniques, such as binned likelihood or χ^2 fitting, utilize a histogram representation of the data, which is an approximation. UEML, conversely, allows us to individually consider each spectral data point with its exact energy, as represented by the index i in Eqs. 9.9, 9.10, and 9.11. Of course, it is also the most computationally intensive. This is not a problem for an analysis such as this one, which is not run frequently and repetitively.

Parameters and Three-Step Iterative Fitting

The line shape fit model as described in Section 9.3 is very complicated, and it is important to facilitate its convergence with appropriate parameter initialization and constraints. Our PDF, given in Eqs. 9.1 – 9.8, contains a total of 13 free parameters. Seven of these describe the photopeak and are defined on the channel-dataset pair level, while the remaining six have “global,” or tower-level significance. We initialize each of these values in one of three ways:

1. Using theoretical values. This is the logical approach for fundamental physical values, namely the energies of the photopeak (at ~ 2615 keV), the X-ray escape peak (at ~ 2585 keV), and the coincidence peak (at ~ 2687).
2. Estimating from the data. We do this for values inherent to the detector/experiment. The resolution parameter σ is taken from the RMS deviation of the fit-region histogram, which is dominated by the main Gaussian of the 2615-keV photopeak. The event multiplicities N_{sig} , N_{comp} , and N_{bkg} are estimated from integrals of the appropriate portions of the histogram, as illustrated in Fig. 9.2. These three numbers give totals of counts, which are directly affected by detector exposure. Our actual fit function is written in terms of event *rates*— n_{sig} , n_{comp} , and n_{bkg} —to avoid this dependence.

$$\begin{aligned} n_{\text{sig}}^{d,c} &= N_{\text{sig}}^{d,c} / \varepsilon_{d,c} \\ n_{\text{comp}} &= N_{\text{comp}} / \sum_{d,c} \varepsilon_{d,c} \\ n_{\text{bkg}} &= N_{\text{bkg}} / \sum_{d,c} \varepsilon_{d,c} , \end{aligned} \tag{9.12}$$

where $\varepsilon_{d,c}$ is the exposure of each channel-dataset pair in kilogram-years.

3. Empirical tuning. The energy and amplitude ratios of the photopeak shoulders, as well as the amplitude ratios of the X-ray escape and the coincidence peaks, are not possible to predict theoretically or to reliably estimate from the un-fitted histogram. We first

selected the initial values of these parameters by eye, and then tuned them slightly by trial and error. In the end, the subpeaks specifically are tuned with a more intricate iterative scheme described below.




Figure 9.2: The CUORE calibration spectrum between 2530 and 2720 keV—the line shape fit region. Shaded are the three regions we integrate to initialize the background (green), signal (red), and multi-Compton (blue) rate parameters. The full-detector spectrum is an illustration only—integration is done on tower or channel-dataset pair level (see Table 9.3).

The parameter ranges are originally set by eye, and then again empirically tuned such that the fit converges well, but parameters do not frequently end up railing at their upper or lower limits. When specifying the parameter constraints, we do so with absolute numbers for those parameters initialized to theoretical or trial-and-error values (techniques 1 and 3 above). These ranges are summarized in Table 9.3. For the event rates and peak width—initialized from the spectrum histogram (technique 2)—limits are defined relatively.

The two subpeaks of the 2615-keV line present a challenge. They constitute a very subtle correction to the PDF, and the statistics of the phenomenon they model are low because we define them on the channel-dataset level with the rest of the photopeak. They can easily end up converging not on the peak shoulders as intended, but “folded into” the bulk of the photopeak, the X-ray escape line on the lower-energy side, or even a random background fluctuation on the higher-energy side. Thus, the subpeaks need to be well-initialized and constrained. However, due to the sheer number of free parameters in our complicated model, we still risk the overall fit having imperfect convergence—even if it looks good—when these parameters are kept free. This problem is solved if we fix the two subpeak energy ratios

Table 9.3: Summary of the CUORE spectral line shape fit parameters and their constraints. “Significance” refers to the level at which the parameter is defined—by channel-dataset pair (“Ch-Ds”, for the photopeak) or for the entire tower (“Global”, for background). A range “from data” means the value’s constraints depend on the initialization value taken from an integral of a histogram of the data. The three-step iterative fitting approach optimizes the photopeak shoulders.

Parameter	Significance	Range: Iteration 1	Range: Iteration 2	Range: Iteration 3
n_{sig}	Ch-Ds	from data	fixed from It. 1	from data
μ	Ch-Ds	2608 – 2622 keV	fixed from It. 1	2608 – 2622 keV
σ	Ch-Ds	from data	fixed from It. 1	from data
R_{lo}	Ch-Ds	0.0 – 0.7	0.0 – 0.7	0.0 – 0.7
R_{hi}	Ch-Ds	0.0 – 0.9	0.0 – 0.9	0.0 – 0.9
ρ_{lo}	Ch-Ds	fixed 0.9962	0.992 – 0.9999	fixed from It.2
ρ_{hi}	Ch-Ds	fixed 1.0055	1.0007 – 1.015	fixed from It.2
R_{Xray}	Global	10^{-4} – 10^{-2}	fixed from It. 1	10^{-4} – 10^{-2}
R_{coin}	Global	10^{-10} – 10^{-2}	fixed from It. 1	10^{-10} – 10^{-2}
ρ_{coin}	Global	1.0256 – 1.0294	fixed from It. 1	1.0256 – 1.0294
n_{comp}	Global	from data	fixed from It. 1	from data
n_{bkg}	Global	from data	fixed from It. 1	from data
m_{bkg}	Global	–1 – 0	fixed from It. 1	–1 – 0

ρ_{lo} and ρ_{hi} . However, the manifestation of shoulders in the photopeak varies noticeably channel-to channel. High-resolution bolometers have prominent subpeaks, while in low-resolution cases the main peak tends to swallow them. These parameters are defined with channel-dataset pair significance, and if we are to fix them, it must be to a unique value for each one. We address this issue by performing a three-step iterative fit:

1. During the first iteration, the subpeak energy ratios are fixed to approximate, global values determined empirically from parameter distributions. This fit converges well for the entire detector, and is in fact usable on its own. It is, however, visibly non-ideal for some channel-dataset pairs.
2. Next, we fix all parameters to their results from Iteration 1 except the four ratios defining the two subpeaks. Those are now *free*, and we run the fit to determine a good

set of unique subpeak parameter values for each channel-dataset pair. This heavily constrained fit, of course, also converges and looks decent.

3. Finally, because the changes to the subpeak parameter values have skewed the rest of the PDF, we re-run the fit with the shoulders fixed, but now with “proper” values taken from Iteration 2. The remaining parameters shift slightly to compensate, and this is the final result.

The parameter ranges and fixing in our three-step approach are summarized in Table 9.3.

9.5 Results




Figure 9.3: CUORE 2615 keV ^{208}Tl line shape fit result for one representative tower (number 14). The fit components are shown with dashed blue lines: (a) the triple-Gaussian photopeak, (b) the multi-Compton smeared step, (c) the six-Gaussian ~ 30 -keV X-ray escape peak, (d) the linear background, and (e) the Gaussian $(2615 - 1022 + 583 + 511)$ keV coincidence peak.

To model the line shape and detector response of CUORE for use in physics analyses, we fit calibration data around the ^{208}Tl 2615-keV γ peak with the model PDF from Section 9.3. Utilizing the three-step iterative technique described in Section 9.4, we performed a simultaneous unbinned extended maximum likelihood (UEML) fit on each of CUORE’s 19 towers between 2530 and 2720 keV. The same quality cuts were applied to the calibration data as the background data for $0\nu\beta\beta$ analysis, with the exception of multiplicity. This selection helps to remove cosmic muons, surface α decay events where the nuclear recoil and α particle are detected in adjacent bolometers, and multi-crystal Compton-scattering γ -rays. Calibration

data is quite “hot,” however, with many random coincidences, so including the multiplicity cuts was found to significantly reduce the statistics of our fit. The same poorly-performing channels were cut and thus, our 19 fits together include 1811 channel-dataset pairs, or $\sim 92\%$ of active¹ channels. In other words, with all 19 fits converging successfully, we have 1811 triple-Gaussian line shapes for fitting the ^{130}Te $0\nu\beta\beta$ -decay region of interest.

The combined fit result from tower 14, which is representative of the other 18 fits, is plotted in Fig. 9.3. Agreement of the fit with data is given by pulls,

$$\text{pull} = \frac{y_{\text{fit}} - y_{\text{data}}}{\sigma_{\text{fit}}}, \quad (9.13)$$

plotted overhead. The dashed blue lines show the separate components of the fit PDF given in Eqs. 9.1 through 9.8. The photopeak, captioned (a) in the figure, represents the sum of 96 triple-Gaussian fits, one for each channel-dataset pair in that tower. The line shapes of two typical channels from tower 14 are given in Fig. 9.4. The remaining components (b) through (e) have global significance, and are actually fit to the data we see in Fig. 9.3. This is why the background portions of the red curves in the by-channel plots in Fig. 9.4 appear somewhat presumptuous. The forms of the secondary peaks and the roll-off of the multi-Compton plateau in Fig. 9.3 look slightly smeared because, while they are defined with tower-level parameters, said parameters are ratios to the photopeak. Thus, though there is only one Compton *fit*, there are again 96 individual curves with slightly different roll-offs right below the 2615-keV line’s mean μ (see Section 9.3 and Eq. 9.7. The same is true of the positions of the X-ray escape and coincidence peaks.




Figure 9.4: CUORE line shape fit results from two typical channels: **(a)** 699 and **(b)** 705, from tower 14 (Fig. 9.3). The photopeak parameters are unique for every channel, while the background ones are common for the entire tower.

¹984 of CUORE’s 988 bolometers (99.6 %) have functional electrical connections, and we consider them “live.” We hope to increase the number of channels used in analysis through detector optimization in the future.

The 19 tower-wise fits are completely independent from each other, but we add them together to construct the summary graph in Fig. 9.5. The combined red curve is overlaid on the full CUORE calibration spectrum in the 2530 – 2720 keV domain. Besides neatly combining into one picture the entirety of the line shape fit results, this plot also allows us to better gauge visually the systematic agreement of our model with the true detector response. Instead of evaluating the pulls like we do in Fig. 9.3, we draw the ratio of the curve to the histogram. We choose to do so because, with the $\sim 19\times$ higher statistics, the statistical uncertainties on each data bin become so small that the value of the pull becomes too large to give good intuitive understanding, particularly at the top of the peak. We can see by eye and from the ratio plot that the semi-effective fit model discussed in detail in Section 9.3 is indeed a very good and largely unbiased approximation of CUORE’s behavior. The largest amount of systematic deviation can be seen in the high-energy tail, particularly in the significant excess of physical events between 2645 and 2665 keV. We explored covering these with additional fit function components, but decided against further complicating the PDF with additional degrees of freedom until a scientific explanation for the excess has been agreed upon. Overall, the high-energy deviations do not affect the photopeak (a) itself, which is the only part of this fit to go on to see use in other analyses.




Figure 9.5: Sum of the 19 tower-level line shape fit results overlaid with full (1811 channel-dataset pairs) data. The ratio of the data to the curve is plotted above. The fit components are shown with dashed blue lines: (a) the triple-Gaussian photopeak, (b) the multi-Compton smeared step, (c) the six-Gaussian ~ 30 -keV X-ray escape peak, (d) the linear background, and (e) the Gaussian (2615 – 1022 + 583 + 511) keV coincidence peak.

Energy Resolution

In addition to providing us with the spectral line shape PDF, these fits reveal the CUORE detector's energy resolution, which is the measured full width at half maximum (FWHM) of a monoenergetic peak. The distribution of the FWHMs of the 2615-keV line fits for all 1811 channel-dataset pairs is given in Fig. 9.6a. We define the effective average energy resolution as the exposure-weighted harmonic mean H of the FWHM's:

$$H = \frac{\sum_{d,c} \varepsilon_{d,c}}{\sum_{d,c} \frac{\varepsilon_{d,c}}{\text{FWHM}_{d,c}}}, \quad (9.14)$$

where d and c are the indices for summing over the datasets and channels, $\varepsilon_{d,c}$ is the exposure of each channel-dataset pair in kilogram-years, and $\text{FWHM}_{d,c}$ is the energy resolution of that pair. In this two-dataset release, we find this to be equal to 8.0 keV, shown by the dashed red line in the figure. We observe some energy-dependence of the resolution, as well as worsening due to pile-up effects in high-rate data. Adjusting for these factors, the effective resolution at the ^{130}Te Q-value of 2528 keV in background data improves to (7.7 ± 0.5) keV [33].




Figure 9.6: **(a)** Distribution of the full width at half maximum (FWHM) of the line shape fits for all 1811 channel-dataset pairs. The average resolution at 2615 keV is 8.0 keV, shown by the dashed line. **(b)** Separate FWHM distributions for CUORE's two datasets, showing resolution improvement thanks to detector optimization.

This is inferior to the CUORE-0 performance of ~ 5 keV FWHM [76, 79]. After the publication of an analysis of the data examined in this thesis, CUORE has been engaged in an ongoing detector optimization campaign to reduce vibrational and electrical noise. Doing so together with analysis refinements, we seek to improve CUORE's performance to this level in the future. We believe this to be possible in light of the already significant gains achieved in about one month of optimization between datasets 1 and 2. As can be seen in Fig. 9.6b, the average FWHM of the 2615-keV line in calibration data decreased from 9.0 to 7.4 keV. Resolution at 2528 keV in background data changed from (8.3 ± 0.4) to (7.4 ± 0.7) keV.

We can see that the FWHM distribution for dataset 2 (blue) is much tighter, with fewer “fuzzy” channels in the tail than that of dataset 1 (red). Finally, we significantly increased the number of “good” channels passing the quality cuts—from 876 to 935 out of 984.

Chapter 10

CUORE Analysis: Fitting the ^{130}Te $0\nu\beta\beta$ Decay ROI

10.1 Introduction

big advantage in hunting for this signature, because it lets us reliably separate out the monoenergetic $Q_{\beta\beta}$ peak from the continuous $2\nu\beta\beta$ spectrum. In fact, the background near $Q_{\beta\beta}$ in CUORE and its predecessors is essentially flat [33, 76, 77, 78, 79].

The $\beta\beta$ sum energy spectrum of $2\nu\beta\beta$ decay, as we discussed in Chapter 3, has a bell-like shape between 0 and the Q-value ($Q_{\beta\beta}$)—the full decay energy—due to a part of the energy being carried away by the undetected neutrinos. In the event of a $0\nu\beta\beta$ decay, the electrons would together carry 100% of decay energy. Thus, the spectral signature we look for when searching for this hypothetical process is a small peak at $Q_{\beta\beta}$, right at the tail end of the bell curve (Fig. 3.5).

The CUORE detector is an array of bolometers—ultra-cold instruments measuring the thermal signatures of absorbed particles (see Chapter 4),—and has inherently excellent energy resolution. This gives us an advantage in hunting for this signature, because we can reliably separate the monoenergetic $Q_{\beta\beta}$ peak from the $2\nu\beta\beta$ continuum—a phenomenon illustrated in Fig. 3.7. In fact, as we can see from Fig. 4.13, the background near $Q_{\beta\beta}$ in the CUORE data is essentially flat [33].

Our analysis centers on fitting a region of interest, or ROI, around $Q_{\beta\beta}(^{130}\text{Te}) = (2527.515 \pm 0.013)$ keV [28, 29, 30] between 2465 and 2575 keV with a PDF that models the background in that region and also allows for the presence of a $0\nu\beta\beta$ decay peak. Specifically, we fit a function comprising a flat background and two triple-Gaussian lines (as described in Chapter 9), for the ^{60}Co 2-photon peak at ~ 2505 keV and at $Q_{\beta\beta}$. In this chapter, we present the methods and results for two alternative approaches to fitting the CUORE ROI: Unbinned Extended Maximum Likelihood (UEML) with RooFit, and Markov Chain Monte Carlo (MCMC) with the Bayesian Analysis Toolkit (BAT).

