

Abstract

A search for neutrinoless double-beta decay in tellurium-130 with CUORE

Jeremy Stein Cushman

2018

The Cryogenic Underground Observatory for Rare Events (CUORE) is a ton-scale cryogenic experiment designed to search for neutrinoless double-beta ($0\nu\beta\beta$) decay in ^{130}Te . The experiment consists of 988 ultracold TeO_2 bolometric crystals, which act as both the $0\nu\beta\beta$ decay sources and detectors, in a close-packed configuration. This dissertation presents a search for $0\nu\beta\beta$ decay with the first two months of CUORE data. An observation of $0\nu\beta\beta$ decay would be direct evidence of lepton number violation and unambiguously prove that neutrinos are Majorana particles.

We analyze the first 83.6 kg·yr of TeO_2 exposure and find no evidence for $0\nu\beta\beta$ decay. We set a limit of $T_{1/2}^{0\nu} > 1.5 \times 10^{25}$ yr (90% C.L.) by combining this exposure with that from two predecessor experiments, CUORE-0 and Cuoricino. With this data, we set the world-leading limit on the rate of $0\nu\beta\beta$ decay in ^{130}Te . The CUORE bolometer array is characterized by an effective energy resolution of (7.7 ± 0.5) keV FWHM and background of (0.014 ± 0.002) counts/(keV·kg·yr) at the $0\nu\beta\beta$ decay Q -value. This is the lowest background level achieved to date in such a large-scale cryogenic experiment, meeting our expectations and requirements for this search.

A good understanding of the detector energy scale through regular calibration is crucial for observation of $0\nu\beta\beta$ decay and other rare processes. The 988 bolometers of the CUORE detector are calibrated with low-activity ^{232}Th sources, which are cooled and routed to their positions inside the cryostat at monthly intervals. This dissertation also discusses the design and implementation of this calibration system.

**A search for neutrinoless double-beta
decay in tellurium-130 with CUORE**

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Jeremy Stein Cushman

Dissertation Director: Karsten M. Heeger

2018

© 2018 by Jeremy Stein Cushman

All rights reserved.

Contents

List of Figures	v
List of Tables	ix
Acknowledgements	x
1 Introduction	1
1.1 Beta decay	1
1.2 Neutrino physics experiments	7
1.3 Open questions	14
2 Double-beta decay	16
2.1 Neutrino oscillation and mass	16
2.2 Neutrinoless double-beta-decay theory	25
2.3 Searching for neutrinoless double-beta decay	33
3 CUORE	40
3.1 Detectors	41
3.2 Before CUORE	49
3.3 Creating the detector towers	53
3.4 Cryostat	61
3.5 Sensitivity	66
4 Detector Calibration System	71
4.1 System overview	72
4.2 Electronic control and monitoring system	96
4.3 Performance in commissioning runs	102
5 Data acquisition and processing	107
5.1 Acquiring data	108
5.2 Processing data	111
5.3 Selecting data	121
6 Neutrinoless double-beta decay analysis	130
6.1 Detection efficiency	131
6.2 Detector response and resolution	136

6.3	Neutrinoless double-beta decay fit	142
7	Conclusion	154
7.1	CUPID: The future for bolometric detectors	154
7.2	Sensitivity and discovery	158
	Bibliography	161
A	Calibration electronics	174
A.1	In the rack	174
A.2	Rack-to-cryostat wiring	191
A.3	On the cryostat	195

List of Figures

1.1	The isobars of mass $A = 130$ with allowed decay modes	4
1.2	Feynman diagrams for $2\nu\beta\beta$ decay and $0\nu\beta\beta$ decay	5
1.3	Energy spectra for $2\nu\beta\beta$ decay and $0\nu\beta\beta$ decay	5
1.4	Schematic of the Cowan–Reines neutrino experiment	9
1.5	Schematic of the Lederman–Schwartz–Steinberger neutrino experiment	10
1.6	Schematic of the Super-Kamiokande detector	13
2.1	Masses of the Standard Model fermions	20
2.2	Masses and flavor composition of the neutrino mass eigenstates	23
2.3	Comparisons of various measurements of the neutrino masses	26
2.4	Black box diagram for the generation of a Majorana neutrino mass	27
2.5	Phase space factors for a variety of $0\nu\beta\beta$ decay candidates	28
2.6	Figure of merit for various $0\nu\beta\beta$ decay candidate isotopes	30
2.7	$0\nu\beta\beta$ decay Q -values and natural isotopic abundances	35
3.1	Schematic of a CUORE bolometer	41
3.2	Representative NTD thermistor pulse	42
3.3	Thermal circuit of a CUORE bolometer	45
3.4	History of $0\nu\beta\beta$ decay experiments with TeO_2 bolometers	50
3.5	Cuoricino detector tower and energy spectrum in the $0\nu\beta\beta$ decay region of interest	51
3.6	Full Cuoricino energy spectrum	52
3.7	CUORE-0 detector tower and energy spectrum in the $0\nu\beta\beta$ decay region of interest	53
3.8	Backgrounds in the Cuoricino and CUORE-0 experiments	54
3.9	Components of a CUORE detector tower	55
3.10	NTD germanium thermistor photograph and sketch	57
3.11	Copper pads and traces on Cu-PEN tapes	58
3.12	Glue dots on a thermistor and heater	59
3.13	Crystal gluing glovebox	60
3.14	Tower assembly glovebox	61
3.15	Wire-bonded heater and thermistor	62
3.16	Rendering of the CUORE cryostat	63
3.17	Rendering of the CUORE cryostat support structure and external shielding	64

3.18	The detector towers installed in the cryostat	64
3.19	CUORE sensitivity over time	69
4.1	Illustration of the DCS in the CUORE cryostat	73
4.2	Rendering of the jig used for the wire electrical discharge machining of the source wires	78
4.3	Schematic and photograph of a source capsule and bottom of a source string	79
4.4	Location of the strings during calibration with respect to the shields and towers	80
4.5	Simulated event rates per bolometer during calibration	81
4.6	Simulated CUORE calibration spectrum, summed over all channels	82
4.7	Rendering of a motion box	84
4.8	Diagram of the motor and spool controlling a single source string	85
4.9	Schematic of DCS guide tubes in the CUORE cryostat	88
4.10	An S-tube assembly in the cryostat	89
4.11	600-mK guide tubes installed below a thermalizer	91
4.12	Helical outer guide tubes installed in the CUORE cryostat	92
4.13	A thermalizer mounted in the CUORE cryostat	93
4.14	Rendering and schematic of the 4-K thermalization hardware	95
4.15	Overview of the DCS control system	97
4.16	Load cell readings for a single string during deployment and extraction	100
4.17	Cryostat base temperature and string position during a full inner-string deployment	104
4.18	Cryostat base temperature and string positions during a 12-string extraction	106
5.1	Example of the CUORE preprocessing on a raw pulse	112
5.2	Scatter plot of heater pulse amplitude vs. pulse baseline	115
5.3	Uncalibrated spectrum for one channel in one dataset	116
5.4	Crystal Ball fit to the 2615 keV line in one channel in one dataset	119
5.5	Example of the CUORE calibration function for a single channel	120
5.6	Example of the CUORE blinding technique	121
5.7	Example of a pulse shape parameter dependence on energy	125
5.8	Cause of jitter between channels and its effect on timing	127
6.1	Signal-to-background figure of merit with various limits on the acceptable pulse shape parameters	133
6.2	Overall fit to the 2615 keV calibration line	139
6.3	Distribution of resolutions at the 2615 keV peak in calibration data	140
6.4	Overall fit to the 2615 keV background line	141
6.5	CUORE physics spectrum, including all cuts	142
6.6	Energy resolution and energy bias	144
6.7	Best fit to the $0\nu\beta\beta$ decay region of interest	145
6.8	Negative log-likelihood for the best-fit model at various signal rates	146