10.2 Fit Function

We designate the ROI around the ^{130}Te double-beta decay Q-value as the portion of CUORE physics spectrum between 2465 and 2575 keV. This is the domain of our fit function, which describes all of the spectral features we find there. Specifically, this is a flat background comprised mainly of degraded α and random coincidence events, the ^{60}Co (1173 + 1442)-keV two-photon coincidence peak, and, potentially, $0\nu\beta\beta$ -decay events at $Q_{\beta\beta} = (2527.515 \pm 0.013)$ keV [28, 29, 30]. For both of these lines, we use the PDF from our detector response fit, as described in Chapter 9, Eq. 9.2. This is unique for each of the 1811 channel-dataset pairs, with shape proportions fully constrained. The overall form of the ROI fit function is as follows in Eq. 10.1:

$$f_{\text{ROI}}(E) = \sum_{d,c} \left(\eta_C \eta_\gamma^d \varepsilon^{d,c} R_{\beta\beta} \cdot \text{PDF}_{\text{LS}}^{d,c}(\beta\beta; E) \right. \\ \left. + \eta_\gamma^d \varepsilon^{d,c} R_{\text{Co}} e^{-t/\tau_{\text{Co}}} \cdot \text{PDF}_{\text{LS}}^{d,c}(\text{Co}; E) \right. \\ \left. + \varepsilon^{d,c} R_{\text{bkg}}^d \Delta E \right), \quad (10.1)$$

where $\text{PDF}_{\text{LS}}^{d,c}(\beta\beta; E)$ refers to the triple-Gaussian line shape PDF, appropriately scaled and shifted to $Q_{\beta\beta}$ for the $0\nu\beta\beta$ signal:

$$\text{PDF}_{\text{LS}}(\beta\beta; E) = \frac{1}{a\sigma\sqrt{2\pi}} e^{-\frac{(E-\rho_{\beta\beta}\cdot\mu)^2}{2(a\sigma)^2}} \\ + \frac{R_{\text{lo}}}{a\sigma\sqrt{2\pi}} e^{-\frac{(E-\rho_{\beta\beta}\cdot\rho_{\text{lo}}\cdot\mu)^2}{2(a\sigma)^2}} + \frac{R_{\text{hi}}}{a\sigma\sqrt{2\pi}} e^{-\frac{(E-\rho_{\beta\beta}\cdot\rho_{\text{hi}}\cdot\mu)^2}{2(a\sigma)^2}}. \quad (10.2)$$

Analogously, for the ^{60}Co 2-photon coincidence peak,

$$\text{PDF}_{\text{LS}}(\text{Co}; E) = \frac{1}{a\sigma\sqrt{2\pi}} e^{-\frac{(E-\rho_{\text{Co}}\cdot\mu)^2}{2(a\sigma)^2}} \\ + \frac{R_{\text{lo}}}{a\sigma\sqrt{2\pi}} e^{-\frac{(E-q\cdot\rho_{\text{Co}}\cdot\rho_{\text{lo}}\cdot\mu)^2}{2(a\sigma)^2}} + \frac{R_{\text{hi}}}{a\sigma\sqrt{2\pi}} e^{-\frac{(E-q\cdot\rho_{\text{Co}}\cdot\rho_{\text{hi}}\cdot\mu)^2}{2(a\sigma)^2}}. \quad (10.3)$$

The superscript indices d and c indicate summing over the two datasets and all channels. There are a total of just five free parameters in this fit:

- The $0\nu\beta\beta$ decay signal rate $R_{\beta\beta}$, our main parameter of interest, measured in counts/(kg y). Scaling appropriately by isotope mass in the detector, $R_{\beta\beta}$ is converted to the physical decay rate $\Gamma_{0\nu\beta\beta}$:

$$\Gamma_{0\nu\beta\beta} = R_{\beta\beta} \cdot \frac{\mathcal{M}_{\text{TeO}_2}}{aN_A}, \quad (10.4)$$

where $\mathcal{M}_{\text{TeO}_2}$ is the molar mass of TeO_2 (~ 0.16 kg/mol), a is the ^{130}Te isotopic abundance ($\sim 34\%$), and N_A is Avogadro's number ($6.022 \cdot 10^{23}$ molecules/mol).

- The ^{60}Co two-photon event rate R_{Co} in the detector, again in counts/(kg y).
- The ^{60}Co peak position quenching factor q (in the exponents of Eq. 10.3). This takes on a value very close to 1, but is allowed to float due to some uncertainty about the reconstruction of 2- γ sum events always seen in CUORE-type bolometers [76, 77, 78, 79]. We strongly believe that this is an artifact of position dependence, and hence do not apply any correction to the $Q_{\beta\beta}$ peak where the two electrons will be absorbed very close to each other.
- The dataset-dependent background rate R_{bkg}^d in counts/(kg y), giving two parameters. In future CUORE analyses, with many more datasets included and the detector operating in a single consistent configuration, we expect R_{bkg} to be a single parameter like the other two rates. However, because detector performance and cut efficiencies changed so much with the optimization campaign carried out between datasets 1 and 2, we deemed it more appropriate to fit two separate rates on this occasion.

The remaining parameters are constants, summarized in Table 10.1. η_C is the double-beta event containment efficiency of CUORE—a global probability of detecting both electrons from a $\beta\beta$ decay—equal to $(88.345 \pm 0.085)\%$. η_γ^d is the dataset-dependent combined cut efficiency of $(85.67 \pm 3.42)\%$ for dataset 1 and $(93.96 \pm 2.89)\%$ for dataset 2. These two values are so different thanks to performance improvement resulting from detector hardware optimization carried out between the two datasets. $\varepsilon^{d,c}$ is the exposure in kilogram years, unique for every channel-dataset pair.

We take into account the natural decay of ^{60}Co in the exponential term in Eq. 10.1—the isotope has a relatively short half-life of about 5.3 y, which can have a non-negligible effect. $\tau_{\text{Co}} = 2.4 \times 10^8$ is the ^{60}Co lifetime in seconds, while t is the approximate time, which we take as the mean event time of the dataset in our likelihood fit or as the start time of the dataset in the Markov Chain Monte Carlo fit (see below in Sections 10.3 and 10.4).

Both the $0\nu\beta\beta$ decay and ^{60}Co 2- γ peaks are described by triple-Gaussian functions. These are unique to every channel-dataset pair, and are taken from the results of the detector response analysis discussed at length in Chapter 9. Because our detector energy calibration is not necessarily perfect, we define the position of $Q_{\beta\beta}$ using a ratio with respect to the 2615-keV line fitted mean μ :

$$\rho_{\beta\beta} = \frac{2527.515\text{keV}}{2614.511\text{keV}}.$$

We have an equivalent constant for ^{60}Co 2- γ coincidence events,

$$\rho_{\text{Co}} = \frac{2505.720\text{keV}}{2614.511\text{keV}},$$

Table 10.1: Values of constant parameters in the ROI fit function in Eqs. 10.1 – 10.3. Line shape parameters with channel-dataset pair significance (ch-ds) are taken directly from the fit result in Chapter 9.

Parameter	Value (dataset 1)	Value (dataset 2)	Description
η_C	$(88.345 \pm 0.085)\%$	$(88.345 \pm 0.085)\%$	$\beta\beta$ containment efficiency
η_γ	$(85.67 \pm 3.42)\%$	$(93.96 \pm 2.89)\%$	Combined cut efficiency (dataset)
ε	—	—	Exposure in kg·y (ch-ds)
t	—	—	Dataset time for ^{60}Co decay
τ_{Co}	2.4×10^8 s	2.4×10^8 s	^{60}Co lifetime
a	$(91.5 \pm 4.6)\%$	$(100.0 \pm 9.3)\%$	Resolution scaling factor (dataset)
σ	—	—	σ from 2615-keV calibration fit (ch-ds; Eq. 9.2, Fig. 9.6)
μ	—	—	3-Gauss. μ of 2615-keV fit (ch-ds; Eq. 9.2)
$\rho_{\beta\beta}$	2527.515/2614.511	2527.515/2614.511	$Q_{\beta\beta}$ w.r.t. 2615 keV
ρ_{Co}	2505.720/2614.511	2505.720/2614.511	^{60}Co 2- γ sum w.r.t. 2615 keV
ρ_{lo}	—	—	3-Gauss. left subpeak energy ratio (ch-ds; Eq. 9.2)
ρ_{hi}	—	—	3-Gauss. right subpeak energy ratio (ch-ds; Eq. 9.2)
R_{lo}	—	—	3-Gauss. left subpeak integral ratio (ch-ds; Eq. 9.2)
R_{hi}	—	—	3-Gauss. right subpeak integral ratio (ch-ds; Eq. 9.2)

but with yet an additional multiplicative factor q , a non-unity quenching factor for multi-photon events, which is a free parameter as described above. Additionally, we observe that CUORE’s energy resolution exhibits a level of energy (Fig. 10.1) and rate dependences, with the former resulting from pile-up effects. We account for this in the ROI fit with the dataset-dependent scaling factor a , which is equal to $(91.5 \pm 4.6)\%$ and $(100.0 \pm 9.3)\%$ in datasets 1 and 2, respectively. The values were determined by fitting the 2615-keV line shape PDF to other peaks in the ^{232}Th calibration spectrum, and to ^{208}Tl , ^{40}K , and single-photon ^{60}Co lines in background data. Again, they differ between the two datasets thanks to detector optimization. Pre-scaling resolution σ , 2615-keV photopeak mean μ , and the subpeak amplitude and energy ratios (R_{lo} , R_{hi} , ρ_{lo} , and ρ_{hi}) are taken directly from line shape fit results, all with channel-dataset pair significance.




Figure 10.1: CUORE effective energy resolution as a function of energy in datasets 1 and 2. Calculation and figure by J. Cushman, from [140].

10.3 UEML ROI Fitting Technique with RooFit

Within the CUORE collaboration, in the interest of confirming robustness and reproducibility, three analyses fitting the “physics” data from datasets 1 and 2 between 2465 and 2575 keV have been performed. All of these worked with an identical version of the data, containing 155 events in the ROI after the full set of cuts on pulse shape, pulse multiplicity (keeping only events with single pulses in a window in a single crystal in the entire detector), and removing “bad intervals.” In all three cases, the fit was carried out using the model given in Eqs. 10.1 – 10.3 in Section 10.2 or one mathematically equivalent to it, but using three different code bases. Two are Unbinned Extended Maximum Likelihood (UEML) fits with the RooFit Toolkit for Data Modeling [132, 133], one presented in [140] and the second

discussed in detail in this section. These highly congruent analyses correspond to the main result published in CUORE's recent Physical Review Letters article, [33]. The third analysis, also consistent and mentioned in the paper, is a fully-Bayesian approach addressed in Section 10.4.

The process of fitting the ROI with the UEML technique is not dissimilar to the spectral line shape fit described in Chapter 9. However, the event rate in background data is much lower due to the absence of the calibration sources, and we are fitting away from any major γ lines. As such, fit statistics are low and we perform a single global fit on the entire CUORE detector, without dividing it by tower or some other type of channel grouping. As we have already mentioned in Section 10.2, there are only five free parameters in this case, and we define our likelihood function as

$$\mathcal{L}(\text{Data}|\boldsymbol{\theta}) = \prod_{d,c} \left(\frac{e^{-\lambda_{d,c}} \lambda_{d,c}^{N_{d,c}}}{N_{d,c}!} \prod_{i=1}^{N_{d,c}} \frac{f_{\text{ROI}}(\text{Data}_{d,c,i}; \Gamma_{0\nu\beta\beta}, \boldsymbol{\theta}_d)}{\lambda_{d,c}} \right). \quad (10.5)$$

Here, $\Gamma_{0\nu\beta\beta}$ is our main physical parameter of interest, the rate of apparent $0\nu\beta\beta$ decay events detected, while the vector $\boldsymbol{\theta}$ represents the other four free variables of the fit (R_{Co} , two values of R_{bkg} , and ^{60}Co quenching factor q). N and λ are the true and expected numbers of events in the ROI, respectively. As before, subscripts d and c indicate the dataset and channel of a given event. Unlike in the detector response analysis, here we do not have any channel-dependent free parameters, and only the background rate depends on dataset. Multiplying over both remains necessary, however, given that the line shape function and the exposure are unique for every channel-dataset pair.

Maximizing the likelihood essentially means converging on the $0\nu\beta\beta$ decay event rate that has the highest plausibility given the fit. Mathematically, this is solving the following differential equation system:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \theta_i} &= 0 \\ \frac{\partial \mathcal{L}}{\partial \Gamma_{0\nu\beta\beta}} &= 0. \end{aligned} \quad (10.6)$$

RooFit does this numerically, using the MINUIT package. In fact, as is the common practice, it is actually performing the mathematically-equivalent calculation of iteratively minimizing the negative log-likelihood (NLL):

$$\begin{aligned} \text{NLL} &= -\ln \mathcal{L}(\text{Data}|\Gamma_{0\nu\beta\beta}) \\ &= \sum_{d,c} \left(- \sum_{i=1}^{N_{d,c}} \ln \left(\frac{f(\text{Data}_{d,c,i}; \boldsymbol{\theta}_d, \Gamma_{0\nu\beta\beta})}{\lambda_{d,c}} \right) - \ln \lambda_{d,c} + \lambda_{d,c} + \ln N_{d,c}! \right). \end{aligned} \quad (10.7)$$

Unbinned likelihood fits handle each event individually, using its exact energy value. As we have already mentioned in Section 9.4, this gives them a great advantage over binned likelihood, χ^2 , and some other approaches in cases with very low statistics. With only 155 events in the entire 110-keV window of the ROI, this analysis is a good example of such a scenario.

With just five free parameters, constraints and initializations are not a very significant issue in this fit the way they are for the detector response, in terms of convergence. There is, however, an analytical question when it comes to the rate of $0\nu\beta\beta$ decay events $R_{\beta\beta}$. Physically, this value can only be zero or positive. In terms of this spectrum, this means we would have an uninterrupted flat background or a peak. In that sense, it is logical to constrain $R_{\beta\beta}$ to be greater than or equal to zero. Statistical fluctuations do occur, however. For example, we saw a downward dip at the Q-value in the Cuoricino data [77, 78]. In such a scenario, the fit would converge on a zero rate, but with a strongly asymmetric positive-only uncertainty. In other words, our result would be biased. We can avoid this by freeing the parameter to take on negative values. Of course, we could not interpret negative decay rates as such in a physical sense, and would consider only the positive range for setting limits on the decay half-life. Such limits would be very similar for either case, and the parameter constraint question is almost a philosophical one. In the end, to make the fullest exploration of our data possible, we perform both versions of the fit.

10.4 MCMC ROI Fitting Technique with BAT

BAT—the Bayesian Analysis Toolkit—is another software package for fitting and statistics compatible with CERN ROOT [141]. As its name suggests, BAT is designed for performing fully-Bayesian analyses. First used for $0\nu\beta\beta$ decay analysis in the context of the GERDA ^{76}Ge experiment, it is very well suited for investigations like ours.

The Bayesian approach to statistics, named after its originator, the Reverend Thomas Bayes, is based on an evidential interpretation of probability. Probabilities are understood as “degrees of belief” in a hypothesis, which is not simply either true or false. Rather, any hypothesis has some probability of being true, between zero and one. In light of new data, we can update our prior hypothesis to a new, more believable, posterior hypothesis, by means of Bayes’ theorem:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \quad (10.8)$$

$P(A|B)$ is the conditional probability of scenario A being true given B , and vice versa for $P(B|A)$. $P(A)$ and $P(B)$ are the “marginal probabilities,” and refer to the chances of events A and B occurring independently of each other.

When undertaking statistical modeling in the Bayesian way, we must specify a “prior distribution” of any unknown parameters, which is then updated using experimental results to get the “posterior distribution.” This, in turn gives us the answer in the form of an

estimator and its uncertainty. Using BAT to fit the CUORE data with the same model described in Eqs. 10.1 – 10.3, we again use the natural logarithm of the likelihood function (Eq. 10.7, but non-negative). Applying Bayes' Theorem from Eq. 10.8 to the log-likelihood of the data given the function parameters, we extract the log-likelihood of a set of parameter values given the data:

$$\ln\mathcal{L}(\boldsymbol{\theta}|\text{Data}) = \ln\mathcal{L}(\text{Data}|\boldsymbol{\theta}) + \sum_{i=1}^{M=5} \left(\ln\pi(\theta_i) \right), \quad (10.9)$$

where $\pi(\theta_i)$ is the prior probability of parameter θ_i and we have dropped a constant coming from normalization.

Prior probability distributions are assigned by the analyst, making reasonable choices using existing knowledge and, in some cases, earlier experimental results. Naturally, for $0\nu\beta\beta$ decay, an undiscovered process, we take a uniform prior. Though we do include a quenching factor to slightly adjust the position of the ^{60}Co 2- γ line due to non-linear reconstruction of the energy of multi-photon events, we can assume it to be close to its theoretically correct location. For the background and ^{60}Co rates, we choose fairly broad ranges that seem reasonable based on previous measurements and CUORE performance predictions. To be maximally consistent with the UEML fit, we assign flat prior distributions to all parameters. These are wide enough to allow up to an estimated 5σ deviation in the posterior for all four of the rates, though the cobalt quenching's range is much tighter. Since the priors are proper distributions, we do not have parameter initialization with specific values the way we do in the likelihood fit. As with the RooFit analysis, we evaluate two versions that either allow or forbid negative values of rates. A Bayesian purist might suggest that the positive-only approach is more correct, because we have a priori knowledge that decay rates cannot be negative. The counter-argument of a biased result in the event of a downward statistical fluctuation still applies.

The output of Eq. 10.9 gives us the posterior probability density function (PDF), for a given set of parameter values. With BAT, we converge on the most probable ones using the Markov Chain Monte Carlo (MCMC) technique [141, 142, 143]. A Monte Carlo technique, as we know, is one based on random number generation constrained by some PDF. A Markov Chain is a sequence of random numbers, within a limiting distribution, in which the next ($t+1$) element depends on the current (t), but not any other prior history. With our MCMC applied to spectral modeling, we produce a Markov Chain limited by our posterior PDF. BAT achieves this with the Metropolis algorithm, the original and most popular approach to this technique [141, 144]. In effect, this is somewhat like a guided random walk. Using an arbitrary symmetrical probability density $g(x|y)$ as a “proposal distribution”, we choose a candidate x' for the next state of the system that is currently in state x_t . Next, we calculate an acceptance ratio defined as $\text{PDF}(x')/\text{PDF}(x_t)$, and compare it to a generated uniform random number on $[0, 1]$. If said number is greater than the acceptance ratio, the candidate x' is rejected, $x_{t+1} = x_t$, and the process repeats again from the same starting state. Otherwise, $x_{t+1} = x'$ and we continue from the new state. This goes on until we reach

the maximum limit of iterations, hopefully having converged. In our case, x represents the set of parameters in the model, and we define the PDF with Eq. 10.9. In practice, it is common to run multiple MCMCs. In the analysis presented in this paper, we use 8 Markov Chains of $1 \cdot 10^6$ samples each.

10.5 ROI Fit Results

Table 10.2: ROI fit results, allowing negative values of $0\nu\beta\beta$ decay rate parameter. Note that $\Gamma_{0\nu\beta\beta}$ is the *best fit rate*, not a measurement or limit.

Parameter	RooFit value	BAT value
$R_{\text{bkg}}(\text{D.S. 1})$ (cts./keV/kg/y)	$(1.46^{+0.20}_{-0.19}) \cdot 10^{-2}$	$(1.46 \pm 0.20) \cdot 10^{-2}$
$R_{\text{bkg}}(\text{D.S. 2})$ (cts./keV/kg/y)	$(1.59^{+0.19}_{-0.18}) \cdot 10^{-2}$	$(1.60 \pm 0.18) \cdot 10^{-2}$
R_{Co} (cts./kg/y)	$(2.3^{+0.9}_{-0.8}) \cdot 10^{-1}$	$(2.3 \pm 0.8) \cdot 10^{-1}$
q	1.0003 ± 0.0005	1.0003 ± 0.0005
$R_{\beta\beta}$ (cts./kg/y)	$(-1.31^{+0.49}_{-0.37}) \cdot 10^{-1}$	$(-1.31 \pm 0.42) \cdot 10^{-1}$
$\Gamma_{0\nu\beta\beta}$ (10^{-25} y^{-1})	$-1.01^{+0.38}_{-0.29}$	-1.02 ± 0.33

The fit performed with the RooFit UEMML technique, with the “unbiased” model that allows for $\Gamma_{0\nu\beta\beta} \leq 0$, is plotted in Fig. 10.2. Fig. 10.3a shows the equivalent plot of the corresponding BAT MCMC fit, with the curve on top of a histogram of the data with asymmetrical Poisson errors. The adjacent Fig. 10.3b illustrates the same result, but with the error on the fit curve (yellow 68% band). We calculate the error from the spread of the numerous iterations of the MCMC (drawn in black, downsampled $\sim 1000\times$ to reduce computing time). The parameters are summarized in Table 10.2. We observe excellent agreement between the two different methodologies, as shown in Fig 10.4. Specifically, all parameter values match to two specific digits or more and are well within each other’s 1σ error intervals. Furthermore, though RooFit quotes asymmetrical Frequentist confidence intervals while BAT uses a symmetrical definition, the error bars’ values are quite close. The symmetrical BAT uncertainties always lie in between RooFit’s smaller and larger bounds. We understand this to be a good validation of the appropriateness of our model to the data we observe, in spite of relatively few events in the ROI.

CUORE’s other RooFit-based ROI fit has results fully consistent with ours [140]. This is not surprising, since said fit employs an essentially equivalent model PDF and the same software package and statistical approach. That said, this is a very good confirmation of the technical correctness of our analysis. The agreement between our two fits and the one featured in [33] and [140] is illustrated by Fig. 10.4. We observe an interesting phenomenon




Figure 10.2: ROI RooFit UEML fit result. The blue vertical band shows $Q_{\beta\beta} = (2527.515 \pm 0.013)$ keV. $R_{\beta\beta}$ is allowed be negative. Pulls (residuals in units of σ) indicating the agreement between the fit and the data, are plotted above




Figure 10.3: ROI BAT MCMC fit result, allowing negative values of $R_{\beta\beta}$, with (a) Poisson error bars are shown on the data and pulls (residuals in units of σ) plotted above; (b) error on the fit determined from a the (down-sampled) MCMC.

with regards to the peak shapes in our two fits: with very close fit parameters and integrals, the lines in the BAT fit are slightly more prominent, yet narrow. This is not expected, since the peak shape for both is taken from the results discussed in Chapter 9, and is fully constrained. This behavior merits further investigation, but as of this writing we believe it emerges from slight differences between the analytic normalization we utilize in the BAT PDF construction and the automated numerical normalization implemented in RooFit’s RooAbsPdf class [145].




Figure 10.4: Fit agreement: comparing our BAT MCMC fit (blue) and the “official” result from [33] with our RooFit UEML fit (red in the bottom frame).

Table 10.3: ROI fit results, constraining the $0\nu\beta\beta$ decay rate parameter to positive values. Note that $\Gamma_{0\nu\beta\beta}$ is the *best fit rate*, not a measurement or limit.

Parameter	RooFit value	BAT value
$R_{\text{bkg}}(\text{D.S. 1})$ (cts./keV/kg/y)	$(1.35^{+0.20}_{-0.19}) \cdot 10^{-2}$	$(1.35 \pm 0.19) \cdot 10^{-2}$
$R_{\text{bkg}}(\text{D.S. 2})$ (cts./keV/kg/y)	$(1.49^{+0.18}_{-0.17}) \cdot 10^{-2}$	$(1.49 \pm 0.18) \cdot 10^{-2}$
R_{Co} (cts./kg/y)	$(2.5^{+0.9}_{-0.8}) \cdot 10^{-1}$	$(2.5 \pm 0.8) \cdot 10^{-1}$
q	1.0003 ± 0.0005	1.0003 ± 0.0005
$R_{\beta\beta}$ (cts./kg/y)	$0.00^{+0.02}_{-0.00}$	0.00 ± 0.02
$\Gamma_{0\nu\beta\beta}$ (10^{-25} y^{-1})	$0.00^{+0.12}_{-0.000}$	0.00 ± 0.13




Figure 10.5: ROI RooFit UEML fit result. The blue vertical band shows $Q_{\beta\beta} = (2527.515 \pm 0.013)$ keV. $R_{\beta\beta}$ is constrained to positive values. Pulls (residuals in units of σ) indicating the agreement between the fit and the data, are plotted above




Figure 10.6: ROI BAT MCMC fit result, constraining $R_{\beta\beta}$ to positive values, with (a) Poisson error bars are shown on the data and pulls (residuals in units of σ) plotted above; (b) error on the fit determined from the (down-sampled) MCMC.

The “physical” model fits, wherein we constrain $\Gamma_{0\nu\beta\beta}$ to positive values, are drawn in Figs. 10.5 (RooFit UEML fit) and 10.6 (BAT MCMC). Fig. 10.6a shows the BAT fit compared to the data with asymmetrical Poisson error bars, while Fig. 10.6b depicts the statistical uncertainty on the fit as determined from the downsampled MCMC output. The fit parameters are given in Table 10.3. As before, we see agreement to two or more significant figures in all parameters, and congruent uncertainties on these values. Compared to the “unbiased” model, background rates $R_{\text{bkg}}(\text{D.S. 1})$ and $R_{\text{bkg}}(\text{D.S. 2})$ are lower. This is expected, since the statistical fluctuation near $Q_{\beta\beta}$ is pulling them down. The ^{60}Co 2-photon event rates are higher to compensate for this. In light of this, it is the author’s opinion that the fit permitting negative $0\nu\beta\beta$ decay rate values is a slightly more appropriate choice. Because it is unaffected by the overall rates, the ^{60}Co quenching factor is identical.

Chapter 11

CUORE Analysis: Limits on the $0\nu\beta\beta$ Half-Life of ^{130}Te

11.1 Introduction

A common occurrence in rare event searches is a lack of discovery of the process of interest. Standard practice under such circumstances is to set an exclusion limit on some portion of the relevant parameter space. In other words, we make a probabilistic statement that our experiment’s sensitivity is sufficient to have observed the behavior, with some level of confidence, were it to exist within a range of physical bounds. Outside those constraints, observation by our detector is not likely and the process cannot be ruled out.

In this specific case, we place a new lower limit on the half-life of the $0\nu\beta\beta$ -decay channel of ^{130}Te , $T_{1/2}^{0\nu}$. Indeed, CUORE’s first run consisting of 86.3 kg·y TeO_2 or 24 kg·y ^{130}Te exposure in two datasets taken over the summer of 2017, does not observe any evidence of this exotic nuclear process. As we discuss in Chapter 10, there is actually a rather prominent downward statistical fluctuation in that region. However, this result *does* represent the strongest limit to-date for this isotope.

In this chapter, we examine four different limit setting techniques and their results. We set Bayesian and Frequentist limits using two alternative treatments of a scan of the profile negative log-likelihood (NLL) from our RooFit UEML fit (Sections 10.3 and 10.5). We set a “fully Bayesian” limit with our BAT analysis (Sections 10.4 and 10.5) by marginalizing. Finally, in the interest of robustness and comparing “apples to apples,” we also generate a profile NLL, and a Bayesian limit therefrom, with our BAT fit. We observe very good agreement between the three Bayesian fits, while the frequentist result can be understood in the light of that value’s different physical meaning.

11.2 Profile NLL Curves from the RooFit Result

Scanning the RooFit Likelihood

To set limits using the result of our RooFit Unbinned Extended Maximum Likelihood (UEML) fit, we must first construct a profile NLL curve. The NLL function, given in Eq. 10.7, depends on all five of the fit parameters: two background rates R_{bkg} , the ^{60}Co 2- γ peak quenching factor q and rate R_{Co} , and the $0\nu\beta\beta$ -decay rate $R_{\beta\beta}$. Only the last of these, which corresponds directly to the decay rate through the neutrinoless channel $\Gamma_{0\nu\beta\beta}$ (scaled by isotope mass, Eq. 10.4), interests us in the context of limit setting. We fix the remaining four “nuisance” parameters to their best fit values, and then numerically evaluating the likelihood function for various values of $\Gamma_{0\nu\beta\beta}$.

In the case of the fit where the rate is allowed to take on negative values (Fig. 10.2 and Table 10.2), we scan between $-1.5 \cdot 10^{-25}$ and $+10.0 \cdot 10^{-25} \text{ y}^{-1}$ in steps equal to $1/10$ of the best fit rate’s uncertainty. This covers a region both below and above the NLL curve’s minimum at $(-1.01^{+0.38}_{-0.29}) \cdot 10^{-25} \text{ y}^{-1}$. The domain is asymmetrical, going much farther in the positive direction, because it is the physical, positive decay rate values that are important in the upcoming limit calculations. For the scenario with a rate constrained to positive values, we instead start the scan from zero. The resulting profile NLL curves are plotted in Figs. 11.1 and 11.2.




Figure 11.1: The CUORE profile NLL curve, with and without systematic uncertainties, if $\Gamma_{0\nu\beta\beta}$ is allowed to take on negative values in the ROI fit.




Figure 11.2: The CUORE profile NLL curve, if $\Gamma_{0\nu\beta\beta}$ is constrained to positive values only in the ROI fit.

Including Systematic Uncertainties

When fitting the model to the data, we intentionally consider only uncorrelated, or statistical, uncertainties. To not over-estimate the limit, we must also include the systematic error. This is done at the level of the NLL curve. CUORE systematics from five sources—energy resolution, calibration, spectral line shape, background shape, and detector efficiency—were evaluated from a toy Monte Carlo study. All of these uncertainties have absolute value component, while some also have a contribution defined as a fraction of the $0\nu\beta\beta$ lifetime. The numbers are published in [33] and are summarized in Table 11.1.

Table 11.1: Systematic uncertainties in the CUORE $0\nu\beta\beta$ decay analysis [33].

Source	Absolute syst. uncertainty (y^{-25})	Relative syst. uncertainty (%)
Energy resolution	—	1.5
Energy reconstruction	—	0.2
Line shape	0.02	2.4
Background shape	0.05	0.8
Efficiency	—	1.8

We can evaluate the combined systematics χ^2 as such:

$$\chi_{\text{syst.}}^2 = \left(\frac{\Gamma_{0\nu\beta\beta} - \Gamma_{0\nu\beta\beta}^{\text{min.}}}{\sigma_{\text{syst}}} \right)^2, \quad (11.1)$$

where

$$\sigma_{\text{syst.}}(\Gamma_{0\nu\beta\beta}) = \sum_{i=1}^5 \left(\sigma_i^{\text{abs.}} + \Gamma_{0\nu\beta\beta} \sigma_i^{\text{rel.}} \right). \quad (11.2)$$

Wilks' theorem states that for sufficiently high sample sizes, assuming Gaussian statistical uncertainties, the value $-2 \times \text{NLL}$ is one-dimensional χ^2 -distributed [146]. With just 155 events in the CUORE ROI, such an approximation is not ideal, but we consider it to be sufficiently close for error analysis. This approach allows us to mathematically understand the NLL curve in terms of the χ^2 statistic as follows:

$$\chi_{\text{stat.}}^2 = -2(\text{NLL} - \text{NLL}_{\text{min.}}), \quad (11.3)$$

where we can shift the curve such that $\text{NLL}_{\text{min.}} = 0$. Now, we can combine our likelihood scan with the systematic $\chi_{\text{syst.}}^2$ as defined by Eq. 11.1, and get out an NLL curve taking into account systematic uncertainty [147]:

$$\text{NLL}_{\text{stat.}+\text{syst.}} = -\frac{1}{2}\chi_{\text{stat.}+\text{syst.}}^2 = -\frac{1}{2} \left(\frac{1}{\chi_{\text{stat.}}^2} + \frac{1}{\chi_{\text{syst.}}^2} \right)^{-1}. \quad (11.4)$$

Our CUORE negative log-likelihood curve with systematics included is shown, together with the original statistics-only scan, in Fig. 11.1. The limit-setting technique, described below in Section 11.3, is the same regardless of which curve we choose to use.

We do not perform an equivalent calculation with the curve in Fig. 11.2, because our chosen approach relies on the existence of a local minimum in the χ^2 plot. Thanks to the negative statistical fluctuation in our data, the analysis constraining $\Gamma_{0\nu\beta\beta}$ to positive values only has a trivial minimum in the profile NLL, at the edge at zero. The curve with systematics cannot be negative, and is forced to zero at that point as well, which is an unphysical artifact that would affect the limit we calculate. As such, we will work only with the “unbiased” version of our RooFit-based profile likelihood analysis when dealing with systematics. We are contented with this, because statistically we find the two versions to be fully consistent with each other (see below in Section 11.3).

Combining CUORE with Earlier Experiments

CUORE, the latest in a series of TeO_2 bolometric detectors, is both the largest and most radiopure device thus far. Its backgrounds are the lowest, and the $24 \text{ kg} \cdot \text{y}$ ^{130}Te exposure from the first two datasets is the largest taken up to this point. Nonetheless, we can strengthen our limit on the $0\nu\beta\beta$ decay half-life of this isotope by combining our measurement with earlier ones.

The Cuoricino and CUORE-0 experiments served as test beds, demonstrators, and prototypes for CUORE. The former had an exposure of 19.75 kg·y, while the latter—9.8 kg·y. Thanks to lower background in CUORE-0 [76, 79], the two detectors had similar sensitivities. Prior to the publication of [33], the final CUORE-0 result combined with that of Cuoricino represented the strongest limit on this value [76, 79]. Continuing the tradition of [138], we also include the measurement from the Three Tower Test (TTT). This was a detector technology investigation into passive background reduction techniques comprising a total of 12 CUORE-style bolometers operated prior to the assembly of CUORE-0, in 2009 – 2010. As a technical study not intended to actually make a competitive $0\nu\beta\beta$ decay measurement, TTT has the smallest exposure of just 1.2 kg·y [58, 122, 138]. Nonetheless, its contribution is non-negligible.

Overall, the measurement and analysis techniques of these smaller TeO_2 bolometers are very similar to CUORE. Both feature an UEML fit to the ^{130}Te $0\nu\beta\beta$ decay region of interest analogous to the one we describe in Chapter 10, from which NLL curves were extracted to set limits. These curves are directly compatible with ours as described herein, and we can combine the CUORE likelihood scan with them with straight forward addition:

$$\text{NLL}_{\text{total}}(\Gamma_{0\nu\beta\beta}) = \text{NLL}_{\text{Q}}(\Gamma_{0\nu\beta\beta}) + \text{NLL}_{\text{Q0}}(\Gamma_{0\nu\beta\beta}) + \text{NLL}_{\text{Qino}}(\Gamma_{0\nu\beta\beta}) + \text{NLL}_{\text{TTT}}(\Gamma_{0\nu\beta\beta}) . \quad (11.5)$$

The “Q” subscripts in Eq. 11.5 refer to the conventional abbreviations of the different detectors in the CUORE lineage. The final result is shown in Fig. 11.3.

11.3 Bayesian and Frequentist Limits from the Profile NLL

The profile NLL curves in Figs. 11.1 – 11.3 show the negative natural logarithm of the likelihood of the various values of $\Gamma_{0\nu\beta\beta}$ given the data. Understanding this in terms of Bayesian inference, we interpret said likelihood as a probability of the data being in its observed state given the value of $\Gamma_{0\nu\beta\beta}$. Treating it as such, we can integrate the likelihood to set a limit on the rate.

We begin by trivially converting the NLL scan to an actual profile likelihood:

$$\mathcal{L}_{\text{profile}}(\Gamma_{0\nu\beta\beta}) = e^{-1 \cdot \text{NLL}(\Gamma_{0\nu\beta\beta})} . \quad (11.6)$$

We take a positive decay rate as our prior—just like in the fully-Bayesian analysis with BAT discussed in Section 10.4 and below—and naturally do not consider non-physical negative values of $\Gamma_{0\nu\beta\beta}$. We integrate the function in Eq. 11.6 from 0 to infinity to find the full area under the curve, which corresponds to 100% probability, and normalize it. Technically this is not possible, because our original NLL scan does not go above $1.0 \cdot 10^{-24} \text{y}^{-1}$. However, at this point the likelihood asymptotes to zero and the result is sufficiently close. Next, we integrate again, but only as far as the rate that gives us the probability corresponding to the




Figure 11.3: The profile NLL curves from the CUORE, CUORE-0, Cuoricino, and TTT experiments, together with their combination. Systematic uncertainties are included in all curves.

credible interval (CI) that we seek. In keeping with convention for $0\nu\beta\beta$ decay searches, we quote 90% CI limits. The rate value Γ_{LIM} , in the integral upper bound, is in fact the rate limit, as shown in Eq. 11.7.

$$\frac{\int_0^{\Gamma_{\text{LIM}}} \mathcal{L}_{\text{profile}}(\Gamma_{0\nu\beta\beta}) d\Gamma_{0\nu\beta\beta}}{\int_0^{\infty} \mathcal{L}_{\text{profile}}(\Gamma_{0\nu\beta\beta}) d\Gamma_{0\nu\beta\beta}} = 0.90 . \quad (11.7)$$

This technique is illustrated, using CUORE result with systematic uncertainties from Fig. 11.1, in Fig. 11.4. Numerically, we perform our integration using a rectangular approximation with a sampling frequency higher than that used to perform the original NLL scan.