6.9	Distribution of half-life limits and best-fit negative log-likelihoods from pseudo-experiments	147
6.10	Negative log-likelihood curve including systematic uncertainties and predecessor experiments	150
6.11	Energy spectra from CUORE-0 and CUORE	152
6.12	Limits on the effective Majorana mass from CUORE	153
7.1	CUPID effective Majorana mass sensitivity	158
A.1	DCS setup overview	176
A.2	Rack diagram	177
A.3	Motion control box layout and photo	181
A.4	Motion control box back panel	181
A.5	Motion control box power connection	182
A.6	Motion control box PCB schematic	182
A.7	UMI feedback cable	183
A.8	TM Contact wiring schematic	183
A.9	UMI potentiometer and preamp signal wiring	184
A.10	Motion control box gate valve wiring	184
A.11	Gate valve wiring schematic	185
A.12	Motor power inside the motion control box	186
A.13	Rack breakout box back panel	187
A.14	Rack breakout box signal wiring	188
A.15	Rack breakout box linear actuator power wiring	189
A.16	Temperature monitoring system overview	189
A.17	LakeShore DIN input pinout	190
A.18	600 mK thermometry breakout	190
A.19	Cable A: Encoders and sensors	192
A.20	Cable B: Preamp signals	192
A.21	Cable C: Motor power signals	193
A.22	Cable E: Gate valve control	193
A.23	Cable F: Thermalizer thermometry and TM contact signals	193
A.24	Cable G: Linear actuator power	194
A.25	Cable H: Linear actuator potentiometer signals	194
A.26	Cable K: 600 mK thermometry	194
A.27	Motion box PCB enclosure and wiring	196
A.28	Motor grounding scheme	197
A.29	Motion box PCB schematic	197
A.30	Encoder cable schematic	198
A.31	Preamp-to-motion box connection	198
A.32	Preamp output-to-PCB connection	199
A.33	Drive spool internal wiring	199
A.34	Proximity sensor input wiring	199
A.35	Proximity sensor output wiring	200
A.36	Gate valve wiring	200

A.37 Fischer connector wiring for thermalizer thermometers	201
A.38 Fischer connector wiring for 600 mK thermometers	202
A.39 Fischer feedthrough in-vacuum PCB schematic	202
A.40 SIP connector diagram	203
A.41 Potentiometer wiring breakout	204
A.42 Linear actuator wiring	204

List of Tables

2.1	Recent results from selected $0\nu\beta\beta$ decay experiments	36
2.2	Final projected sensitivity of current and upcoming $0\nu\beta\beta$ decay experiments	36
3.1	Bulk contamination levels of the CUORE detector components	55
4.1	Cooling power budgeted for the calibration system at all thermal stages of the cryostat	75
5.1	Peaks used for calibrating the CUORE detectors	116
5.2	Functions used to determine the dependence of the the pulse shape parameters on energy	125
6.1	Cuts used in the final $0\nu\beta\beta$ decay analysis	134
6.2	Selection efficiencies for the cuts used in the $0\nu\beta\beta$ decay analysis	135
6.3	Fit results for the prominent γ background lines	143
6.4	Systematic uncertainties used in the $0\nu\beta\beta$ decay analysis	149
A.1	Wiring from the rack to the cryostat	191
A.2	Connectors on the rack-to-cryostat cables	192

Acknowledgements

I would like to begin by thanking my advisor, Prof. Karsten Heeger, for taking me on as his graduate student, and for his invaluable advice, guidance, and support during my research. I am grateful to have had an advisor who allowed me to explore new ideas and who always supported my endeavors, no matter how promising or futile. In addition to the physics, I am lucky to have had Karsten's example in learning how to be a strong leader and how to effectively communicate my work.

I would also like to thank Prof. Reina Maruyama for her support throughout the past several years, both in my physics analysis work and in general. Reina's deep knowledge of a variety of fields of physics and her dedication to our group has been a great asset to all of the work that we've done for CUORE, and for me personally.

I would like to thank the other members of my committee, Prof. Keith Baker, Prof. Bonnie Fleming, and Prof. Francesco Iachello, who have each provided valuable guidance and feedback on my work and on CUORE.

This dissertation is built on the many years of work of the 250+ current and former members of the CUORE collaboration, many of whom I've worked with directly. It takes a great deal of teamwork and coordination to build such a huge experiment with physicists, engineers, and technicians scattered around the globe, and it has been a pleasure working with such a friendly, cooperative, and brilliant group of scientists.

I would like to particularly thank the past and present members of the CUORE group at Yale: Tom Wise, for his incredible expertise, judgement, and advice in en-

gineering and experimental design; Ke Han, for his unbelievably broad knowledge of CUORE and the field; Kyungeun Lim, for her fantastic analysis skills and friendly advice; Danielle Speller, for her amazing insightfulness and thoroughness; and Christopher Davis, for his great intuition and cheerful optimism through many, many long nights commissioning and operating the calibration system at LNGS.

Outside of CUORE, I would like to thank my colleagues, classmates, and friends in Wright Laboratory, at Yale, and beyond, for great discussions on physics, music, politics, and life.

Long before I was a graduate student, I had the loving support and advice of my family, both of my parents and my sister, Rebecca. Throughout my whole life, my parents have let me explore and dive into my passions, always encouraging me and guiding me but letting me explore the world for myself and make my own decisions. And Rebecca's passion for making the world a better, more inclusive place has a huge impact on how I see my work and my life today.

And finally, I would like to thank my partner, Natalie Heer. Natalie's encouragement, guidance, intuition, and advice, both as a fellow scientist and as a friend, has been an incredibly important part of my work, and of my life.

Chapter 1

Introduction

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

— Wolfgang Pauli

Neutrinos are everywhere. Trillions of them are passing through this page every second. Yet despite their abundance, neutrinos are remarkably difficult to study. They are extremely light, electrically neutral, and not subject to the strong nuclear force. They can traverse entire galaxies interacting with nothing. But neutrinos play a key role in the universe, created in stellar nucleosynthesis, in supernovae, and immediately after the Big Bang. And understanding them is critical to understanding the evolution of the universe and why we exist today.

1.1 Beta decay

By the early twentieth century, it was well established that beta (β) decay involves the transformation of one element into another, and that the particle emitted in the decay is an electron. It was also understood that when a nucleus decays, it releases the same energy every time. As expected, when alpha (α) decay was observed, each

emitted α particle appeared at one of a few defined energies. Yet in β decay, the electron always seemed to carry away a different fraction of the total energy of the decay. This conundrum was the source of great confusion and controversy in the science community [1].

Such was the state of nuclear physics in 1930, when a number of prominent physicists gathered for a conference in the German town of Tübingen. Wolfgang Pauli was not in attendance on account of his being “indispensable” at a ball in Zürich, but he sent along his thoughts in a letter to the conference attendees [2]. Perhaps, he offered, the continuous energy spectrum of β decay could be explained if a yet-undetected neutral particle was emitted along with the electron in the decay. This neutral particle would then escape with a varying fraction of the energy of the decaying nucleus. By his own admission, it was a “desperate remedy.” But by 1934, Enrico Fermi had devised a full theory of β decay incorporating Pauli’s new particle [3], which he named the neutrino, and neutrino physics had begun.

1.1.1 Theoretical predictions

When Pauli postulated the existence of the neutrino in 1930, atoms were thought to contain only a nucleus of protons surrounded by electrons. The field was progressing rapidly, however, and by 1932 the neutron had been discovered; Werner Heisenberg and others quickly produced significantly more accurate models of the nucleus involving both protons and neutrons [4]. Fermi’s theory of β decay unified the neutrino with this new model of the atom, explaining β decay as the decay of a neutron to a proton, accompanied by the emission of both an electron and a neutrino.