The profile likelihood function can also be understood from a frequentist perspective, directly as a measure of the plausibility of the $\Gamma_{0\nu\beta\beta}$ decay parameter value given the data. With this approach, it is appropriate to treat the profile NLL as being equivalent to $-\frac{1}{2}\chi^2$ —in the way that we discussed in Section 11.2 when adding systematic uncertainties. As before, this is an approximation, because our event rate in the ROI is not very high and the assumption that statistical uncertainties are Gaussian is imperfect. The p -value of the χ^2 distribution can be used to find a confidence interval. Starting from the profile NLL minimum, which is the best fit $\Gamma_{0\nu\beta\beta}$ value, we move up the curve to the right. The decay rate at which profile NLL increases by the 10th percentile of the χ^2 distribution ($\chi^2 = 2.7056$) represents a limit on $\Gamma_{0\nu\beta\beta}$ at a 90% confidence level (CL). This calculation is commonly known as a “Rolke limit,” and combines the “ $\ln\mathcal{L} + \frac{1}{2}$ ” method with the profile likelihood




Figure 11.4: An illustration of integrating the profile likelihood of $\Gamma_{0\nu\beta\beta}$ (CUORE, stat. + syst.) to set a 90% CI Bayesian limit.

as described in [148]. The physical meaning of such a limit is that, were we to perform a very large number of equivalent experiments, 90% of these hypothetical tests would cover the (unknown) true value. Note that this interpretation is different from the Bayesian limit described earlier, where the *credible* interval represents a one-sided range within which the true $0\nu\beta\beta$ decay rate falls with a 90% probability.

In our analysis, we calculate both the Bayesian and Rolke limits from four different profile NLL curves:

1. The CUORE result, with the fit constrained to $\Gamma_{0\nu\beta\beta} \geq 0 \text{ y}^{-1}$ —Fig. 11.2.
2. The CUORE result with an “unbiased fit,” allowing $\Gamma_{0\nu\beta\beta}$ to take on negative values 11.1.
3. Version (2), with the addition of systematic uncertainties—Fig. 11.1.
4. The combined likelihoods of CUORE (“unbiased” version, with systematics), CUORE-0, Cuoricino, and TTT—Fig. 11.3.

We convert decay rate limits into the more informative half-life equivalents. The results are summarized in Table 11.2.

There are good arguments—“unbiased” versus “physical”—for and against allowing the decay rate to take on negative values in our fit. Not considering systematic uncertainties, however, we see that the NLL curve in Fig. 11.1, for rate values greater than zero, is nearly

Table 11.2: 90% CI/CL limits on the $0\nu\beta\beta$ half-life of ^{130}Te from RooFit profiled NLL.

	Bayesian lim. (10^{25} y)	Rolke lim. (10^{25} y)
CUORE (stat. only, $\Gamma_{0\nu\beta\beta} \geq 0$)	1.38	2.14
CUORE (stat. only)	1.38	2.14
CUORE (stat. + syst.)	1.35	2.11
CUORE + CUORE-0 + Cuoricino + TTT	1.53	2.33

identical to that in Fig. 11.2. We have excellent agreement—to four significant digits—between both Bayesian and Frequentist limits from the two approaches. Thanks to this congruence, we can safely choose just one of the curves when accounting for systematics. Naturally, when we do so the limits are weaker than their equivalents that include statistical errors only. The limit augmented with the results of previous experiments is noticeably stronger. There is an apparent discrepancy between the Frequentist and Bayesian limits. In fact, this is not a physical disagreement, and stems from the different statistical meaning of these two values as earlier described.

11.4 Bayesian Limits with BAT

In a fully Bayesian analysis like this one, the most natural way to set a limit is through marginalization. As we discuss in detail in Section 10.4 of the previous chapter, BAT (the Bayesian Analysis Toolkit) [141] uses a Markov Chain Monte Carlo (MCMC) method to update the model parameters given the ROI data, sampling the log-likelihood as stated in Eq. 10.9. Once $\ln\mathcal{L}(\boldsymbol{\theta}, \Gamma_{0\nu\beta\beta} | \text{Data})^1$ is finalized, we extract the posterior distribution of $\Gamma_{0\nu\beta\beta}$ by integrating over the remaining free parameters $\boldsymbol{\theta}$ and normalizing the result. Such distributions for our two fit scenarios—with the decay rate free to take on negative values or constrained to the physical range—are shown in Figs. 11.5 and 11.6, respectively.

We find the distance of the mode of this distribution from zero—the null hypothesis for the $0\nu\beta\beta$ decay rate. Were this value significantly positive—a signal from this rare decay channel—we would use it to claim evidence or discovery. It is not, so instead we use the mode and distribution to set a limit. In the case with $\Gamma_{0\nu\beta\beta}$, this is very straight forward. As we can see, the mode is at zero, fully consistent with the null hypothesis. We integrate the normalized posterior distribution from the mode up to the rate value that would give us a probability of 0.9. This rate is the 90% CI marginalized limit on $\Gamma_{0\nu\beta\beta}$, which we then convert to the ^{130}Te $0\nu\beta\beta$ decay half-life. Thanks to the downward statistical fluctuation at the Q-value in the ROI spectrum, if we allow the decay rate parameter to take on negative values

¹We have separated out the parameter of interest, $\Gamma_{0\nu\beta\beta}$. $\boldsymbol{\theta}$ now refers to the remaining free parameters—background rates, ^{60}Co rate, and ^{60}Co quenching.




Figure 11.5: $\Gamma_{0\nu\beta\beta}$ posterior distribution from marginalization, with the fit parameter allowed to take on negative values.

the mode of the posterior distribution is $-1 \cdot 10^{-25} \text{ y}^{-1}$. Knowing that negative decay rates are not physical, we do not adjust our null hypothesis downwards. As such, instead of integrating up from the mode, we re-normalize just the positive portion of the distribution and integrate from zero. The final result is very similar to that from the positively-constrained fit.

In the interest of robustness and comparison with our RooFit-based analysis, we also set a profiled limit. We get a profile curve by taking the posterior log-likelihood, fixing all parameters except $\Gamma_{0\nu\beta\beta}$ to their best fit values, and scanning in the same way we did with the RooFit result in Section 11.2. To get a profile NLL, we take the negative of this curve and shift it downward such that it has a zero value at the minimum at $-1.016 \cdot 10^{-25} \text{ y}^{-1}$, for the fit where the rate is permitted to go negative. For the positive-only fit, we vertically align the NLL with the negative one. The two results are plotted in Fig. 11.7. Far away from the minimum, where the statistics in the MCMC are lower, the curves begin to get noisy. As a result, the profiling result from our BAT analysis is limited in its precision and relegated to a cross-check role. Away from the tails, the two profile NLLs appear to agree well with each other and with the corresponding RooFit result. We integrate the curves in the same way as described in Section 11.2 to set a Bayesian limit.

Our 90% CI Bayesian limits on the $0\nu\beta\beta$ decay half-life of ^{130}Te from the BAT analysis are given in Table 11.3. We have carried out marginalization and profiling calculations for




Figure 11.6: $\Gamma_{0\nu\beta\beta}$ posterior distribution from marginalization, with the fit parameter allowed constrained to the positive, physical range.

both versions of our fit. All four results are consistent with each other to two significant digits, though the profiled values are both somewhat higher than their marginalized counterparts. The $1.42 \cdot 10^{25}$ y number is particularly outlying. This is a consequence of the noise we see in the NLL curves, particularly in the case of the scan that covers negative rates. Marginalization is more appropriate and reliable with this fitting approach, and so we quote $T_{1/2}^{0\nu} \geq 1.38 \cdot 10^{25}$ y (stat. only) as the result of the BAT analysis.

Table 11.3: Marginalized and profiled 90% CI limits on the $0\nu\beta\beta$ decay half-life of ^{130}Te from the BAT analysis.

	Marginalized limit (10^{25} y)	Profiled limit (10^{25} y)
CUORE (stat. only, $\Gamma_{0\nu\beta\beta} \geq 0$)	1.38	1.39
CUORE (stat. only)	1.38	1.42




Figure 11.7: Profile NLL curves derived from the posterior log-likelihoods of the two BAT fits—with and without constraining $\Gamma_{0\nu\beta\beta}$ to negative values. The RooFit profile NLL from Fig. 11.1 (stat. only) is plotted for comparison.

11.5 Conclusions and Comparison with Published Result

We evaluate the Bayesian lower limit on the $0\nu\beta\beta$ decay half-life of ^{130}Te from CUORE data with statistical uncertainty with a profile NLL in a RooFit-based analysis, and through marginalization in a fully-Bayesian analysis with BAT. As described in Sections 11.3 and 11.4, both calculations were done with two models: one that allows for negative values of $\Gamma_{0\nu\beta\beta}$, and the other where the rate is constrained to physical, positive values. All four of these results agree within three significant digits, and are consistent with

$$T_{1/2}^{0\nu} \geq 1.4 \cdot 10^{25} \text{ y (90\% CI, stat. only).} \quad (11.8)$$

Factoring in systematic uncertainties in our RooFit-based profiled calculations, this is reduced to

$$T_{1/2}^{0\nu} \geq 1.3 \cdot 10^{25} \text{ y (90\% CI, stat. + syst.).} \quad (11.9)$$

The Rolke-type Frequentist limits corresponding to Eqs. 11.8 and 11.9, reported to two significant figures are both consistent with

$$T_{1/2}^{0\nu} \geq 2.1 \cdot 10^{25} \text{ y (90\% CL).} \quad (11.10)$$

When we combine CUORE’s first results with those of its predecessor experiments—CUORE-0, Cuoricino, and TTT—we get:

$$T_{1/2}^{0\nu} \geq 1.5 \cdot 10^{25} \text{ y (90\% CI, Bayesian)} \quad (11.11)$$

$$T_{1/2}^{0\nu} \geq 2.3 \cdot 10^{25} \text{ y (90\% CL, Rolke).} \quad (11.12)$$

The values in Eqs. 11.11 and 11.12 represent, on par with the analysis in [140], the strongest limit on the $0\nu\beta\beta$ decay half-life of ^{130}Te to date.

A total of three analyses determining limits on this quantity have been carried out with the data from CUORE datasets 1 and 2 during the preparation of the official result published in [33]. In addition to the two discussed in this thesis, J. Cushman conducted another RooFit-based investigation with a UEML fit and profile likelihood limits, which is an earlier version of that presented in [140]. All three analyses were developed and performed side by side, ensuring consistency between statistical techniques and code bases. The equivalent limits from all of the treatments matched up to within two significant digits. We report 90% CI Bayesian and Rolke limits of $T_{1/2}^{0\nu} \geq 1.5 \cdot 10^{25} \text{ y}$ and $T_{1/2}^{0\nu} \geq 2.2 \cdot 10^{25}$, respectively, from CUORE data combined with CUORE-0 and Cuoricino [33]. While [140] achieves an improved sensitivity over this result thanks to a refinement of the systematic uncertainties on efficiency, we do so via the inclusion of TTT measurements. In both cases, the change manifests on the level of the third significant digit of the Bayesian limit. The yellow-shaded band on the exclusion plot for effective Majorana mass $m_{\beta\beta}$ in Fig. 3.4 in Chapter 3 corresponds generally to a $1.5 \cdot 10^{25} \text{ y}$ half-life limit.

Chapter 12

Conclusions

This dissertation showcases my personal contributions to the CUORE experiment. Nonetheless, it presents a wide cross section of the many diverse tasks that we have had to complete in order to achieve its successful operation. The general approach of searching for the $0\nu\beta\beta$ decay of ^{130}Te with cryogenic TeO_2 bolometers has been thoroughly honed with the CUORE-0 and Cuoricino prototype detectors. Being nearly 20 times larger, however, CUORE necessitated a number of novel solutions to be implemented. The CUORE cryostat, in particular, is one of a kind, and is the largest and most powerful dilution refrigerator known at the time of this writing. In fact, CUORE is by far the biggest and most massive sub-Kelvin bolometer array, and is the world's first ton-scale instrument of this type. CUORE's size also presented a challenge during installation: to first order, its 19 towers would take 19 times as long to mount as CUORE-0, risking significant recontamination. Our radon-free installation environment represents a notable innovation that will likely see future use and development in the field of low-background physics.

Scientifically, the limit on the $0\nu\beta\beta$ decay half-life of ^{130}Te evaluated in this thesis— $T_{1/2}^{0\nu} \geq 1.5 \cdot 10^{25}$ y (90% CI, Bayesian), and $T_{1/2}^{0\nu} \geq 2.3 \cdot 10^{25}$ y (90% CL, Frequentist)—is the most stringent to date. It is evaluated from only two months of live time (86.3 kg·y TeO_2 or 24 kg·y ^{130}Te exposure). In the future, we plan to operate CUORE for a total live time of about five years. Besides significantly increasing statistics, we are aiming to improve detector performance (thanks to noise, energy resolution is still inferior to that of CUORE-0), as well as the analysis (including lowering energy thresholds). We hope to achieve our ^{130}Te half-life target sensitivity of $\sim 9 \cdot 10^{25}$ y with the full run. The corresponding sensitivity to the effective Majorana mass of the neutrino $m_{\beta\beta}$ is indicated by the blue-gray band in Fig. 12.1.

Thanks to its large isotope mass and good energy resolution, CUORE is a highly competitive $0\nu\beta\beta$ decay search. Nonetheless, as we can clearly see in Fig. 12.1, it will only slightly scrape the main portion of the parameter space for this process. The key to further improvement lies in the bolometric detectors themselves. The materials characterization discussed in Appendix A offers only a glimpse into ongoing bolometer R&D efforts. As I explain in Section 3.4 and Eq. 3.5, background plays a limiting role in experimental sensitivity, and the




Figure 12.1: The sensitivity of CUPID to the effective Majorana mass of the neutrino ($m_{\beta\beta}$) [149], compared to CUORE [33, 34] and state of the art experiments using other isotopes [35, 36, 37]. Nuclear matrix elements from [38, 39, 40, 41, 42, 43, 44, 45, 46]. The red and green bands represent the normal and inverted hierarchy parameter spaces, smeared due to the unknown value of the Majorana phases.

lack of active background rejection is CUORE’s biggest weakness. This can be addressed with an optical bolometer, which uses scintillation or Cherenkov light from the crystal to discriminate different types of events. Scintillating rare event bolometers are already being operated, but at scales much smaller than CUORE. CUPID—the CUORE Upgrade with Particle Identification—is a proposed ton-scale optical bolometer that will succeed CUORE in the same facility and cryostat [150]. CUPID will also feature enriched crystals, allowing us to increase isotope mass without scaling up the detector (and backgrounds from structural components). Running for about 10 years, CUPID would be sufficiently sensitive to cover the entirety of $0\nu\beta\beta$ decay parameter space in the event of an inverted neutrino mass hierarchy [149] (gray band in Fig. 12.1). R&D for CUPID, particularly its low-temperature light detectors, is currently being actively pursued at UC Berkeley, LBNL, and a number of collaborating institutions around the world.

Bibliography

- [1] E. Fermi. “Tentativo di una Teoria Dei Raggi β ”. In: *Il Nuovo Cimento* 11 (1934), p. 1.
- [2] C. L. Cowan Jr. et al. “Detection of the Free Neutrino: a Confirmation”. In: *Science* 124.3212 (1956), p. 103.
- [3] S. L. Glashow. “Partial-Symmetries of Weak Interactions”. In: *Nuclear Physics* 22.4 (1961), p. 579.
- [4] S. Weinberg. “A Model of Leptons”. In: *Physical Review Letters* 19.21 (1967), p. 1264.
- [5] A. Salam. “Elementary Particle Physics: Relativistic Groups and Analyticity”. In: *Eighth Nobel Symposium*. Ed. by N. Svartholm. Stockholm: Almquist and Wiksell, 1968, p. 367.
- [6] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Physical Review Letters* 13.16 (1964), p. 321.
- [7] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Physical Review Letters* 13.16 (1964), p. 508.
- [8] C. R. Hagen G. S. Guralnik and T. W. B. Kibble. “Global Conservation Laws and Massless Particles”. In: *Physical Review Letters* 13.20 (1964), p. 585.
- [9] user Marcokarlo. *File:Standard Model of Elementary Particles modified version.svg*. Wikimedia Commons. https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles_modified_version.svg. accessed 8 May 2018.
- [10] Б. М. Понтекорво. “Мезоний и антимезоний”. In: *Журнал Экспериментальной и Теоретической Физики* 33 (1957), p. 549.
- [11] B. Pontecorvo. “Mesonium and Antimesonium”. In: *Journal of Experimental and Theoretical Physics* 33 (1957), p. 429.
- [12] M. Nakagawa Z. Maki and S. Sakata. “Remarks on the Unified Model of Elementary Particles”. In: *Progress of Theoretical Physics* 28.5 (1962), p. 870.
- [13] B. Pontecorvo. “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”. In: *Journal of Experimental and Theoretical Physics* 26.5 (1967), p. 984.
- [14] С. П. Михеев и А. Ю. Смирнов. “Резонансное усиление осцилляций в веществе и спектроскопия солнечных нейтрино”. In: *Ядерная Физика* 42.6 (1985), p. 1441.