This formal theory of β decay set the stage for a flurry of theoretical work about this decay and the neutrino in general. Fermi’s theory dealt only with β^- decay, but in the same year that his theory was finalized, Irène and Frédéric Joliot-Curie announced the discovery of positron emission as a new form of radioactivity [5]. By

this time, Paul Dirac's theory of antimatter was established [6] and positrons were understood to be anti-electrons [7]. Fermi's assistant, Gian Carlo Wick, applied this theory to positron emission, explaining it as a combination of β^+ decay and electron capture (EC) and introducing the antineutrino [8].

In the years following the discovery of β^+ decay and the postulation of the antineutrino, Ettore Majorana formulated a new approach to describe Dirac's theory of matter and antimatter, which had special relevance for the neutrino. In Majorana's theory, unlike in Dirac's, there is no need to infer the existence of antiparticles for electrically-neutral particles, such as the neutrino. Majorana notes that if the neutrinos are, in fact, particles described by his new formalism (now known as "Majorana particles"), the theory of β decay still holds, with both β^- and β^+ decay accompanied by the emission of a neutrino [9].

Meanwhile, Maria Goeppert-Mayer was formalizing a theory of double-beta ($\beta\beta$) decay; $\beta\beta$ decay involves the simultaneous emission of two electrons (and two neutrinos) from an atomic nucleus. In a 1935 paper, she explained the theory and predicted that such decays could have extremely long half-lives [10]. The idea was motivated by the fact that multiple isobars (i.e., elements with the same number of nucleons but different atomic number) appear in nature for a variety of mass numbers, even though theoretically only the lowest-mass isobar of each should be stable. Goeppert-Mayer correctly predicted that there exist nuclei that cannot emit a single β particle because the resulting isobar would have a greater energy, but that can emit two β particles simultaneously, bringing them to a lower-energy isobar (see [Figure 1.1](#)). Because this is a second-order process, the decay is sufficiently suppressed so as to make the $\beta\beta$ emitter appear stable over geological time scales.

A few years following Goeppert-Mayer's theory of $\beta\beta$ decay, Wendell H. Furry had a new proposal. Majorana had shown that if the neutrino is a Majorana particle then the results of β decay remain unchanged. However, the same need not be true for

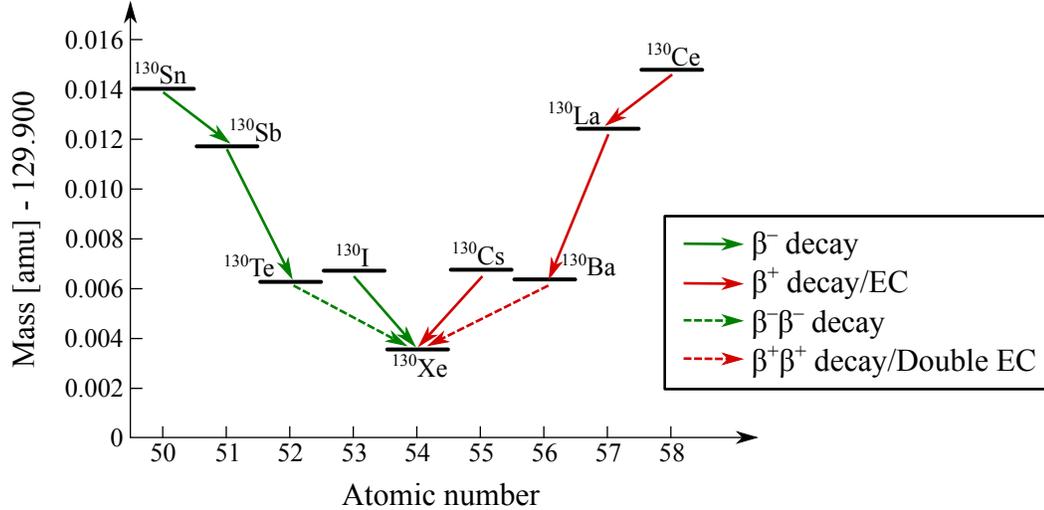


Figure 1.1: The isobars of mass $A = 130$ with allowed decay modes. Because the mass of ^{130}I is greater than that of ^{130}Te , ^{130}Te cannot undergo single β decay, but it can undergo $\beta\beta$ decay to ^{130}Xe . Likewise, ^{130}Ba cannot decay to the higher-mass ^{130}Cs but can decay directly to ^{130}Xe . (Not shown is the rare β^- decay of ^{130}Cs to ^{130}Ba .) Atomic masses from AME2012 [11].

$\beta\beta$ decay. In particular, Furry showed that a completely different decay mechanism can occur in the case that the neutrino is Majorana, with potentially very different half-lives [12]. In this process, now called neutrinoless double-beta ($0\nu\beta\beta$) decay, no neutrinos are emitted, and the decay is accompanied by the emission of only β particles (electrons). As Furry envisioned it, a virtual neutrino would be emitted along with one electron and absorbed along with the emission of the second electron (see Figure 1.2).

While the continuous spectrum of β decay prompted the birth of neutrino physics, the sum of the energies of the electrons emitted in $0\nu\beta\beta$ decay is actually constant. The full energy of the decay is the difference in mass between the initial and final state nuclei, commonly called the Q -value of the decay. In single beta decay and two-neutrino double-beta ($2\nu\beta\beta$) decay, the neutrinos carry away part of the energy released in the decay. In neutrinoless double-beta ($0\nu\beta\beta$) decay, there are no neutrinos in the final state, so all of the energy is carried by the electrons (see Figure 1.3).

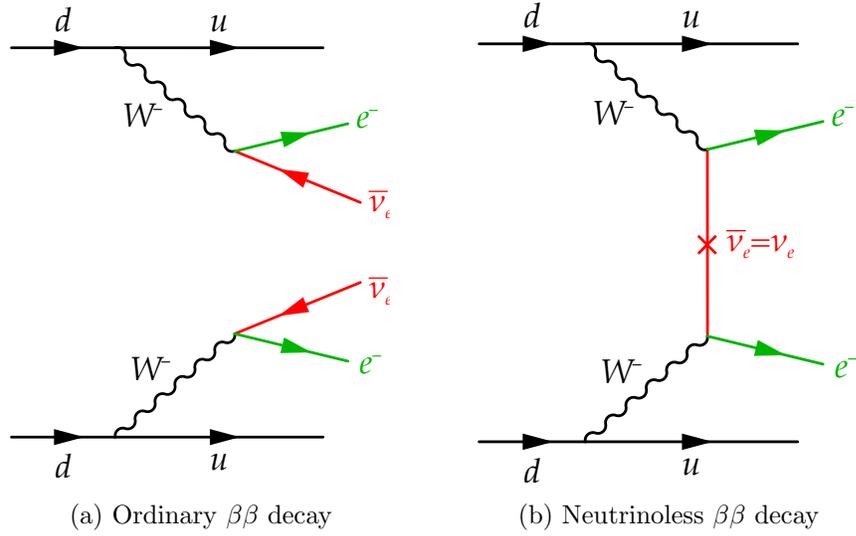


Figure 1.2: Feynman diagrams for ordinary $2\nu\beta\beta$ decay and hypothesized $0\nu\beta\beta$ decay. The process shown for $0\nu\beta\beta$ decay, light Majorana neutrino exchange, is one of several theorized possibilities.

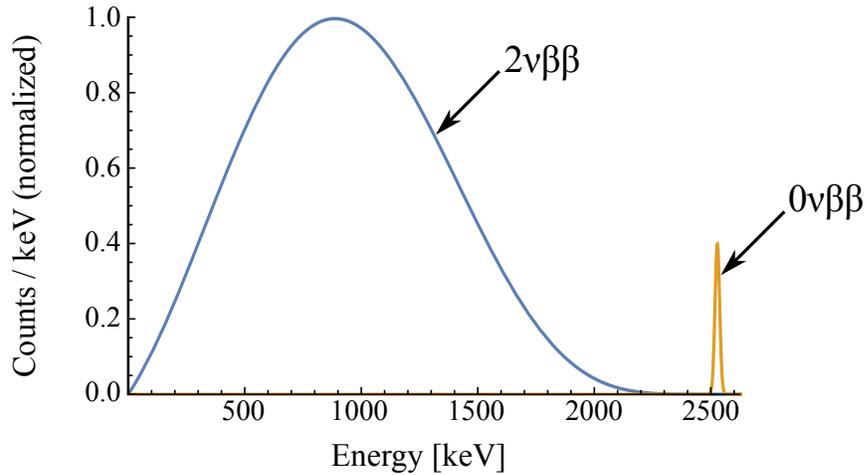


Figure 1.3: Energy spectra for ordinary $2\nu\beta\beta$ decay and hypothesized $0\nu\beta\beta$ (neutrinoless) decay of ^{130}Te . The shape of the $2\nu\beta\beta$ decay spectrum follows the Primakoff–Rosen approximation [13]. The $0\nu\beta\beta$ decay peak is smeared slightly in energy, as it would appear in any experiment, and greatly exaggerated in magnitude.

1.1.2 Early searches for double-beta decay

Because of a lack of understanding of the nature of the neutrino and the weak interaction at the time, Furry and others believed that the $0\nu\beta\beta$ decay rate could be significantly greater than the ordinary double-beta ($2\nu\beta\beta$) decay rate. In particular, Goeppert-Mayer herself estimated half-lives of over 10^{17} years for $2\nu\beta\beta$ decay, and her formula gives half-lives of 10^{25} years for isobars with a mass difference of 0.002 amu (1.9 MeV) [10]. Furry's calculations showed that $0\nu\beta\beta$ decay could be more frequent than $2\nu\beta\beta$ decay "by a factor which ranges from 10^5 to 10^{15} or more" [12]. He was mistaken, but his predictions convinced experimentalists to take up the charge and begin searching for $\beta\beta$ decay. Some early experiments erroneously claimed success [14, 15], but all were orders of magnitude away from true observation.

The first true detection of $\beta\beta$ decay was indirect, coming from geochemical experiments. In 1950, Inghram and Reynolds correctly attributed an excess of ^{130}Xe in old tellurium ore to the $\beta\beta$ decay of ^{130}Te and calculated a half-life of 1.4×10^{21} years [16], within a factor of 2 of the currently accepted value for $2\nu\beta\beta$ decay of ^{130}Te [17]. Indirect experiments are unable to distinguish $0\nu\beta\beta$ decay from $2\nu\beta\beta$ decay, as both result in the same decay product, and instead detect the sum of the decay rates of the two processes. Inghram and Reynolds' finding of such a long half-life thus dashed any hope of $0\nu\beta\beta$ decay having half-lives accessible to experiments at the time.

Almost forty years later, in 1987, $\beta\beta$ decay was directly observed for the first time by Elliott, Hahn, and Moe [18]. In order to overcome enormous backgrounds to detect such a rare process, they used a time projection chamber, which allowed them to track the emitted electrons and select only events in which the electrons were emitted approximately back-to-back and with the correct energy. They measured the energy spectrum of the electrons emitted in the decay of ^{82}Se and found a small excess of events over their background consistent with $2\nu\beta\beta$. Their result, $1.1_{-0.3}^{+0.8} \times 10^{20}$ years, was consistent with geochemical measurements that had been made in the years

since Inghram and Reynolds’ first result, strengthening the case that they had, in fact, observed $2\nu\beta\beta$.

Since 1987, many experiments have detected $2\nu\beta\beta$ directly in a variety of isotopes [17, 19–25]. All of those observed have half-lives of approximately 10^{19} to 10^{21} years. Despite decades of searches, no experiments have yet detected $0\nu\beta\beta$ decay.

1.2 Neutrino physics experiments

When the neutrino was hypothesized to explain an anomaly in the β decay spectrum, almost nothing was known about its properties, other than that it must be electrically neutral and could not be very heavy. In fact, many physicists, including Pauli, thought that the neutrino would never be directly detected. Fortunately, Pauli was incorrect on this count; a series of highly successful neutrino experiments over the past 60 years have revealed a host of fascinating properties of the neutrino. Much, however, still remains unknown.

1.2.1 Direct detection

The neutrino was initially only known to be produced by the β -like decays: β^- decay ($n \rightarrow p + e^- + \bar{\nu}_e$), β^+ decay ($p \rightarrow n + e^+ + \nu_e$), and electron capture ($p + e^- \rightarrow n + \nu_e$). Thus, it was natural for searches for the neutrino to center around these processes, or analogous ones. In particular, a process known as “inverse beta decay” was proposed. This process is essentially β^+ decay induced by the presence of an antineutrino; symbolically, this is $\bar{\nu}_e + p \rightarrow e^+ + n$.

The first requirement for observing inverse beta decay was a large flux of neutrinos, since the interaction of an incoming neutrino with a proton was (correctly) predicted to have an extremely small cross-section. Nuclear reactors were relatively new at

the time, but had developed greatly during World War II. They derive energy from the fission of large elements into neutron-rich daughters, which themselves generally undergo β^- decay to achieve a more stable neutron-proton ratio. These β^- decays result in the emission of antineutrinos in large numbers, making the space next to a nuclear reactor an excellent location for a neutrino detector.

Cowan and Reines realized such an experiment to observe inverse beta decay in 1953 [26]. The detector was a 300-liter tank of liquid scintillator doped with cadmium salt and surrounded by photo-multiplier tubes (PMTs). When an incoming neutrino interacts with a proton (i.e., a hydrogen nucleus) in this type of liquid scintillator, it produces a positron and a free neutron. The positron almost immediately annihilates with a nearby electron, emitting two easily detectable gamma rays. The free neutron is very likely to be absorbed by a ^{113}Cd nucleus, as the isotope ^{113}Cd has a large neutron capture cross-section. This creates $^{114}\text{Cd}^*$, an excited cadmium nucleus, which de-excites almost instantly, releasing several gamma rays in the process. This absorption and decay takes place, on average, several microseconds following the inverse beta decay and positron annihilation. The delayed coincidence signal of positron annihilation followed by neutron capture is used to identify the neutrino (see [Figure 1.4](#)).

Cowan and Reines' first neutrino detected at a reactor near Hanford, Washington, gave hints that the neutrino had really been observed, but it was their second, upgraded experiment that provided incontrovertible evidence [27]. They moved to a more powerful reactor at Savannah River, South Carolina, rebuilt their detector with a new design, and went underground to shield the detector from cosmic rays and from gamma rays and neutrons from the reactor itself. They recorded 3 neutrinos per day in their new detector, kicking off an era of neutrino direct detection experiments.