- [15] S. P. Mikheev and A. Yu. Smirnov. “Resonance Enhancement of Oscillations in Matter and Solar Neutrino Spectroscopy”. In: *Soviet Journal of Nuclear Physics* 42.6 (1985), p. 913.
- [16] L. Wolfenstein. “Neutrino Oscillations in Matter”. In: *Physical Review D* 17.9 (1978), p. 2369.
- [17] B. T. Cleveland et al. “Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector”. In: *The Astrophysical Journal* 496.1 (1998), p. 505.
- [18] R. Davis Jr. “Solar Neutrinos. II. Experimental”. In: *Physical Review Letters* 12.11 (1964), p. 303.
- [19] A. Gando et al. “Constraints on θ_{13} from a Three-Flavor Oscillation Analysis of Reactor Antineutrinos at KamLAND”. In: *Physical Review D* 83.5 (2011), p. 052002.
- [20] Q. R. Ahmad et al. “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory”. In: *Physical Review Letters* 87.7 (2001), p. 071301.
- [21] Q. R. Ahmad et al. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. In: *Physical Review Letters* 89.1 (2002), p. 011301.
- [22] K. Eguchi et al. “First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance”. In: *Physical Review Letters* 90.2 (2003), p. 021802.
- [23] S. F. King and C. Luhn. “Neutrino Mass and Mixing with Discrete Symmetry”. In: *Reports on Progress in Physics* 76.5 (2013), p. 056201.
- [24] E. Majorana. “Teoria Simmetrica dell’Elettrone e del Positrone”. In: *Il Nuovo Cimento* 14 (1937), p. 171.
- [25] M. Goeppert-Mayer. “Double Beta-Disintegration”. In: *Physical Review* 48.6 (1935), p. 512.
- [26] A. A. Hahn S. R. Elliott and M. K. Moe. “Direct Evidence for Two-Neutrino Double-Beta Decay in ^{82}Se ”. In: *Physical Review Letters* 59.18 (1987), p. 2020.
- [27] M. Wang et al. “The Ame2012 Atomic Mass Evaluation”. In: *Chinese Physics C* 36 (2012), p. 1603.
- [28] M. Redshaw et al. “Masses of ^{130}Te and ^{130}Xe and Double-Beta-Decay Q Value of ^{130}Te ”. In: *Physical Review Letters* 102.21 (2009), p. 212502.
- [29] D. Scielzo et al. “Double- β -Decay Q Values of ^{130}Te , ^{128}Te , and ^{120}Te ”. In: *Physical Review C* 80.2 (2009), p. 025501.
- [30] S. Rahaman et al. “Double-Beta Decay Q Values of ^{116}Cd and ^{130}Te ”. In: *Physics Letters B* 703.412 (2011), p. 395.
- [31] G. Racah. “Sulla Simmetria Tra Particelle e Antiparticelle”. In: *Il Nuovo Cimento* 14 (1937), p. 322.

- [32] W. H. Furry. “On Transition Probabilities in Double Beta-Disintegration”. In: *Physical Review* 56.12 (1939), p. 1184.
- [33] C. Alduino et al. “First Results from CUORE: A Search for Lepton Number Violation via $0\nu\beta\beta$ Decay of ^{130}Te ”. In: *Physical Review Letters* 120.13 (2018), p. 132501.
- [34] R. Ardito et al. (CUORE Collaboration). “CUORE: A Cryogenic Underground Observatory for Rare Events”. In: arXiv:hep-ex/0501010 (2005).
- [35] M. Agostini et al. “Improved Limit on Neutrinoless Double- β Decay of ^{76}Ge from GERDA Phase II”. In: *Physical Review Letters* 120.13 (2018), p. 132503.
- [36] R. Arnold et al. “Search for Neutrinoless Double-Beta Decay of ^{100}Mo With the NEMO-3 Detector”. In: *Physical Review D* 89.11 (2014), 111101(R).
- [37] A. Gando et al. “Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen”. In: *Physical Review Letters* 117.8 (2016), p. 082503.
- [38] S. T. Petcov A. Meroni and F. Šimkovic. “Multiple CP Non-Conserving Mechanisms of $(\beta\beta)0\nu$ -Decay and Nuclei with Largely Different Nuclear Matrix Elements”. In: *Journal of High Energy Physics* 2013 (2013), p. 25.
- [39] J. Menéndez et al. “Disassembling the Nuclear Matrix Elements of the Neutrinoless $\beta\beta$ Decay”. In: *Nuclear Physics A* 818.3-4 (2009), p. 139.
- [40] F. Šimkovic et al. “ $0\nu\beta\beta$ and $2\nu\beta\beta$ Nuclear Matrix Elements, Quasiparticle Random-Phase Approximation, and Isospin Symmetry Restoration”. In: *Physical Review C* 87.4 (2013), p. 045501.
- [41] M. T. Mustonen and J. Engel. “Large-Scale Calculations of the Double- β Decay of ^{76}Ge , ^{130}Te , ^{136}Xe , and ^{150}Nd in the Deformed Self-Consistent Skyrme Quasiparticle Random-Phase Approximation”. In: *Physical Review C* 87.6 (2014), p. 064302.
- [42] J. Kotila J. Barea and F. Iachello. “ $0\nu\beta\beta$ and $2\nu\beta\beta$ Nuclear Matrix Elements in the Interacting Boson Model with Isospin Restoration”. In: *Physical Review C* 91.3 (2015), p. 034304.
- [43] J. Hyvärinen and J. Suhonen. “Nuclear Matrix Elements for $0\nu\beta\beta$ Decays with Light or Heavy Majorana-Neutrino Exchange”. In: *Physical Review C* 91.2 (2015), p. 024613.
- [44] A. Neacsu and M. Horoi. “Shell Model Studies of the ^{130}Te Neutrinoless Double- β Decay”. In: *Physical Review C* 91.2 (2015), p. 024309.
- [45] T. R. Rodríguez and G. Martínez-Pinedo. “Energy Density Functional Study of Nuclear Matrix Elements for Neutrinoless $\beta\beta$ Decay”. In: *Physical Review Letters* 105.25 (2010), p. 252503.
- [46] T. R. Rodríguez N. López Vaquero and J. L. Egido. “Shape and Pairing Fluctuation Effects on Neutrinoless Double Beta Decay Nuclear Matrix Elements”. In: *Physical Review Letters* 111.14 (2013), p. 142501.

- [47] А. Д. Сахаров. “Нарушение СР-инвариантности, С-асимметрия и барионная асимметрия Вселенной”. In: *Письма в Журнал Экспериментальной и Теоретической Физики* 5.1 (1966), p. 32.
- [48] A. D. Sakharov. “Violation of CP Invariance, C Asymmetry, and Baryon Asymmetry of the Universe”. In: *Soviet Physics Uspekhi* 34.5 (1991), p. 392.
- [49] E. W. Kolb and M. S. Turner. *The Early Universe*. Chicago, IL: Westview Press, 1990.
- [50] P. A. R. Ade et al. “Planck 2015 Results. XIII. Cosmological Parameters”. In: *Astronomy and Astrophysics* 594 (2016), A13.
- [51] A. S. Barabash. “Average and Recommended Half-Life Values for Two-Neutrino Double Beta Decay”. In: *Nuclear Physics A* 935 (2015), p. 52.
- [52] J. Ouellet. “The Coldest Cubic Meter in the Known Universe”. In: arXiv:1410.1460 [physics.ins-det] (2014).
- [53] C. Kittel. *Introduction to Solid State Physics*. John Wiley & Sons, 2005.
- [54] C. Arnaboldi et al. “Production of High Purity TeO₂ Single Crystals for the Study of Neutrinoless Double Beta Decay”. In: *Journal of Crystal Growth* 312.20 (2010), p. 2999.
- [55] F. Alessandria et al. “CUORE Crystal Validation Runs: Results on Radioactive Contamination and Extrapolation to CUORE Background”. In: *Astroparticle Physics* 35.12 (2012), p. 839.
- [56] M. Barucci et al. “Measurement of Low Temperature Specific Heat of Crystalline TeO₂ for the Optimization of Bolometric Detectors”. In: *Journal of Low Temperature Physics* 123.5-6 (2001), p. 303.
- [57] C. Alduino et al. (CUORE Collaboration). “CUORE-0 Detector: Design, Construction and Operation”. In: *Journal of Instrumentation* 11 (2016), P07009.
- [58] A. D. Bryant. “A Search for Neutrinoless Double Beta Decay of ¹³⁰Te”. PhD thesis. Berkeley, CA: University of California, Berkeley, Spring 2010.
- [59] E. E. Haller et al. “NTD Germanium: A Novel Material for Low Temperature Bolometers”. In: *Neutron Transmutation Doping of Semiconductor Materials*. Ed. by R. D. Larrabee. Boston, MA: Springer, 1984, p. 21.
- [60] K. M. Itoh E. E. Haller and J. W. Beeman. “Neutron Transmutation Doped (NTD) Germanium Thermistors for Submillimetre Bolometer Applications”. In: *Submillimetre and Far-Infrared Space Instrumentation, Proceedings of the 30th ESLAB Symposium held in Noordwijk, 24-26 September 1996*. Ed. by E. J. Rolfe and G. Pilbratt. Paris: European Space Agency, 1996, p. 115.
- [61] *Araldite Rapid*. Araldite Professional Adhesives. <http://www.go-araldite.com/products/epoxy-adhesives/araldite-rapid-2-x-15ml-tube>. accessed 10 May 2018.

- [62] C. Arnaboldi et al. “A Programmable Multichannel Antialiasing Filter for the CUORE Experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A* 617.1-3 (2010), p. 327.
- [63] P. Carniti et al. “A Low Noise and High Precision Linear Power Supply with Thermal Foldback Protection”. In: *Review of Scientific Instruments* 87 (2016), p. 054706.
- [64] C. Arnaboldi et al. “Very Low Noise AC/DC Power Supply Systems for Large Detector Arrays”. In: *Review of Scientific Instruments* 86 (2015), p. 124703.
- [65] C. Arnaboldi et al. “A Front-End Electronic System for Large Arrays of Bolometers”. In: *Journal of Instrumentation* 13 (2018), P02026.
- [66] A. Giachero. “Characterization of Cryogenic Bolometers and Data Acquisition System for the CUORE Experiment”. PhD thesis. Genoa: Università degli Studi di Genova, 2008.
- [67] S. Di Domizio. “Search for Double Beta Decay to Excited States with CUORICINO and Data Acquisition System for CUORE”. PhD thesis. Genoa: Università degli Studi di Genova, 2009.
- [68] S. Copello. “Dark Matter Induced Annual Modulation Analysis in CUORE-0”. PhD thesis. Genoa: Università degli Studi di Genova, 2017.
- [69] M. Vignati. “Model of the Response Function of CUORE Bolometers”. PhD thesis. Rome: Sapienza Università di Roma, Dec. 2009.
- [70] A. Drobizhev et al. “Thermal Conductivity Measurements of PTFE and Al_2O_3 Ceramic at Sub-Kelvin Temperatures”. In: *Cryogenics* 85 (2017), p. 63.
- [71] V. Singh et al. “Specific Heat of Teflon, Torlon 4203 and Torlon 4301 in the Range of 30–400 mK”. In: *Cryogenics* 67 (2015), p. 15.
- [72] K. Alfonso et al. “A High Precision Pulse Generation and Stabilization System for Bolometric Experiments”. In: *Journal of Instrumentation* 13 (2018), P02029.
- [73] D. R. Artusa et al. (CUORE Collaboration). “Initial Performance of the CUORE-0 Experiment”. In: *The European Physics Journal C* 74 (2014), p. 2956.
- [74] C. Alduino et al. (CUORE Collaboration). “The Projected Background for the CUORE Experiment”. In: *The European Physics Journal C* 77 (2017), p. 543.
- [75] D. R. Artusa et al. (CUORE Collaboration). “Searching for Neutrinoless Double-Beta Decay of ^{130}Te with CUORE”. In: *Advances in High Energy Physics* 2015 (2015), p. 879871.
- [76] K. Alfonso et al. (CUORE Collaboration). “Search for Neutrinoless Double-Beta Decay of ^{130}Te with CUORE-0”. In: *Physical Review Letters* 115.10 (2015), p. 102502.
- [77] E. Andreotti et al. (Cuoricino Collaboration). “ ^{130}Te Neutrinoless Double-Beta Decay with CUORICINO”. In: *Astroparticle Physics* 34.11 (2011), p. 822.

- [78] C. Arnaboldi et al. (Cuoricino Collaboration). “Results from a Search for the $0\nu\beta\beta$ -Decay of ^{130}Te ”. In: *Physical Review C* 78.3 (2008), p. 035502.
- [79] C. Alduino et al. (CUORE Collaboration). “Analysis Techniques for the Evaluation of the Neutrinoless Double- β Decay Lifetime in ^{130}Te with CUORE-0”. In: *Physical Review C* 93.4 (2016), p. 045503.
- [80] C. Alduino et al. (CUORE Collaboration). “Measurement of the Two-Neutrino Double Beta Decay Half-life of ^{130}Te with the CUORE-0 Experiment”. In: *The European Physics Journal C* 77 (2017), p. 13.
- [81] G. Benato et al. “Radon Mitigation During the Installation of the CUORE $0\nu\beta\beta$ Decay Detector”. In: *Journal of Instrumentation* 13 (2018), P01010.
- [82] C. Rusconi. “Optimization of the Bolometric Performances of the CUORE-0/CUORE and LUCIFER Detectors for the Neutrinoless Double Beta Decay Search”. PhD thesis. Como: Università degli Studi dell’Insubria, July 2011.
- [83] E. Bucchieri et al. “An Assembly Line for the Construction of Ultra-Radio-Pure Detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A* 768 (2014), p. 130.
- [84] *ABB Robotics: Industrial Robots*. ASEA Brown Boveri. <http://new.abb.com/products/robotics/industrial-robots>. accessed 2 May 2018.
- [85] J. S. Cushman et al. “The detector calibration system for the CUORE cryogenic bolometer array”. In: *Nuclear Instruments and Methods in Physics Research Section A* 844 (2017), p. 32.
- [86] F. Orio S. Di Domizio and M. Vignati. “Lowering the Energy Threshold of Large-Mass Bolometric Detectors”. In: *Journal of Instrumentation* 6 (2011), P02007.
- [87] C. Alduino et al. (CUORE collaboration). “Low Energy Analysis Techniques for CUORE”. In: *The European Physics Journal C* 77 (2017), p. 857.
- [88] E. Gatti and P. F. Manfredi. “Processing the Signals from Solid-State Detectors in Elementary-Particle Physics”. In: *La Rivista del Nuovo Cimento* 9.1 (1986), p. 1.
- [89] P. C. Mahalanobis. “On the Generalized Distance in Statistics”. In: *Proceedings of the National Institute of Science of India* 2.1 (1936), p. 49.
- [90] G. Cowan et al. “Asymptotic Formulae for Likelihood-Based Tests of New Physics”. In: *The European Physical Journal C* 71 (2011), p. 1554.
- [91] G. Ventura and L. Risegari. *The Art of Cryogenics: Low-Temperature Experimental Techniques*. Oxford: Elsevier, 2008.
- [92] F. Pobell. *Matter and Methods at Low Temperatures*. Berlin: Springer-Verlag, 2007, p. 149.
- [93] *Cryorefrigerator Specification Sheet: PT415-RM with CP1110*. http://www.cryomech.com/specificationsheet/PT415RM_ss.pdf. Cryomech, Inc., 2011; accessed 26 March 2018.

- [94] V. Singh et al. “The CUORE cryostat: Commissioning and Performance”. In: *Journal of Physics: Conference Series* 718.6 (2016), p. 062054.
- [95] C. Ligi et al. “The CUORE Cryostat: A 1-Ton Scale Setup for Bolometric Detectors”. In: *Journal of Low Temperature Physics* 184.3-4 (2016), p. 590.
- [96] *Aurubis Shapes*. Aurubis AG. <https://www.aurubis.com/en/products/page-shapes>. accessed 7 May 2018.
- [97] C. E. Chase and G. O. Zimmerman. “Measurement of P-V-T Relations and Critical Indices of ^3He ”. In: *Journal of Low Temperature Physics* 11.5-6 (1973), p. 551.
- [98] *Teragon’s Summary of Cryogen Properties*. <http://www.trgn.com/database/cryogen.html>. San Francisco, CA: Teragon Research, 2013; accessed 26 March 2018.
- [99] A. A. Tarasov E. I. Mikulin and M. P. Shkrebyonock. “Low-Temperature Expansion Pulse Tubes”. In: *Advances in Cryogenic Engineering*. Ed. by R. W. Fast. Vol. 29. Boston, MA: Springer, 1984, p. 629.
- [100] R. Radebaugh. “Cryocoolers: the State of the Art and Recent Developments”. In: *Journal of Physics: Condensed Matter* 21.16 (2009), p. 164219.
- [101] A. Th. A. M. de Waele and J. G. M. Kuerten. “Thermodynamics and Hydrodynamics of ^3He - ^4He Mixtures”. In: *Progress in Low Temperature Physics*. Ed. by D. F. Brewer. Vol. 13. Amsterdam: Elsevier, 1992, p. 167.
- [102] *PT415 Cryorefrigerator with Remote Motor Option: Cold Head Outline Drawing*. http://www.cryomech.com/coldhead/PT415RM_ch.pdf. Cryomech, Inc., 2007; accessed 26 March 2018.
- [103] A. D’Addabbo et al. “An Active Noise Cancellation Technique for the CUORE Pulse Tube Cryocoolers”. In: arXiv:1712.02753 [physics.ins-det] (2017).
- [104] *LNx Series Motor and Drives*. <http://www.premotn.com/ln2xmnl.pdf>. San Jose, CA: Precision Motion Controls, accessed 26 March 2018.
- [105] *Leiden Cryogenics: Leader in Low Temperature Techniques*. <http://www.leiden-cryogenics.com/company>. accessed 26 March 2018.
- [106] *Alcatel/Adixen ATH-3200M Turbo Pump*. <https://www.provac.com/products/alcatel-adixen-ath-3200m-turbo-pump>. accessed 26 March 2018.
- [107] C. Bucci. Personal Communication. October 2017.
- [108] “How to Support the Coldest Cubic Meter in the Universe with Vibration Isolators”. In: *AzoNano* (2010; accessed 26 March 2018). <https://www.azonano.com/article.aspx?ArticleID=4675>.
- [109] *Negative-Stiffness Vibration Isolators: How they Work*. https://www.minusk.com/content/technology/how-it-works_passive_vibration_isolator.html. Minus K, Inc., accessed 26 March 2018.