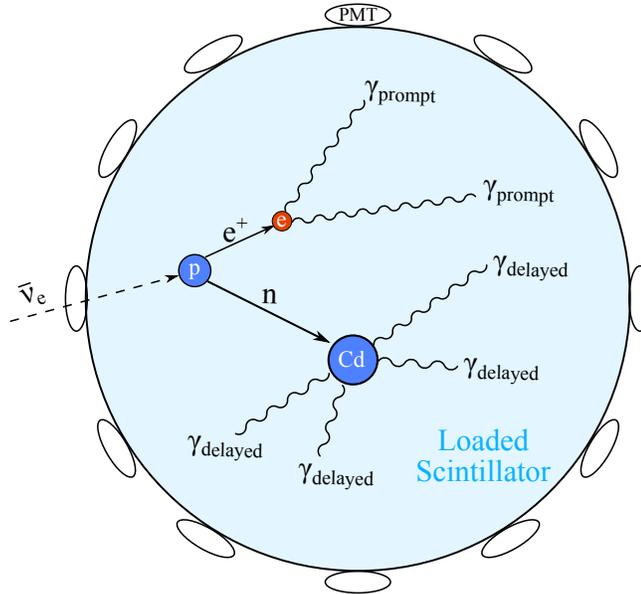


Figure 1.4: Schematic of the first Cowan–Reines neutrino experiment, the first experiment to detect free neutrinos. An incoming antineutrino interacts with a proton inside a liquid scintillator volume, turning into a neutron and emitting a positron. The positron annihilates quickly, producing gamma rays that are easily detected by PMTs outside the scintillator. The neutron is absorbed by a cadmium nucleus several microseconds later, producing a delayed signal in the PMTs.

1.2.2 Beams and flavors

The next major discovery in the field of neutrino physics came in 1962. By this point, pions were known to decay to muons along with the emission of a neutrino ($\pi^+ \rightarrow \mu^+ + \nu$ and $\pi^- \rightarrow \mu^- + \bar{\nu}$), and muons were known to decay to electrons with the emission of two neutrinos ($\mu \rightarrow e + \nu + \bar{\nu}$). If the neutrinos involved in these decays were the same as the neutrinos involved in β decay, however, a third decay mode should be possible: $\mu \rightarrow e + \gamma$ [28]. That decay was not observed, which had convinced many that there may be two different neutrinos at play.

Lederman, Schwartz and Steinberger designed an experiment to measure whether the neutrinos produced in pion decay and those produced in β decay are really the same particle [29]. For their experiment, they created the first neutrino beam. By colliding protons from an accelerator with a beryllium target, they produced a large

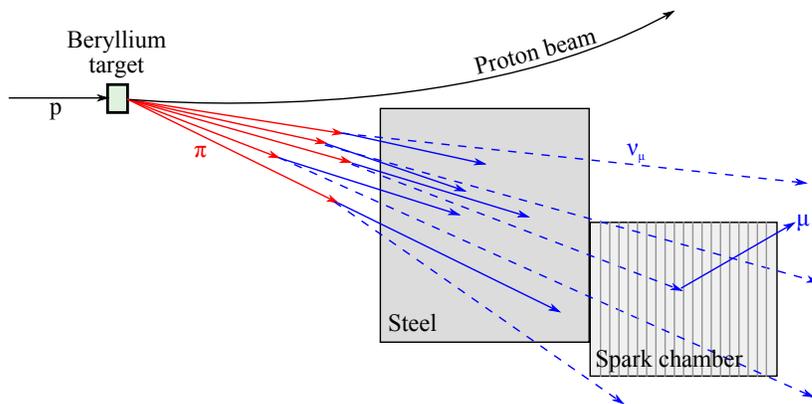


Figure 1.5: Schematic of the Lederman–Schwartz–Steinberger neutrino experiment, the first experiment to create a neutrino beam and the first to show that there are multiple types of neutrinos.

flux of pions, which decay in flight to muons, emitting a neutrino or antineutrino depending on the muon charge. The beam of muons and neutrinos was sent through a large iron shield, which stopped almost all of the muons, and then into a spark chamber that can detect and distinguish muon and electron tracks (see [Figure 1.5](#)). The experiment saw a large number of muons created by neutrinos interacting in the detector volume and very few electrons, confirming the existence of two different neutrinos: the electron neutrino (ν_e) and muon neutrino (ν_μ).

Fifteen years later, the tau lepton was discovered, implying the existence of a third neutrino (ν_τ). The tau is much more difficult to detect than the muon, traveling a fraction of a millimeter before decaying, and thus the direct detection of the tau neutrino did not come until much later. In 2001, the DONUT experiment reported the observation of 4 tau neutrino interactions in nuclear emulsion targets over an expected background of 0.34 events, confirming its existence [30]. The tau neutrino was the second-to-last standard model particle to be detected, followed only by the Higgs boson in 2012 [31, 32].

1.2.3 Oscillations

The potential for a new source of neutrinos for experiments came in the 1960s, when solar models predicted that neutrinos from the Sun might be detectable on Earth. There are several sources of solar neutrinos, emitted at a range of energies and a variety of rates. But those from ${}^8\text{B}$ decay, while only a tiny fraction of solar neutrinos, have a significantly greater energy and are thus significantly more detectable than other, more numerous, neutrinos from the Sun. The most promising detection method for these neutrinos was the reaction ${}^{37}\text{Cl}(\nu_e, e^-){}^{37}\text{Ar}$ ($Q = -0.81$ MeV); that is, an electron neutrino interacts with a ${}^{37}\text{Cl}$ nucleus, which transforms into an ${}^{37}\text{Ar}$ nucleus, absorbing 0.81 MeV and emitting an electron. Neutrinos from ${}^8\text{B}$, with energies up to ~ 15 MeV, can easily provide enough energy for this reaction, and even have enough energy to leave the ${}^{37}\text{Ar}$ nuclei in excited states, vastly increasing the cross section for ${}^8\text{B}$ neutrino interactions with ${}^{37}\text{Cl}$.

John Bahcall produced the first detailed calculations on the detection potential for this process and concluded that it should be possible to detect ${}^8\text{B}$ solar neutrinos on Earth [33]. Based on his calculations, Ray Davis proposed an experiment. A huge tank containing 380 m^3 of C_2Cl_4 (commonly known as “perc” and available in large quantities for its use in dry cleaning) would be placed deep underground at the Homestake Gold Mine in Lead, South Dakota. Solar neutrinos would penetrate the earth above (and below) and interact with the ${}^{37}\text{Cl}$ in the detector, but other cosmic rays would not. Periodically, the minuscule amount of argon generated by this interaction would be collected from the tank and counted to determine the rate of neutrinos from the Sun.

Davis’ experiment was a success, announcing the first detection of ${}^8\text{B}$ solar neutrinos in 1968 [34]. But there was a discrepancy between the rate of neutrinos observed in his experiment and the predictions based on the solar model by Bahcall, a discrepancy that lasted for three decades despite significant theoretical and experimental im-

provements [35]. Experiments, seemingly inexplicably, observed only approximately a fraction the predicted flux of solar neutrinos.

This “solar neutrino problem” persisted for years, and it was not the only unexpected and unexplained problem in neutrino physics. The so-called “atmospheric neutrino anomaly” had also emerged, as the Kamiokande and IMB water Cherenkov experiments both noted a deficit in the number of atmospheric muon neutrino events observed relative to the number of electron neutrino events [36, 37]. Physicists began to reexamine the process of neutrino oscillation, originally proposed by Bruno Pontecorvo in 1957 [38] and extended to multiple neutrinos by Maki, Nakagawa, and Sakata with the discovery of the muon neutrino in 1962 [39]. In this framework, neutrinos can oscillate between flavors as they propagate through space. The solar neutrino problem could then be explained by the oscillation of electron neutrinos into undetected muon and tau neutrinos. The existence of neutrino oscillation would have another significant implication: for neutrinos to oscillate they must have mass, in conflict with the Standard Model assumption that they are massless.