- [110] E. Fiorini. “2000 Years-Old Roman Lead for Physics”. In: *AStroParticle ERA net* (2010; accessed 26 March 2018). http://212.71.251.65/aspera/images/stories/news/PAPER_VERSIONS/asperanewsletter0610.pdf.
- [111] C. Enss and D. McCammon. “Physical Principles of Low Temperature Detectors: Ultimate Performance Limits and Current Detector Capabilities”. In: *Journal of Low Temperature Physics* 151.1-2 (2008), p. 5.
- [112] *Introduction: PID Controller Design*. University of Michigan. <http://ctms.engin.umich.edu/CTMS/index.php?example=Introduction§ion=ControlPID>. accessed 26 March 2018.
- [113] J. Bechhoefer. “Feedback for Physicists: A Tutorial Essay on Control”. In: *Reviews of Modern Physics* 77.3 (2005), p. 783.
- [114] J. Engert et al. “Noise Thermometry at Low Temperatures: MFFT Measurements Between 1.6 K and 1 mK”. In: *Journal of Physics: Conference Series* 400.5 (2012), p. 052003.
- [115] A. Kumar N. Samkharadze and G. A. Cáthy. “A New Type of Carbon Resistance Thermometer with Excellent Thermal Contact at Millikelvin Temperatures”. In: *Journal of Low Temperature Physics* 160.5-6 (2010), p. 246.
- [116] *Device Specifications: NI 6321*. <http://www.ni.com/pdf/manuals/374461b.pdf>. National Instruments Corporation, 2015; accessed 26 March 2018.
- [117] *Keithley Precision DC and AC+DC Low Noise Current Sources*. Tektronix, Inc. <https://www.tek.com/keithley-low-level-sensitive-and-specialty-instruments/keithley-ultra-sensitive-current-sources-seri>. accessed 26 March 2018.
- [118] *Connecting Instruments via GPIB*. <https://www.ni.com/getting-started/setup-hardware/instrument-control/gpib-connect>. National Instruments Corporation, accessed 26 March 2018.
- [119] *Integral (Reset) Windup, Jacketing Logic and the Velocity PI Form*. Control Guru: Practical Process Control. <https://controlguru.com/integral-reset-windup-jacketing-logic-and-the-velocity-pi-form/>. 2015; accessed 26 March 2018.
- [120] M. T. Tham. *Discretised PID Controllers: Part of a set of study notes on Digital Control by M. Tham*. University of Newcastle Upon Tyne, Chemical Engineering and Advanced Materials, 2014.
- [121] *AVS-47B AC Resistance Bridge*. RV-Elektroniikka Oy PICOWATT. <http://www.picowatt.fi/avs47b/avs47b.html>. 2018; accessed 26 March 2018.
- [122] F. Alessandria et al. (CUORE Collaboration). “Validation of Techniques to Mitigate Copper Surface Contamination in CUORE”. In: *Astroparticle Physics* 45 (2013), p. 13.
- [123] M. Clemenza et al. “Radon-Induced Surface Contaminations in Low Background Experiments”. In: *European Physics Journal C* 71 (2011), p. 1805.

- [124] *HEPA-ULPA Filters*. American Air Filter International. <https://www.aafintl.com/en/commercial/browse-products/commercial/hepa-ulpa-filters>. accessed 7 May 2018.
- [125] *Hospi-Gard IsoClean and IsoClean with Ultraviolet Light*. Envirco. https://www.envirco-hvac.com/products/index.aspx?prod=HospiGard_IsoClean. accessed 7 May 2018.
- [126] *Winch 5 Item No. 240 User Manual*. http://wahlberg.dk/media/240.805.008_Manual_Winch5.pdf. Wahlberg Motion Design, 2017; accessed 26 March 2018.
- [127] *Winch 10 Item No. 241 User Manual*. http://wahlberg.dk/media/241.805.012_Manual_Winch10.pdf. Wahlberg Motion Design, 2017; accessed 26 March 2018.
- [128] *DMX 101: A DMX 512 Handbook*. <http://cdb.s3.amazonaws.com/ItemRelatedFiles/10191/dmx-101-handbook.pdf>. Elation Professional, 2008; accessed 26 March 2018.
- [129] *ANSI E1.27-1-2006 (R2016): Entertainment Technology Standard for Portable Control Cables for Use with ANSI E1.11 (DMX512-A) and USITT DMX512/1990 Products*. CP/2003-1028r5.2. http://tsp.esta.org/tsp/documents/docs/E1-27-1_2006R2016.pdf. Entertainment Services and Technology Association, 2016; accessed 26 March 2018.
- [130] *MagicQ*. ChamSys. <https://secure.chamsys.co.uk/magicq>. accessed 7 May 2018.
- [131] C. Alduino et al. (CUORE Collaboration). “Search for Neutrinoless β^+ EC Decay of ^{120}Te with CUORE-0”. In: *Physical Review C* 97.5 (2018), p. 055502.
- [132] W. Verkerke and D. Kirkby. “The RooFit Toolkit for Data Modeling”. In: arXiv:physics/0306116 [physics.data-an] (2003). Talk from the 2003 Computing in High Energy Physics (CHEP03), La Jolla, CA, USA, March 2003.
- [133] *RooFit Users Manual v2.91*. https://root.cern.ch/download/doc/RooFit_Users_Manual_2.91-33.pdf. CERN, 2008; accessed 26 March 2018.
- [134] G. F. Knoll. *Radiation Detection and Measurement*. New York, NY: John Wiley & sons, Inc., 2000.
- [135] M. H. Crosthwait. *Review of the Gamma Spectrum Structure*. <http://www.people.vcu.edu/~mhcrosthwait/clrs322/ReviewoftheGammaSpectrum.htm>. Virginia Commonwealth University, accessed 2 April 2018.
- [136] J. B. Kortright and A. C. Thompson. “1.2 X-Ray Emission Energies”. In: *X-Ray Data Booklet*. Ed. by A. C. Thompson et al. <http://cxro.lbl.gov/x-ray-data-booklet>. Berkeley, CA: Lawrence Berkeley National Laboratory, 2009, pp. 1–8.
- [137] M. A. Carrettoni. “Data Analysis for Neutrinoless Double Beta Decay”. PhD thesis. Milan: Università degli Studi di Milano Bicocca, Mar. 2011.
- [138] J. L. Ouellet. “The Search for $0\nu\beta\beta$ Decay in ^{130}Te with CUORE-0”. PhD thesis. Berkeley, CA: University of California, Berkeley, Spring 2015.

- [139] L. P. Ekström S. Y. F. Chu and R. B. Firestone. *208Tl: Table of Isotopes Decay Data*. <http://nucldata.nuclear.lu.se/toi/nuclide.asp?iZA=810208>. 1997; accessed 2 April 2018.
- [140] J. S. Cushman. “A Search for Neutrinoless Double-Beta Decay in Tellurium-130 with CUORE”. PhD thesis. New Haven, CT: Yale University, 2018.
- [141] D. Kollar A. Caldwell and K. Kröninger. “BAT - The Bayesian Analysis Toolkit”. In: *Computer Physics Communications* 180 (2009), p. 2197.
- [142] S. Karlin and H. M. Taylor. *A First Course in Stochastic Processes*. Cambridge, MA: Academic Press, 2012.
- [143] S. Richardson W. R. Gilks and D. Spiegelhalter (eds.) *Markov Chain Monte Carlo in Practice*. London: Chapman and Hall, 1996.
- [144] N. Metropolis et al. “Equation of State Calculations by Fast Computing Machines”. In: *The Journal of Chemical Physics* 21.6 (1953), p. 1087.
- [145] *RooAbsPdf Class Reference*. <https://root.cern.ch/doc/master/classRooAbsPdf.html>. CERN, 2018; accessed 9 April 2018.
- [146] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. In: *The Annals of Mathematical Statistics* 9.1 (1938), p. 60.
- [147] B. H. Hooberman. “Two Complementary Strategies for New Physics Searches at Lepton Colliders”. PhD thesis. Berkeley, CA: University of California, Berkeley, Spring 2009.
- [148] A. M. Lopez W. A. Rolke and J. Conrad. “Limits and Confidence Intervals in the Presence of Nuisance Parameters”. In: *Nuclear Instruments and Methods in Physics Research Section A* 551.2-3 (2005), p. 493.
- [149] D. R. Artusa et al. “Exploring the Neutrinoless Double Beta Decay in the Inverted Neutrino Hierarchy with Bolometric Detectors”. In: *The European Physics Journal C* 74 (2014), p. 3096.
- [150] The CUPID Interest Group. “CUPID: CUORE (Cryogenic Underground Observatory for Rare Events) Upgrade with Particle Identification”. In: arXiv:1504.03599 [physics.ins-det] (2015). CUPID white paper.
- [151] T. Scott and M. Giles. “Dislocation Scattering in Teflon at Low Temperatures”. In: *Physical Review Letters* 29.10 (1972), pp. 642–643.
- [152] J. P. Le E. D. Marquardt and R. Radebaugh. “Cryogenic Materials Properties Database”. In: *Cryocoolers 11*. Ed. by R. G. Ross. Boston, MA: Springer-Verlag, 2002.
- [153] R. Berman. “The Thermal Conductivity of some Polycrystalline Solids at Low Temperatures”. In: *Proceedings of the Physical Society, Section A* 65.12 (1952), p. 1029.
- [154] G. Deutscher S. Alterovitz and M. Gershenson. “Heat Capacity and Thermal Conductivity of Sintered Al_2O_3 at Low Temperatures by the Heat Pulse Technique”. In: *Journal of Applied Physics* 46 (1975), p. 3637.

- [155] *Raw Materials*. McMaster-Carr. <https://www.mcmaster.com/#raw-materials/=14xkupj>. accessed 7 November 2016.
- [156] E. T. Swartz and R. O. Pohl. “Thermal Boundary Resistance”. In: *Reviews of Modern Physics* 61.3 (1989), p. 605.
- [157] I. Peroni et al. “Thermal Conductivity of Manganin Below 1 K”. In: *Nuclear Physics B - Proceedings Supplements* 78.1–3 (1999), p. 573.
- [158] J. R. Olson. “Thermal Conductivity of Some Common Cryostat Materials Between 0.05 and 2 K”. In: *Cryogenics* 33.7 (1993), p. 729.
- [159] J. Cheeke. “The Kapitza Resistance and Heat Transfer at Low Temperatures”. In: *Journal de Physique Colloques* 31.C3 (1970), p. 129.
- [160] M. A. Gubrud. “Scanning Tunneling Microscopy at MilliKelvin Temperatures: Design and Construction”. PhD thesis. College Park, MD: University of Maryland, 2010.
- [161] W. Reese A. C. Anderson and J. C. Wheatley. “Thermal Conductivity of Some Amorphous Dielectric Solids Below 1 °K”. In: *Review of Scientific Instruments* 34 (1963), p. 1386.
- [162] B. I. Halperin P. W. Anderson and C. M. Varma. “Anomalous Low-Temperature Thermal Properties of Glasses and Spin Glasses”. In: *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics* 25.1 (1972), p. 1.
- [163] W. A. Phillips. “Tunneling States in Amorphous Solids”. In: *Journal of Low Temperature Physics* 7.3-4 (1972), p. 351.
- [164] Y. S. Touloukian et al. *Thermophysical Properties of Matter vol. 2: Thermal Conductivity, Nonmetallic Solids*. New York, NY: Plenum Press, 1970.
- [165] D. M. Finlayson and P. J. Mason. “Structure Scattering and the Thermal Conductivity of a Semicrystalline Polymer”. In: *Journal of Physics C: Solid State Physics* 18.9 (1984), p. 1791.
- [166] G. J. Morgan and D. Smith. “Thermal Conduction in Glasses and Polymers at Low Temperatures”. In: *Journal of Physics C: Solid State Physics* 7.4 (1973), p. 649.
- [167] R. B. Stephens. “Low-Temperature Specific Heat and Thermal Conductivity of Non-crystalline Dielectric Solids”. In: *Physical Review B* 8.6 (1973), p. 2896.
- [168] *Heat Transfer*. <https://www.mathworks.com/help/pde/heat-transfer-and-diffusion-equations.html>. MathWorks, accessed 16 May 2017.
- [169] L. P. Ekström S. Y. F. Chu and R. B. Firestone. *Summary Drawing for A=130*. <http://nucleardata.nuclear.lu.se/toi/sumframe.htm>. 1997; accessed 2 May 2018.

Appendix A

Thermal Conductivity Measurements of PTFE and Al_2O_3 Ceramic at Sub-Kelvin Temperatures

This appendix chapter is derived from [81].

A.1 Introduction

The operation and results of CUORE and its predecessors, including those presented in this thesis, evidences the great potential of bolometric detectors for rare event searches. However, due to the very low operating temperatures of these devices (10–100 mK for state of the art examples), their design and operation also present unique technical challenges. Among these are the characterization and selection of appropriate structural materials.

As discussed in Sections 4.2 and 4.3 of Chapter 4, the bolometer’s absorber is coupled by a weak link to a heat sink. The weak link’s thermal conductivity $K(T)$, needs to be low, so that event heat energy is not immediately dissipated to the bath and to cause a temperature rise in the absorber that is high and long enough to make a good measurement. These considerations, combined with wide availability, easy machining, and good radiopurity motivates the widespread use of PTFE— $(\text{C}_2\text{F}_4)_n$, a soft plastic commonly known by the trademark name Teflon—in many instruments, including CUORE and Cuoricino [57, 78]. Given its popularity in cryogenics, low temperature thermal properties of PTFE have been studied extensively. The thermal conductivity has been measured down to 170 mK, revealing an approximately quadratic temperature dependence that is consistent with theoretical models of amorphous solids [151]. However, the parameter that truly optimizes a weak link is thermal diffusivity,

$$\alpha(T) = K(T)/C(T), \tag{A.1}$$

through which heat capacity $C(T)$ also comes into play. While we want the weak link to have high impedance to heat dissipation from the absorber, it should not store heat itself. Such

storage causes the undesirable effects of lengthening the bolometer pulses when the stored heat diffuses back into the absorber and sensor, and decreases their amplitude when a large portion of the heat dissipates directly to the bath. PTFE’s high heat capacity compared to typical bolometric absorber materials [152] makes it less than optimal in this role, especially since its low structural strength necessitates relatively large parts. Efforts at constructing a thermal model of the CUORE bolometers in the form of a numerically evaluated system of classical differential equations reveal the strong effect PTFE’s low diffusivity has on pulse shape [69].

Alumina interests us, because it has low heat capacity and strong mechanical rigidity. It is prepared by sintering (compacting under high pressure and temperature) corundum/sapphire crystallites. Use of alumina in cryogenic experiments is currently less extensive than PTFE, as is the characterization of its thermal properties at ultra-low temperatures. Measurements of its thermal conductivity exist down to ~ 1 K and higher. These exhibit approximately quadratic or cubic temperature dependencies, as well as wide sample-to-sample variation in absolute value due to manufacturing differences [153, 154]. We measure the thermal conductivity of alumina in the ~ 0.1 – 1.3 K range, and PTFE in the 0.17 – 0.43 K range. Our PTFE measurements are carried out at significantly finer temperature increments than currently available results, useful for our efforts at bolometer thermal modeling, as well as validating our experimental technique for the more poorly understood alumina.

A.2 Experimental Setup

Technique

We perform our measurements using the “integrated thermal conductivity method” [91]. This involves applying known powers P_{applied} on one end of the sample and measuring that end’s temperature T , while maintaining the other end at a known constant base temperature T_0 .

The total thermal power $P(T)$ applied to the sample is equal to the integral of the thermal conductivity of the sample over the measured temperature gradient multiplied by a geometric factor

$$P(T) = A/l \cdot \int_{T_0}^T K(T') dT', \quad (\text{A.2})$$

where A is the cross-sectional area of the sample, l is its length, and $K(T)$ is thermal conductivity. Notice that maintaining T_0 constant makes the thermal power a function of the upper limit of the integral, simplifying analysis.

Ideally, there would be no gradient at all across the sample while the heater is switched off. However, in practice we often observe a parasitic heat load. Sources can include leakage current from the heater power supply, thermal radiation from the sample’s surroundings, vibrational noise, and RF noise/pickup. Making the assumption that this unwanted contribution is constant, we include it as an additional parasitic power, $P_{\text{parasitic}}$:

$$P(T) = P_{\text{applied}}(T) + P_{\text{parasitic}}. \quad (\text{A.3})$$

The presence of a constant offset does not affect our analysis, since we extract thermal conductivity by differentiation, which negates the effect of a constant parasitic heat load:

$$K(T) = l/A \cdot \frac{dP(T)}{dT} = l/A \cdot \frac{dP_{\text{applied}}(T)}{dT}. \quad (\text{A.4})$$

We perform the above calculation by fitting our $P_{\text{applied}}(T)$ data with a power law plus a negative constant offset for $P_{\text{parasitic}}$ (separately for each sample and dataset), and differentiating the result analytically. Uncertainties are propagated through.

Hardware

The low temperature conditions are achieved using an Oxford Instruments Triton 400 dilution refrigerator. This cryostat combines a standard ^3He - ^4He dilution unit with a “cryogen-free” pulse tube cooler setup to reach temperatures as low as 6 mK. Nominal cooling power at 100 mK is 400 μW . In practice, for these measurements temperatures are in the hundreds of millikelvin due to the heat loads involved.

To create the temperature gradients needed to measure thermal conductivity, one end of an elongated sample is thermalized to the mixing chamber plate (coldest stage) of the cryostat, while the other end hangs freely and is equipped with a heater. A pair of thermometers measure the temperatures of both ends.