Neutrino oscillation was definitively discovered in 1998 by Super-Kamiokande (“Super-K”) [40]. Super-K observed neutrinos originating in Earth’s atmosphere, mainly generated through $\pi^+ \rightarrow \mu^+ + \nu_\mu$ followed by $\mu^+ \rightarrow e^+ + \bar{\nu}_\mu + \nu_e$ and the equivalent π^- decay. As Super-K could measure the energy and incident angle of the neutrinos, it was able to distinguish between downwards-going neutrinos, which were created just above the detector in the atmosphere, and upwards-going neutrinos, which were created in the atmosphere up to thousands of kilometers away and had passed through the earth before reaching the detector (see [Figure 1.6](#)). In the end, the experiment saw significantly fewer upwards-going muon neutrinos than would be expected from simulation, while the number of electron neutrinos was in line with predictions. This indicated that the upwards muon neutrinos had oscillated into tau neutrinos as they propagated through the earth. After ruling out a number of

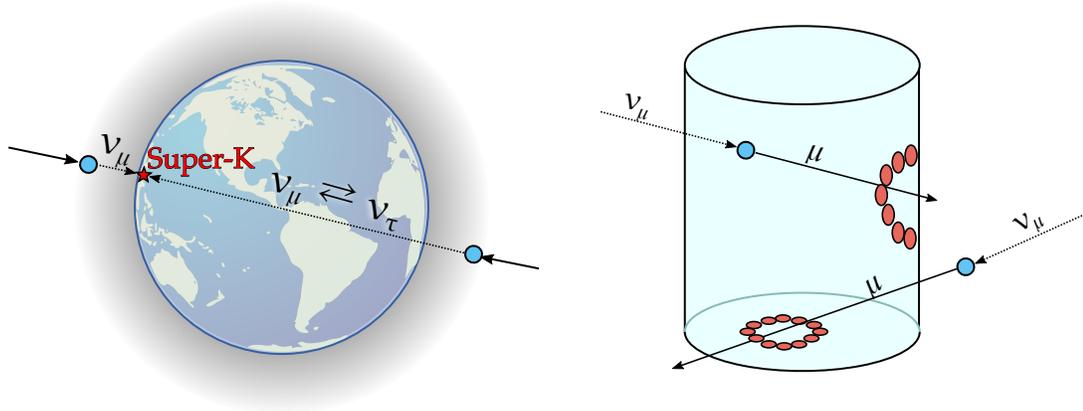


Figure 1.6: *Left:* Super-K detects muon neutrinos generated in the atmosphere as a cosmic ray (solid line) collides with a nucleus in the atmosphere (blue circle). Super-K can distinguish between those generated right above the experiment and those that had passed through the earth and oscillated. *Right:* The Super-K detection scheme. A 22-kiloton tank of water is surrounded on all sides by photomultiplier tubes (PMTs). After a neutrino interaction in the water, Super-K is able to distinguish upwards-going neutrinos (those that had traveled through the earth) from downwards-going neutrinos by looking at where the Cherenkov light from the secondary particles lands on the PMTs. Muons can be distinguished from electrons by analyzing the shape of the Cherenkov light rings.

backgrounds, Super-K was able to claim a discovery.

The solar neutrino problem was officially resolved by the SNO experiment shortly afterwards. SNO, unlike earlier solar neutrino experiments, was sensitive to three different processes: $\nu_e + d \rightarrow p + p + e^-$, in which only electron solar neutrinos participate; $\nu_x + d \rightarrow p + n + \nu_x$, in which all three neutrinos participate; and $\nu_x + e^- \rightarrow \nu_x + e^-$, in which all three participate but electron neutrinos dominate. In 2001 and 2002, the SNO collaboration confirmed that not only were solar neutrinos from ${}^8\text{B}$ decay a mix of the three types of neutrinos [41], but they were almost exactly one-third electron neutrinos [42]. This discovery of neutrino flavor transformation confirmed that neutrinos must have mass and brought experiments into good agreement with solar models [43].

1.3 Open questions

The wide energy range and disparate distance scales of neutrinos generated in different processes has lead to a proliferation of neutrino experiments in the past two decades. Experiments with neutrinos from beams ($E \sim 1 - 10$ GeV, $d \sim 10 - 10^6$ m), nuclear reactors ($E \sim 1 - 10$ MeV, $d \sim 10 - 10^5$ m), the atmosphere ($E \sim 1 - 100$ GeV, $d \sim 10^3 - 10^7$ km), and the sun ($E \sim 1 - 10$ MeV, $d \sim 10^8$ km) have explored oscillations between all three families of neutrinos. These experiments have confirmed the results of Super-K and SNO, putting to rest any doubt that neutrinos oscillate and have mass, and have enhanced our understanding of neutrino oscillations greatly. But a great deal is still unknown about the neutrino, and there are still several open questions in the field, including:

1. Are neutrinos Dirac or Majorana particles? That is, are neutrinos and antineutrinos actually different particles or the same particle, as hypothesized by Majorana? This question is a direct motivator of this dissertation, and we return to this in [Section 2.1.1](#).
2. Relatedly, is lepton number a fundamental symmetry of nature? That is, can the total number of leptons (electrons, muons, taus, and their accompanying neutrinos) change in particle interactions or decays? If neutrinos are Majorana particles and lepton number is not conserved, can neutrino oscillations help explain the dominance of matter over antimatter in the early universe?
3. What are the masses of the different neutrinos? Oscillation experiments only measure the difference in mass between neutrino states, but the absolute masses, and the order of the masses, remains unknown. We return to this question, and the question of how to define the mass of an oscillating particle, in [Section 2.1.3](#).
4. How many neutrino species are there? Is there a fourth type of neutrino that

does not interact with matter except through gravity? Several anomalies in recent neutrino experiments point to this possibility.

The only currently feasible method to address whether neutrinos are Dirac or Majorana particles is to search for $0\nu\beta\beta$ decay. Not only would this decay prove conclusively that neutrinos are Majorana particles, but it would also be the first observation of lepton number violation, with profound implications for the Standard Model of particle physics. And through our knowledge of nuclear structure, the half-life measured in $0\nu\beta\beta$ decay can give us hints of the absolute effective neutrino mass, which remains a mystery.

Chapter 2

Double-beta decay

“The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble.”

— Paul Dirac

We turn now to a quantitative discussion of neutrino masses and neutrinoless double-beta ($0\nu\beta\beta$) decay.

2.1 Neutrino oscillation and mass

Neutrino oscillation is the process by which a neutrino created with a specific lepton flavor (ν_e , ν_μ , or ν_τ) can have a different lepton flavor when detected. This is due to the fact that the neutrino is created in a weak-interaction (flavor) eigenstate and propagates through space in a mass eigenstate, but the two bases are not the same. As a result, as the neutrino travels through space, the probability of observing it in a given flavor state oscillates.

The matrix that relates the weak eigenstates to the mass eigenstates of the neutrino is known as the PMNS matrix, after Pontecorvo, Maki, Nakagawa, and Sakata,

who formalized the theory of neutrino oscillation [38, 39]. It can be expressed as

$$\begin{bmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{bmatrix} = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix}. \quad (2.1)$$

The matrix is unitary¹, and can be parameterized with three Euler angles and six phases. Three of the six phases can be removed by redefining the phase of the charged lepton fields. If the neutrino is a Dirac fermion, then 2 of the 3 remaining phases can be removed by redefining the phases of the neutrino fields; conversely, if the neutrino is a Majorana fermion, then no more can be removed, since the phase of the fields is physically observable in the Majorana mass terms. A standard parameterization of the PMNS matrix is

$$\begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{bmatrix} \begin{bmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (2.2)$$

where $c_{ij} = \cos \theta_{ij}$, $s_{ij} = \sin \theta_{ij}$, δ is the CP-violating phase, and α_1 and α_2 are the Majorana phases. A nonzero value for δ would imply a difference in oscillation probabilities for neutrinos and antineutrinos. Nonzero values of α_1 and α_2 cannot be measured by neutrino oscillation, but are a factor in $0\nu\beta\beta$ decay.