Four 3” (7.6 cm) long alumina samples are measured: a pair of cylindrical rods 0.25” (6.35 mm) in diameter (henceforth referred to as samples A and B), a prism with a 0.25” \times 0.25” (6.35 \times 6.35 mm²) square cross section (sample C), and a thicker cylinder 0.75” (1.9 cm) in diameter (sample D). All of these samples are “Nonporous Alumina Ceramic Rods” acquired from McMaster Carr, with a quoted porosity of 0% and purity of 96–99.8% Al_2O_3 [155]. The first PTFE sample (sample 1) used is a 2.75” (7.0 cm) long rectangular prism with a 0.25” \times 0.25” (6.35 \times 6.35 mm²) square cross section. The second (sample 2) was cut from the same brick of PTFE as the first, with the same cross-sectional area and about half the length. The PTFE was also acquired from McMaster Carr, under the name “Ultra-Clean PTFE” [155].

The thermometers used for all of these measurements are Lakeshore Cryotronics RX-102B-CB RuO_2 devices (numbers 2 and 9 in Figs. A.1a – A.2). These temperature sensors are sensitive within our entire measurement range. We calibrate them in separate cooldowns, in reference to a factory-calibrated RuO_2 thermometer and a ^{60}Co nuclear orientation thermometer. The factory thermometer mounts are made from copper, and are bolted to surfaces. Readout of the RuO_2 thermometers is performed with a Lakeshore AC370 resistance bridge. This instrument also includes a PID (Proportional, Integral, Differential) controller, which is used with the factory mixing chamber plate heater and thermometer, also a RuO_2 (not seen in figures), to stabilize the temperature of the cold bath.

We use two heater types. In our first measurements, with alumina samples A-C, these are Lakeshore Cryotronics 50 Ω cartridge heaters (wound NiCr wire, no. 10 in Fig. A.1a). We drive these with a Sorensen LS 18-5 DC voltage source. Given the low resistance of the

cartridge heaters, a room temperature load resistor is added in series. The resulting setup suffers from significant unwanted heating from leakage current, as can be seen in the large parasitic offset of those three datasets in Fig. A.5a. For the thicker alumina sample D and PTFE, we upgrade to 50 k Ω (alumina) and 1 k Ω (PTFE) NiCr thin film heaters (no. 10 in Figs. A.1b and A.2). These are in the form of chips varnished to a small Kapton flexible PCB, with gold wire bonds forming electrical connections to a pair of solder pads. For biasing, we switch to the built-in DC voltage source of a Keithley 6517B electrometer for alumina, and a Keithley 6220 precision DC current source for PTFE. Given the higher resistance of the thin film heaters, no additional load is used. This arrangement provides significantly lower leakage and parasitic heating, as manifested by the smaller offsets in Figs. A.3a and A.5a.

Three different variations of sample holder design are used for alumina samples A-C, alumina sample D, and the PTFE (Figs. A.1a – A.2). In the case of alumina, the ends of the sample are firmly held in vise-like copper clamps with cutouts in the jaws in the shape and size of the sample cross section. The top clamp, for the cold end, is flat with a large surface area for secure mounting and thermalization in the cryostat (no. 4 in Figs. A.1a and A.1b). For samples A-C, the bottom clamp is brick-like, covering the sample from three sides. Below the sample is a 0.25" (6.35 mm) diameter channel perpendicular to the sample, exactly matching the cartridge heater's dimensions, in which it is embedded with Dow Corning high vacuum grease impregnated with copper powder for improved thermal coupling (Fig. A.1a, nos. 8 and 10). The RuO_2 thermometer is attached to the clamp (Fig. A.1a, no. 9). For sample D, the bottom clamp is also a flat design more similar to the top holder (Fig. A.1b, no. 8). The thermometer (no. 9) is bolted to its top surface, and the thin film heater (no. 10) is adhered with GE varnish. Since the heater now directly contacts only one of two jaws, a flexible sheet of copper (no. 11) joins the two jaws of the clamp to ensure even thermalization. The top clamps are generally similar, though for the thick sample the cold end RuO_2 sensor (no. 2 in Figs. A.1a and A.1b) is moved directly to the clamp from its original position nearby. In all cases, samples penetrate the clamps to a depth of 0.25" (6.35 mm). The length of the sample is defined as the distance between the clamp edges, and does not include the embedded portions of the rod.

We find the clamp design to be inadequate for PTFE—differential thermal contraction, combined with the material's soft and slippery texture, causes the holders to separate from the sample during cooling, even when clamps are upgraded with flexible leaf springs. Instead, we drill holes in the samples and use brass threaded rods to make attachments, without thermal grease (Fig. A.2). On the warm end, besides the RuO_2 and the thin film heater on a small copper rectangle on opposite sides of the sample, we attach two copper bricks on the remaining two sides. These are both connected to the heater with copper wire for more even heat distribution. The cold end setup is analogous, except the heater is replaced by an attachment to a T-shaped thermalized base. Of all the elements, the two thermometers are mounted closest to the center of the sample. The length of the PTFE sample is defined by the distance between these. Mounting the thermometer separately from the heater has an additional advantage in reducing difficult-to-characterize parasitic heat load and holder-on-sample thermal boundary resistance [156].




Figure A.1: (a) Experimental setup for the 0.25" diameter cylindrical and 0.25" \times 0.25" square alumina samples. 1—Cu plate, thermalized to mixing chamber; 2—cold end sensor; 3—electrical connections; 4—cold end Cu clamp; 5—sample; 6—thermalization bobbins; 7—Cu shield; 8—heated end Cu clamp; 9—heated end sensor; 10—cartridge heater (embedded) [70]. (b) Experimental setup for the 0.75" diameter cylindrical alumina sample. 1 – 9—same as (a); 10—thin film heater; 11—Cu sheet thermalizing clamp jaws [70].

Shunt Conductance, Thermal Boundary Resistance, and Parasitic Heat Load

Technically, we are not measuring the thermal conductance of the sample on its own. Instead, we measure the effective conductance of a system consisting of the sample connected in series with the copper-on-sample thermal boundary resistance, and in parallel with a “shunt” in the form of the thermometer and heater wiring. The thermometer connections are Manganin alloy, which has electrical resistance nearly independent of temperature. Its thermal conductivity is measured to be about $5 \cdot 10^{-3} - 8 \cdot 10^{-2} \text{ W/m/K}$ in the 0.1–1.0 K range [157]. Heaters are connected with superconducting NbTi wiring, which conducts heat less—about $1 \cdot 10^{-4} - 2 \cdot 10^{-2} \text{ W/m/K}$ in the same range [158]. All wires are 15 – 20 cm long and 36 AWG. Looking at these properties and at our final results (Figs. A.3b and A.5b), we see that the shunt conductance can be neglected.

Determining thermal boundary resistance between the holder and sample is more complicated—




Figure A.2:]

Experimental setup for the PTFE samples, featuring threaded rod attachment of sensors and heater. 1–Cu plate, thermalized to mixing chamber; 2–cold end sensor; 3–electrical connections; 4–Cu cold end base; 5–sample; 7–Cu shield; 9–heated end sensor; 10–thin film heater; 12–thermalization bricks (3 on each end, connected to base/heater by Cu wire).

Wire thermalization bobbins were omitted, as they were found to be unnecessary [70].

it depends on temperature, power applied across the boundary, and material-specific properties [156]. Lack of knowledge of these properties prevents us from doing a precise theoretical calculation, whereas a fully experimental characterization is out of the reach of this study. We use measured and calculated values for similar materials taken from [159] and [160] to make conservative estimates of this contribution. For alumina, we expect a maximum possible boundary temperature gradient of ~ 0.05 mK at 100 mK, where the total gradient is more than ~ 12 mK. At 1 K, the boundary contribution is no greater than ~ 0.5 mK, while total gradient is in the hundreds of millikelvin. With PTFE, we compute a boundary effect resulting in a maximum gradient of ~ 0.6 mK at 170 mK, where the total gradient is ~ 35 mK, and a boundary contribution of ~ 0.8 mK at 430 mK, where the total gradient is ~ 80 mK. This suggests that boundary effects are not significant, and we include their contributions in our systematics.

Besides intentional heating, we experience a parasitic heat load on our samples, the primary sources of which we believe to be leakage current in the sample heater, and a combination of thermal radiation and mechanical vibration. The former is particularly prominent

in the datasets of samples A, B, and C, which use the less optimal Sorensen power supply and low resistance cartridge heaters. We confirm its influence by fully unplugging that connection, which eliminates most of the offset. The switch to high resistance thin film heaters driven by Keithleys for the later measurements significantly reduces the effect. Radiative heating comes from black body radiation from the sample’s surroundings. Our cryostat’s innermost thermal shield is at the 600–700 mK level. To decrease the radiation heat load on the experimental setup, the hardware is mounted inside an additional copper shield thermalized to the cryostat’s mixing chamber plate (Figs. A.1a – A.2). That said, the shield is fabricated from OFHC (oxygen free, high conductivity) copper, which is less than ideal for cryogenic use due to high hydrogen content, and is of large mass. Furthermore, it is exposed directly to the still vessel’s radiation—our cryostat is unable to host a shield on the “100 mK” stage. Thus, equilibrium base (mixing chamber) temperature is on the order of 100 mK, and the resultant heat load dominates the parasitic power we observe in the runs instrumented with thin film heaters and Keithleys. A simple calculation with the Stefan-Boltzmann law is consistent with this. The reason the radiation results in a gradient on the sample is that, while impingement on all surfaces is more or less even, the low thermal conductivity of the sample means that any heat absorbed by it and its hot-end hardware is dissipated much more slowly than that absorbed by the bath. Vibrational heating contributes an unknown amount of heat that is “mixed in” with the radiative component. Given the consistency of our observed conditions with black body emission and absorption, such a mechanical contribution would be of comparable magnitude or less than radiation. Though we cannot decouple the effect of vibrational heating, it is not necessary to extract thermal conductivity, since the system experiences no time-varying impacts or stresses and it would likewise be approximately constant.

The leakage current from the supply is constant while it is on and connected, and the heaters’ resistances are found to be temperature-independent and stable within 2% or less. The temperature of the cold bath and thermal shield is kept stable by PID control, ensuring the radiative load is also constant. Thus, we expect the parasitic heat load to be constant in the temperature range of interest, which is confirmed experimentally.

Bath base temperatures as low as ~ 70 mK are achieved, but are not possible to maintain for higher heater powers, creating the main constraint for the lower end of our measurement range.

Experimental Uncertainties

In our power versus temperature data, we handle statistical and systematic uncertainties separately. For the former, we take the spread of our temperature readings at every nominal power setting. Each spread is actually entangled with correlated errors stemming from thermometer calibration and PID temperature control. Thus, we adjust them with a common scaling factor to get $\chi^2/\text{df} = 1$ during fitting and take covariances into account in

propagation. The systematic uncertainty on applied power derives primarily from heater ¹ and load resistances, and voltage/current uncertainties from the power supply. Systematic temperature errors come from the thermometer calibrations, as well as a contribution from thermal boundary resistance estimates. In Figs. A.3a and A.5a, the statistical uncertainties are indicated by horizontal error bars, while the shaded bands represent systematics. The bands are defined by combining temperature and power uncertainties (conservatively assuming 100% correlation), shifting the data points by the appropriate amount, and repeating the fit. As with statistics, covariances are included in propagation calculations. Error values quoted in the text and tables, as well as the shaded bands in Figs. A.3b and A.5b, represent the combination of statistical and systematic components.

A.3 Results and Discussion

PTFE

We measure the thermal conductivity of PTFE over four temperature ranges. The first sample is measured over three ranges within the interval of 0.18–0.43 K. The second sample is measured over a narrow temperature range of about 0.17–0.21 K. The lower bounds are determined by the fact that, at temperatures below these, the thermal gradients across the samples are large and unstable in time. We prioritize making high resolution measurement with small temperature increments over covering a wide temperature range, which has already been done in an overlapping domain [151, 161].

For each data set we perform a shifted power law fit of the form

$$P_{\text{applied}}(T) \cdot l/A = a \left(\frac{T}{1 \text{ K}} \right)^b + c. \quad (\text{A.5})$$

The data is plotted in Fig. A.3a with the fit results, uncertainties, and correlations given in Table A.3. The thermal conductivity obtained by differentiation is given in Fig. A.3b and Table A.1, with equations of the form

$$K(T) = \mu \left(\frac{T}{1 \text{ K}} \right)^\nu, \quad (\text{A.6})$$

where $\mu = ab$ and $\nu = b - 1$. We observe an approximately quadratic temperature dependence of $K(T)$, and find μ and ν to be strongly correlated. We attribute the $\sim 1\sigma$ discontinuity at 0.37 K in Fig. A.3b to the relatively narrow temperature interval over which we measure.

Since samples 1 and 2 are cut from the same brick of PTFE, they are considered equivalent. We may thus take a weighted average of the data presented in Table A.1, with each μ and ν weighted by the reciprocal square of their respective error. We yield the result:

$$K_w(T) = (5.0 \pm 0.8) \cdot 10^{-3} \left(\frac{T}{1 \text{ K}} \right)^{1.83 \pm 0.11} \text{ W/m/K}. \quad (\text{A.7})$$

¹Heaters are wired in a four-wire configuration and characterized in separate cooldowns.

Our findings are consistent with the tunneling model. This theory describes amorphous solid structure by making the assumption that each atom has access to two neighboring equilibrium sites, modeled as a double well potential, within which tunneling occurs [162, 163]. It is applied to broadly defined amorphous solids, including plastics, by Phillips [163]. His treatment focuses on the tunneling states' frequency-dependent relaxation time, $\Gamma(\omega)$, to recover the thermal conductivity temperature dependence:

$$K(T) \propto T^2. \quad (\text{A.8})$$

$\Gamma(\omega)$ is experimentally shown to account for tunneling states in PE (polyethylene, $(\text{C}_2\text{H}_4)_n$), which is similar to PTFE $((\text{C}_2\text{F}_4)_n)$ in molecular and physical structure [163]. Thus, we would expect analogous tunneling states—the structural units providing two equilibrium positions and the differences between hydrogen and fluorine determining the height of the potential barrier separating them. This makes the tunneling model a reasonable theory to account for the thermal conductivity of PTFE at cryogenic temperatures.

T. Scott and M. Giles measure the thermal conductivity of PTFE from 0.17–4 K and observe a strong T^2 dependence below 1.2 K [151]. Anderson et al. report a $T^{2.4}$ dependence in the range of 0.3–0.7 K [161]. We observe a temperature dependence, as well as absolute value, of thermal conductivity that agree with those reported by T. Scott and M. Giles within the entire measurement range of 0.17–0.43 K. In Eq. A.7, the scaling factor is about a 30% larger than that observed by A.C. Anderson et al. with the temperature power being about 25% smaller. These differences result in an absolute value of about a factor of 2.3 times larger than that found by these authors in the overlapping regime of 0.30–0.43 K. We may account for the disagreement through the variability in samples that occurs in the industrial manufacturing process of PTFE.

Sample	Temperature Range (K)	μ (W/m/K)	ν	corr[μ, ν]
1	0.18–0.26	$(4.1 \pm 1.1) \cdot 10^{-3}$	1.70 ± 0.17	0.991
1	0.27–0.36	$(6.2 \pm 1.5) \cdot 10^{-3}$	1.99 ± 0.20	0.988
1	0.37–0.43	$(7.4 \pm 2.8) \cdot 10^{-3}$	2.05 ± 0.41	0.995
2	0.17–0.21	$(5.0 \pm 1.8) \cdot 10^{-3}$	1.82 ± 0.22	0.994

Table A.1: Results for PTFE thermal conductivity: $K(T) = \mu(T/1 \text{ K})^\nu$ [70].

Alumina

We acquire four alumina datasets, corresponding to the four samples. For samples A–C, all of which are measured above 0.35 K, we perform a single fit using Eq. A.5 over the entire data range. The constant c accounts for the parasitic heat load. Eq. A.5 agrees well with the data. The fits can be seen in Fig. A.5a and are given in Table A.4, together with parameter




Figure A.3: **(a)** PTFE measurements: applied heater power scaled by sample length l and cross-sectional area A as a function of heated end temperature T . Error bars indicate statistical uncertainty. Bands represent systematics. Numbers in legend refer to samples [70]. **(b)** PTFE thermal conductivity determined by differentiation of scaled P_{applied} fit in Fig. A.3a. Shaded bands represent combined propagated statistical and systematic uncertainty. Numbers in legend refer to samples [70].

uncertainties, correlations, and fit domains. The relatively large negative constant offset, seen prominently in Fig. A.5a, stems mainly from leakage current from the heater power supply used for these measurements.

We measure alumina sample D over a temperature range that is both wider and colder (down to 0.1 K). A fit of the form of Eq. A.5 is a visibly poor representation of the data, as can be seen in Fig. A.4a. Examining the graph, there appears to be a kink or transition around 0.3 K, below which the slope is more shallow. Thus, we fix b in Eq. A.5 to the “crystal” value of 4 and add second power law term:

$$P_{\text{applied}}(T) \cdot l/A = a \left(\frac{T}{1 \text{ K}} \right)^4 + r \left(\frac{T}{1 \text{ K}} \right)^s + c. \quad (\text{A.9})$$

Fig. A.4b clearly demonstrates that Eq. A.9 agrees well with the data, giving a quadratic dependence in the second power law term.

This fit is given in Fig. A.5a and Table A.4 together with the results for other samples, with parameter correlations in Matrix A.2. Due to the much cleaner heater power supply, and consequently the lower parasitic heat load, the constant offset is much smaller in the sample D data than for the others. Differentiating all four fits gives alumina thermal conductivity results, plotted in Fig. A.5b and given in Table A.2 with parameter correlations in Matrix A.1. Eq. A.6 describes samples A-C, while sample D follows an equation of the form

$$K(T) = \mu \left(\frac{T}{1 \text{ K}} \right)^3 + \lambda \left(\frac{T}{1 \text{ K}} \right)^\gamma, \quad (\text{A.10})$$

where $\lambda = rs$ and $\gamma = s - 1$.