As a neutrino propagates through space, the phases of the different mass states evolve differently in time, leading to the possibility for the neutrino to be observed in a different flavor state than it was created in. The magnitude of the oscillations between flavors is controlled by the three Euler angles in the PMNS matrix. The frequency of the oscillations is controlled by the squared mass differences between the

1. The matrix is unitary in the standard three-neutrino framework. If the PMNS matrix is measured to be non-unitary, then this would imply the existence of additional non-interacting neutrinos.

neutrino mass states. In a 2-neutrino framework, the probability of ν_x oscillating to ν_y is

$$\begin{aligned}
 P(\nu_x \rightarrow \nu_y) &= \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 c^3 L}{4\hbar E}\right) \\
 &\approx \sin^2(2\theta) \sin^2\left(1.27 \cdot \Delta m^2 [\text{eV}^2] \cdot \frac{L [\text{km}]}{E [\text{GeV}]}\right), \quad (2.3)
 \end{aligned}$$

where θ is a mixing angle, $\Delta m^2 = m_{\nu_x}^2 - m_{\nu_y}^2$ is the difference in squared neutrino masses, L is the baseline of the oscillations, and E is the energy of the neutrino. Because of the large difference in observed values for the neutrino mass splittings, the 2-neutrino framework is often a good approximation for experimental results.

The three Euler angles in the PMNS matrix are best measured with combinations of different types of experiments, with neutrinos from four different sources: the sun, the atmosphere, nuclear reactors, and neutrino beams. Various experiments have used radiochemical techniques [34, 44–46], Cherenkov light [41, 47–49], scintillation [50–57], and other methods to measure these neutrinos. In general, the best measurements of θ_{12} are from a day/night asymmetry and other measurements using solar neutrinos, the best of θ_{23} are from atmospheric and accelerator ν_μ disappearance, and the best of θ_{13} are from reactor $\bar{\nu}_e$ disappearance.

2.1.1 Dirac and Majorana masses

As neutrino oscillation requires that the neutrino masses are nonzero, it is necessary to add neutrino mass terms to the Standard Model Lagrangian. These mass terms can be of two general forms: a Dirac term and a Majorana term. All other known fermions have only Dirac mass terms.

The Dirac mass term for a spinor ψ , expressed in the chiral basis as $\psi = \psi_L + \psi_R$,

is

$$\mathcal{L} = -m_D \bar{\psi} \psi = -m_D (\bar{\psi}_L + \bar{\psi}_R) (\psi_L + \psi_R) = -m_D (\bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L). \quad (2.4)$$

The other two terms in the product, $\bar{\psi}_L \psi_L$ and $\bar{\psi}_R \psi_R$, vanish by parity². As a result, giving the neutrino a Dirac mass requires the introduction of right-handed neutrinos and left-handed antineutrinos into the Standard Model. The weak interaction only acts on left-handed neutrinos and right-handed antineutrinos, so these neutrinos would be completely non-interacting, or sterile.

Majorana mass terms, on the other hand, can be constructed without assuming the existence of sterile neutrinos by using the charge-conjugated field ψ^c . In general, the charge conjugation operator changes particle to antiparticle and reverses the sign of the electric charge. Because neutrinos have no charge, we can construct a field $\psi = \psi_L + \psi_L^c$ that satisfies $\psi = \psi^c$; in this framework, neutrinos are equivalent to antineutrinos (Majorana fermions). Importantly, ψ_L^c is right-handed. This allows us to form the Majorana mass terms

$$\mathcal{L} = -\frac{m_L}{2} (\bar{\psi}_L \psi_L^c + \bar{\psi}_L^c \psi_L) - \frac{m_R}{2} (\bar{\psi}_R \psi_R^c + \bar{\psi}_R^c \psi_R) \quad (2.5)$$

that couple particles and antiparticles. Note that two of these terms allow for right-handed (sterile) neutrinos, but they are not required for the generation of mass. Neutrinos are the only neutral Standard Model fermions, so they are the only particle that can have Majorana mass terms, as these terms would otherwise violate charge conservation.

The most general mass term, a combination of Dirac and Majorana mass terms,

2. Let $P_L = (1 + \gamma^5)/2$ be the left-handed projection and $P_R = (1 - \gamma^5)/2$ the right-handed. Then $2P_L \gamma^0 = (1 + \gamma^5) \gamma^0 = \gamma^0 (1 - \gamma^5) = 2\gamma^0 P_R$. This allows us to write $\bar{\psi}_L \psi_L = \psi_L^\dagger \gamma^0 \psi_L = \psi_L^\dagger \gamma^0 P_L \psi_L = \psi_L^\dagger P_R \gamma^0 \psi_L = (P_R \psi_L)^\dagger \gamma^0 \psi_L = 0$, and likewise for $\bar{\psi}_R \psi_R$.

Rest mass of the fundamental fermions

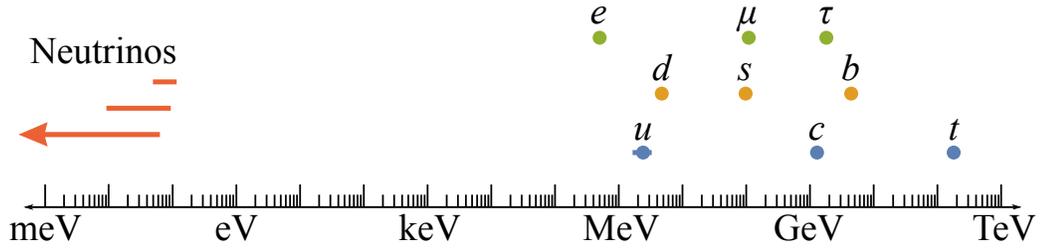


Figure 2.1: Masses of the Standard Model fermions. The absolute value of the neutrino masses are unknown, so a range of possible masses is indicated. The upper limit on the neutrino masses comes from cosmological measurements of the sum of the neutrino masses [58]. Quark and charged lepton masses are from the Particle Data Group [59].

can be expressed as

$$\mathcal{L} = -\frac{1}{2} \begin{pmatrix} \overline{\psi}_L & \overline{\psi}_R^c \end{pmatrix} \mathcal{M} \begin{pmatrix} \psi_L^c \\ \psi_R \end{pmatrix} - \frac{1}{2} \begin{pmatrix} \overline{\psi}_L^c & \overline{\psi}_R \end{pmatrix} \mathcal{M} \begin{pmatrix} \psi_L \\ \psi_R^c \end{pmatrix}, \quad (2.6)$$

with the mass matrix \mathcal{M} given by

$$\mathcal{M} = \begin{pmatrix} m_L & m_D \\ m_D & m_R \end{pmatrix}. \quad (2.7)$$

This gives three possibly distinct terms: two Majorana masses m_L and m_R and a Dirac mass m_D .

2.1.2 Small neutrino masses

From a variety of experiments (see [Section 2.1.3](#)), we know that the neutrino masses are minuscule, significantly smaller than all other known fundamental particles (see [Figure 2.1](#)). A variety of theories exist to explain this imbalance; almost all require that the neutrino have a Majorana mass term.

A small neutrino mass can naturally arise through a process known as the Type-I Seesaw Mechanism. In this process, N right-handed neutrinos are added to the standard model to accompany the three left-handed neutrinos. We extend the mass matrix \mathcal{M} in [Equation 2.7](#) to cover multiple neutrino generations and set $m_L = 0$. This is a plausible scenario, because right-handed neutrinos would be weak isospin singlets (they carry no weak isospin charge) while left-handed neutrinos are doublets, so a Majorana mass term for the left-handed neutrinos would violate weak isospin conservation. This leaves us with the Lagrangian

$$\mathcal{L} = -\frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \bar{\nu}_R^c \end{pmatrix} \begin{pmatrix} 0 & M_D \\ M_D^T & M_R \end{pmatrix} \begin{pmatrix} \nu_L^c \\ \nu_R \end{pmatrix} + \text{h.c.}, \quad (2.8)$$

where ν_L and ν_R are 3-dimensional and N -dimensional vectors of the left-handed and right-handed neutrinos, M_D and M_R are $3 \times N$ and $N \times N$ matrices of Dirac and Majorana masses, and h.c. is the Hermitian conjugate.