Figure A.4: **(a)** Alumina sample D data with single power law fit (Eq. A.5). Fractional differences between fit and measured values are plotted below the main graph [70]. **(b)** Alumina sample D data with a Structure Scattering Model double power law fit (Eq. A.9). Fractional differences between fit and measured values are plotted below the main graph. This fit demonstrates superior agreement compared to the Eq. A.5 fit (Fig. A.4a) [70].

Above 0.3 K, in all our data sets, we recover a close to cubic dependence of thermal conductivity on temperature, which is consistent with crystalline/crystallite structure. The thermal conductivity of crystalline insulators at temperatures below a few kelvin is dominated by lattice vibrations, with phonons being the carriers of heat. In this regime, the phonon mean free path, l , is dependent on the dimensions of the solid in question [53]. Under the Debye approximation, so too is the phonon frequency ω , hence the phonon dispersion relation gives $l(\omega) = \xi\omega^{-n}$, with ξ and n denoting constants. The form of $l(\omega)$ makes our expression for thermal conductivity at low temperatures

$$K(T) = \frac{3\xi N k_B T^{3-n}}{2\pi\theta_D^3} \int_0^\infty \frac{x^{4-n} e^x}{(e^x - 1)^2} dx, \quad (\text{A.11})$$

where $x = \hbar\omega/k_B T$, N is the atomic number density, and θ_D the Debye temperature [164]. Thus,

$$K(T) \propto T^{3-n}. \quad (\text{A.12})$$

Touloukian explains that imperfections can lead to a wide array values for n that are case-specific [164]. Thus, considering the variability between alumina samples that results from the manufacturing process, we are not surprised to observe the aforementioned range of temperature dependencies above 0.3 K.

The behavior seen in Sample D is consistent with the Structure Scattering Model [165]. This involves the process of elastic scattering of phonons in semicrystalline solids due to imperfections and defects that exist on the microscopic level [166]. Scatterings can be both long and short range. The presence of long range correlations alone yields a standard “crystalline” T^3 dependence of thermal conductivity on temperature, while short range correlations yield




Figure A.5: **(a)** Alumina measurements: applied heater power scaled by sample length l and cross sectional area A as a function of heated end temperature. Error bars indicate statistical uncertainty. Bands represent systematics. Letters in legend refer to samples [70]. **(b)** Alumina thermal conductivity determined by differentiation of scaled P_{applied} fits in Fig. A.5a. Shaded bands represent combined propagated statistical and systematic uncertainty. Letters in legend refer to samples [70].

an approximately linear term. These correlation ranges may coexist within semicrystalline solids, resulting in an additive effect [165, 166]:

$$K(T) = c_1 T + c_3 T^3, \quad (\text{A.13})$$

for some constants c_1 and c_3 . In our fit, we fix the higher power value to four after noting that higher temperature measurements follow the Debye model well. Doing so reduces the number of degrees of freedom by one, which is advantageous given that the data set contains only eighteen points. We find that allowing the higher power value to float results in far larger errors on the fit parameters.

The literature exhibits variability in measured thermal conductivities of alumina. Berman reports the thermal conductivity of sintered Al_2O_3 to be proportional to $T^{2.8}$ at 2 K, attributing such behavior to the scattering of phonons off crystallite boundaries [153].

Alterovitz et al. measure two sintered Al_2O_3 samples yielding $T^{1.25}$ and $T^{1.4}$ dependencies between 10 and 40 K. They suspect that their samples were partially amorphous in light of the $T^{\leq 2}$ dependence of other amorphous materials measured by R.B. Stephens [154, 167].

Comparing our data to these results, we note that power dependencies are not inconsistent with the linear and cubic phases that we see, though varied and not in precisely the same temperature regimes. In the overlapping region near 1 K, the magnitude of our measured thermal conductivity of alumina ($\sim 1 \cdot 10^{-2}$ W/m/K) is about an order of magnitude lower than [154] ($\sim 1 \cdot 10^{-1}$ W/m/K). However, at higher temperatures, where studies have been more plentiful, we do observe a fairly significant spread between measurements. For both power dependence and absolute value, fabrication differences can result in great variations from sample to sample.

It is worth mentioning that we have also verified the efficacy of the assumptions on the effective area of the heat flow using a 2-D thermal model for the samples. A 2-D thermal model suffices, considering the cylindrical symmetry for most of our samples. We used the measured thermal conductivity in our 2-D thermal model to estimate the temperature gradient in the sample, for a given applied power. As can be seen from Fig. A.7, the modeled data matches extremely well with the experimental data for alumina sample D, which had the largest cross-sectional area among all the samples. The error bars on the 2-D thermal model comes from the uncertainty in the measured thermal conductivity. A short description about the 2-D thermal model can be found in Section A.5.

Sample	Temp. Range (K)	μ (W/m/K)	ν	λ (W/m/K)	γ	corr[μ, ν]
A	0.37–1.30	$(2.40 \pm 0.16) \cdot 10^{-2}$	2.76 ± 0.03	-	-	0.014
B	0.46–0.78	$(2.59 \pm 0.67) \cdot 10^{-2}$	2.95 ± 0.51	-	-	0.954
C	0.57–0.85	$(2.49 \pm 0.24) \cdot 10^{-2}$	3.33 ± 0.18	-	-	0.633
D	0.10–0.72	$(1.65 \pm 0.18) \cdot 10^{-2}$	3 (fixed)	$(8.59 \pm 3.81) \cdot 10^{-4}$	1.09 ± 0.22	<i>see matrix A.1</i>

Table A.2: Results for alumina thermal conductivity: $K(T) = \mu(T/1 \text{ K})^\nu$ (A, B, C) or $K(T) = \mu(T/1 \text{ K})^3 + \lambda(T/1 \text{ K})^\gamma$ (D) [70].

A.4 Conclusion

We measure the thermal conductivity of PTFE in the 0.17–0.43 K range. The average result is $K_w(T) = (5.0 \pm 0.8) \cdot 10^{-3} (T/1 \text{ K})^{1.83 \pm 0.11} \text{ W/m/K}$. This measurement is consistent with existing results and validates our experimental techniques. We measure the thermal conductivity of alumina in the ~ 0.1 –1.3 K range, over which the value varies from $\sim 9 \cdot 10^{-5}$ to $\sim 4 \cdot 10^{-2} \text{ W/m/K}$. The temperature dependence of thermal conductivity is cubic, with a smaller linear contribution coming out below $\sim 0.3 \text{ K}$, consistent with microcrystalline structure. In the vicinity of 0.1 K, we find our Alumina to be slightly less conductive than PTFE. This ceramic is widely available, and, compared to PTFE, is very rigid and solid, allowing any structural parts to be made smaller to significantly reduce heat capacity. In the process of our measurements, we find it to be much more compatible than PTFE with copper in terms of thermal contraction/expansion—we are able to hold alumina in copper clamps, which fail for PTFE due to differential contraction. As such, it appears to be a good candidate for use as a weak link in future cryogenic bolometers. We plan to test this with prototype detectors in our cryostat.

A.5 Technical Details

Applied Heater Power Fit Results

Sample	Temperature Range (K)	a (W/m)	b	c (W/m)	corr[a, b]
1	0.18–0.26	$(1.52 \pm 0.29) \cdot 10^{-3}$	2.70 ± 0.17	$(-1.57 \pm 0.15) \cdot 10^{-5}$	0.983
1	0.27–0.36	$(2.07 \pm 0.35) \cdot 10^{-3}$	2.99 ± 0.20	$(-4.15 \pm 0.42) \cdot 10^{-5}$	0.977
1	0.37–0.43	$(2.42 \pm 0.60) \cdot 10^{-3}$	3.05 ± 0.41	$(-1.12 \pm 0.19) \cdot 10^{-4}$	0.989
2	0.17–0.21	$(1.77 \pm 0.50) \cdot 10^{-3}$	2.82 ± 0.22	$(-0.84 \pm 0.12) \cdot 10^{-5}$	0.990

Table A.3: Results for the PTFE scaled power fits: $P_{\text{applied}}(T) \cdot l/A = a(T/1 \text{ K})^b + c$ [70].

Sample	Temp. Range (K)	a (W/m)	b	c (W/m)	r (W/m)	s	corr[a, b]
A	0.37–1.30	$(6.38 \pm 0.43) \cdot 10^{-3}$	3.76 ± 0.03	$(-1.5 \pm 0.1) \cdot 10^{-4}$	-	-	-0.105
B	0.46–0.78	$(6.55 \pm 0.93) \cdot 10^{-3}$	3.95 ± 0.51	$(-3.1 \pm 0.9) \cdot 10^{-4}$	-	-	0.838
C	0.57–0.85	$(5.76 \pm 0.44) \cdot 10^{-3}$	4.33 ± 0.18	$(-5.2 \pm 0.5) \cdot 10^{-4}$	-	-	0.240
D	0.10–0.72	$(4.12 \pm 0.45) \cdot 10^{-3}$	4 (fixed)	$(-2.2 \pm 0.6) \cdot 10^{-6}$	$(4.1 \pm 1.4) \cdot 10^{-4}$	2.09 ± 0.22	see matrix A.2

Table A.4: Results for the alumina scaled power fits: $P_{\text{applied}}(T) \cdot l/A = a(T/1 \text{ K})^b + c$ (A, B, C) or $P_{\text{applied}}(T) \cdot l/A = a(T/1 \text{ K})^4 + r(T/1 \text{ K})^s + c$ (D) [70].

Alumina Sample D Correlation Matrices

$$\begin{pmatrix} \text{corr}[\mu, \mu] & \text{corr}[\mu, \lambda] & \text{corr}[\mu, \gamma] \\ \text{corr}[\lambda, \mu] & \text{corr}[\lambda, \lambda] & \text{corr}[\lambda, \gamma] \\ \text{corr}[\gamma, \mu] & \text{corr}[\gamma, \lambda] & \text{corr}[\gamma, \gamma] \end{pmatrix} = \begin{pmatrix} 1.000 & -0.003 & -0.004 \\ -0.003 & 1.000 & 0.979 \\ -0.004 & 0.979 & 1.000 \end{pmatrix}$$

Matrix A.1: Alumina sample D thermal conductivity $K(T) = \mu(T/1 \text{ K})^3 + \lambda(T/1 \text{ K})^\gamma$ parameter correlations. Parameter values and uncertainties are given in Table A.2. ν is fixed, and correlations with it are trivially zero [70].

$$\begin{pmatrix} \text{corr}[a, a] & \text{corr}[a, r] & \text{corr}[a, s] \\ \text{corr}[r, a] & \text{corr}[r, r] & \text{corr}[r, s] \\ \text{corr}[s, a] & \text{corr}[s, r] & \text{corr}[s, s] \end{pmatrix} = \begin{pmatrix} 1.000 & -0.812 & -0.861 \\ -0.812 & 1.000 & 0.965 \\ -0.861 & 0.965 & 1.000 \end{pmatrix}$$

Matrix A.2: Alumina sample D scaled power fit $P_{\text{applied}}(T) \cdot l/A = a(T/1 \text{ K})^4 + r(T/1 \text{ K})^s + c$ parameter correlations. Parameter values and uncertainties are given in Table A.2. b is fixed, and correlations with it are trivially zero. Correlations with c do not affect thermal conductivity calculation [70].

2-D Thermal Model

The steady state heat transfer in a circular cylindrical rod, in the absence of internal heat source, can be described by the following equation:

$$\nabla \cdot (K \nabla T) = 0, \quad (\text{A.14})$$

where K is the thermal conductivity of the rod. Since the sample is axisymmetric, it is convenient to write the above expression in a cylindrical form

$$\frac{1}{r} \frac{\partial}{\partial r} \left(K r \frac{\partial T}{\partial r} \right) - \frac{\partial}{\partial z} \left(K \frac{\partial T}{\partial z} \right) = 0. \quad (\text{A.15})$$

The above equation allows us to analyze the 3-D axisymmetric problem using a 2-D model. A MATLAB routine, the Partial Differential Equation ToolboxTM [168], is used to analyze the model with appropriate boundary conditions. These include keeping one side of the rod at a constant temperature, while continuously adding heat from the other side. We assume that the heat transfer at the outer boundary takes place solely due to radiation exchange. The ambient temperature is fixed at 0.1 K for the same. The output of this routine for alumina sample D is displayed below in Fig. A.6. Fig. A.7 shows the applied power against temperature, as obtained by the thermal model for alumina sample D, overlaid with the data and fit from Fig. A.5a, where we assume 1-dimensional heat flow. The 2-D model agrees with measurements within 5%.




Figure A.6: Temperature gradient along the alumina sample D at 100 mK. The edge at $z = 0$ m was kept at a constant temperature while a constant heat power of $P_{\text{applied}} = 4.5$ nW was used as a boundary condition on the other end ($z = 0.076$ m). Numerical solution to Eq. A.15 was obtained through use of the Partial Differential Equation ToolboxTM and using the measured K for the sample. The vertical lines denote isotherms and clearly shows an uniform heat flow along the sample [70].




Figure A.7: Plot of applied power versus temperature as obtained from the 2-D thermal model and overlaid on the experimental data points for alumina sample D. The horizontal error bars on the 2-D model originates from propagation of the uncertainty in the measured thermal conductivity that was used in the thermal model [70].

Appendix B

The $A = 130$ Isobar




Figure B.1: A detailed chart of the $A = 130$ isobar. ^{130}Te , the $0\nu\beta\beta$ decay candidate isotope studied in the CUORE experiment, is energetically forbidden from undergoing single- β decay to its would-be daughter ^{130}I . Instead, it decays to ^{130}Xe via the rare $2\nu\beta\beta$, and possibly the hypothetical $0\nu\beta\beta$ channels. Symmetrically, ^{130}Ba would undergo a $\beta^+\beta^+$ decay. Figure from the LUND/LBNL Table of Radioactive Isotopes [169].

Appendix C

List of Abbreviations and Acronyms

BAT	B ayesian A nalysis T oolkit—a software package featured in Chapters 10 and 11.
BSM	Physics B eyond the S tandard M odel (SM).
CI	Statistical C onfidence I nterval
CL	Statistical C onfidence L evel
CR	C leanroom. Used here in the context of CUORE assembly and installation.
CR1	Entry/gowning CR in the CUORE hut.
CR2	Assembly CR in the CUORE hut where the gluing of NTDs and heaters to TeO ₂ crystals took place.
CR3	Assembly CR in the CUORE hut where tower assembly and wire bonding took place.
CR4	CR in the CUORE hut used primarily for storage.
CR5	CR in the CUORE hut with access to the cryostat interior.
Cu-PEN	A type of PCB construction, with copper traces on a PEN substrate.
CR6	Special radon-free softwall CR inside CR5, used for CUORE detector installation (Chapter 7).
CUORE	The C ryogenic U nderground O bservatory for R are E vents.
CUPID	C UORE U ppgrade with P article I dentification—a proposed future bolometric $0\nu\beta\beta$ decay search.
DAQ	D ata A cquisition; refers to the electronic hardware and software performing this task.
DCS	D etector C alibration S ystem. Normally refers to CUORE’s unique internal, cryogenic setup. “External DCS” describes the alternative system featured in Chapter 8.
DMX	D igital M ultiplex—as in DMX512, a type of serial bus and control standard.

DR	D ilution R efrigerator—a sub-Kelvin $^3\text{He}/^4\text{He}$ cryogenic device, used to maintain CUORE at ~ 10 mK.
DU	D ilution U nit—the main active component of a DR, through which the helium mixture circulates.
FCS	F ast C ooling S ystem—setup for circulating large volumes of pre-cooled helium gas through the IVC of the CUORE cryostat to reduce cooldown time.
FWHM	F ull W idth at H alf M aximum (of a spectral peak).
GPB	G eneral P urpose I nterface B us, a.k.a. IEEE-488 or HP-IB. A hardware computer control standard/protocol/connector.
HEPA	H igh- E fficiency P articulate A ir filter, commonly used in cleanrooms.
IVC	I nnner V acuum C hamber of the CUORE cryostat, at the 4-K temperature stage.
LBNL	L awrence B erkeley N ational L aboratory.
LNGS	L aboratori N azionali del G ran S asso—Italian for Gran Sasso National Laboratory—an underground facility in Abruzzo, Italy, where the CUORE detector is located.
MCMC	M arkov C hain M onte C arlo—used for fitting in Bayesian analyses.
NLL	N egative (natural) L ogarithm of the L ikelihood function.
NOSV	A very high-purity variety of Electrolytic Tough Pitch copper used in the CUORE frames and cryostat components. Proprietary name of the manufacturer, Aurubis AG [96].
NTD	N eutron T ransmutation D oped Ge thermistors, or the technique of N eutron T ransmutation D oping used to manufacture them.
OVC	O uter V acuum C hamber of the CUORE cryostat—the outermost, 300-K vessel.
PDF	P robability D ensity F unction.
PE	P olyethylene (C_2H_4) _n , a plastic.
PEN	P olyethylene N aphthalate ($(\text{C}_{14}\text{H}_{10}\text{O}_4)_n$), a plastic.
PID	P roportional, I ntegral, D erivative feedback controller (Chapter 6).
PTFE	P olytetrafluoroethylene ($(\text{C}_2\text{F}_4)_n$), a plastic commonly known by its trademark name Teflon.
PTR	P ulse T ube R efrigerator—a closed-cycle cryocooler, used in the CUORE cryostat.
PT	P ulse T ube, synonymous with PTR.
RAS	R adon A batement S ystem, which scrubs radon from air and provides the atmosphere of CR6 (Chapter 7).
RM	R adon M onitor—an sensitive α spectrometer used for monitoring the activity of the atmosphere in CR6 (Chapter 7).
SM	The S tandard M odel of particle physics.

ν SM	A modification/expansion of the SM to include massive neutrinos, pronounced “new SM”.
SQUID	S uperconducting Q uantum I nterference D evice.
UEML	U nbinned E xtended M aximum L ikelihood—a curve-fitting technique.
ULPA	U ltra- L ow P articulate (or P enetration) A ir filter, similar but superior to HEPA. Used for re-circulation in CR6.