To get the mass eigenstates of the neutrino, we diagonalize \mathcal{M} . But first, we make the assumption that the Dirac masses in M_D are much smaller than the Majorana masses in M_R . We expect the Dirac masses to be similar in scale to the masses of the quarks and charged leptons, as both result from the Higgs mechanism. The Majorana masses, on the other hand, are generated independently, and can be at any scale, including the extremely large Grand Unified Theory (GUT) scale [\[60\]](#). The diagonalized form is

$$\mathcal{L} = -\frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \bar{\nu}_R^c \end{pmatrix} U \begin{pmatrix} -M_D M_R^{-1} M_D^T & 0 \\ 0 & M_R \end{pmatrix} U^\dagger \begin{pmatrix} \nu_L^c \\ \nu_R \end{pmatrix} + \text{h.c.} \quad (2.9)$$

after using the approximation for the relative sizes of M_D and M_R [\[61\]](#). This gives 3 small eigenvalues and N large eigenvalues. The mismatch between these eigenvalues (as the heavy neutrino mass eigenvalues increase, the light neutrino mass eigenstate

masses decrease) is what gives the Seesaw Mechanism its name.

We know that the neutrino interaction eigenstates are not the same as the mass eigenstates from experiment, as we have observed neutrinos oscillating in flight. In fact, the upper-left 3×3 portion of matrix U^\dagger in Equation 2.9 is the PMNS matrix that we have measured and which represents the mixing of the 3 observed neutrino flavors.

2.1.3 Neutrino mass observables

Information on the neutrino masses can be obtained only indirectly. In general, there are four complementary approaches to this problem.

The first approach, and the only one to date to have shown unambiguously that neutrinos have nonzero masses, is through the measurement of neutrino oscillation. Oscillation experiments give information on the quantities $\Delta m_{21}^2 (= m_2^2 - m_1^2)$, Δm_{31}^2 , and Δm_{32}^2 . These experiments cannot give any information on the absolute neutrino masses, just the differences between the squares of the masses. In addition, the sign of the mass splitting is only currently observable in effects due to neutrino propagation through matter³. Two mass splittings have been measured: the “solar” mass splitting and the “atmospheric” mass splitting. The solar mass splitting Δm_{21}^2 is measured by a disappearance of electron antineutrinos from nuclear reactors and electron neutrinos from the sun; we know the sign of Δm_{21}^2 due to matter effects in the sun. The atmospheric mass splitting Δm_{32}^2 is measured by a disappearance of muon neutrinos from muon neutrino beams and from the atmosphere; the sign of Δm_{32}^2 is currently unknown.

The ambiguity in the sign of Δm_{32}^2 allows us to define two scenarios: a normal mass hierarchy, in which m_3 is the heaviest neutrino, and an inverted mass hier-

3. As neutrinos propagate through matter, only ν_e and $\bar{\nu}_e$ have charged-current interactions with the electrons in matter. This affects the oscillation probability as it changes the effective potential of ν_e and $\bar{\nu}_e$ with respect to the other neutrino flavors.

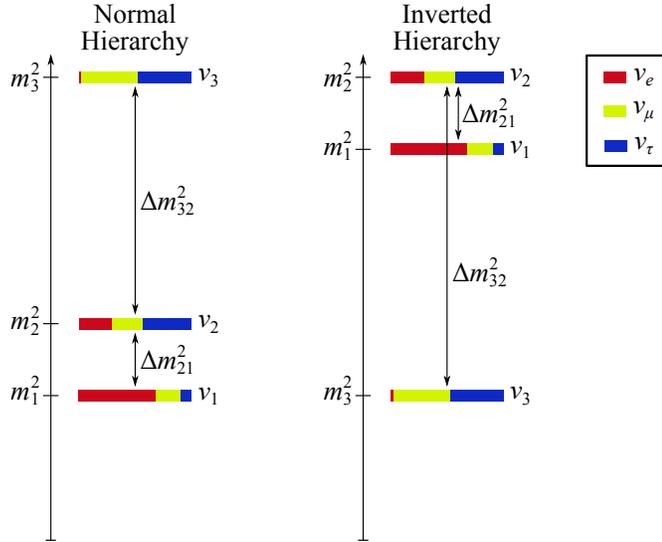


Figure 2.2: Masses and flavor composition of the neutrino mass eigenstates. Oscillation parameters, which determine flavor composition, are from the PDG [59].

archy, in which m_2 is the heaviest. And as it turns out, the solar mass splitting $\Delta m_{21}^2 = (7.53 \pm 0.18) \times 10^{-5} \text{ eV}^2$ is significantly smaller than the atmospheric mass splitting $|\Delta m_{32}^2| = (2.44 \pm 0.06) \times 10^{-3} \text{ eV}^2$ (normal hierarchy) or $|\Delta m_{32}^2| = (2.51 \pm 0.06) \times 10^{-3} \text{ eV}^2$ (inverted hierarchy) [59]. These mass splittings set a lower bound of 50 meV on the mass of the heaviest neutrino mass eigenstate ($\sqrt{\Delta m_{21}^2 + \Delta m_{32}^2}$ in the normal hierarchy and $\sqrt{-\Delta m_{32}^2}$ in the inverted hierarchy, in the limit where the lightest eigenstate is massless). The results of neutrino oscillation experiments are summarized in [Figure 2.2](#).

A second approach to measuring the neutrino masses is to use cosmological measurements. These measurements are insensitive to neutrino flavor and thus measure the quantity $m_{\text{tot}} = \sum m_i$, where the sum is over the light neutrinos. The neutrino masses appear in a variety of cosmological quantities, affecting anisotropies in the cosmic microwave background (CMB) and the formation of large-scale structure in the universe, among other things [62]. Depending on the combination of data used, experiments produce various (model-dependent) limits on m_{tot} . The best claimed and widely accepted is by the Planck Collaboration, $m_{\text{tot}} < 0.23 \text{ eV}$ [58], a combination

of Planck measurements of the CMB temperature, WMAP CMB polarization data, data from various high-resolution CMB measurements, and measurements of baryon acoustic oscillations. This measurement also relies on the accuracy of the standard Λ CDM model; that is, an expanding Universe that obeys the laws of General Relativity and is dominated by cold dark matter (CDM) and a cosmological constant (Λ). Other recent work has suggested that correlation lengths in galaxy clusters exceed that predicted by the Λ CDM model, a result that could be explained by free streaming of light neutrinos with $m_{\text{tot}} = (0.11 \pm 0.03)$ eV [63].

A third approach to neutrino mass measurement is by extremely precise observations of β -decay spectra near its endpoint, just below the Q -value of the decay. In particular, a nonzero neutrino mass means that the emitted electron cannot carry away the full Q -value in energy, as some of the energy is used to create the neutrino. These experiments measure an effective electron neutrino mass of

$$m_\beta = \sqrt{\sum_i |U_{ei}|^2 m_i^2}. \quad (2.10)$$

The world-leading limit in this area comes from the Troitsk and Mainz searches in tritium β decay: $m_\beta < 2.1$ eV (95% C.L.) from Troitsk [64] and $m_\beta < 2.3$ eV (95% C.L.) from Mainz [65]. Results from oscillation experiments indicate that m_β could be as low as 9 meV (normal hierarchy) or 49 meV (inverted hierarchy), so a significant increase in sensitivity could be necessary for a definite discovery.

Finally, we come to $0\nu\beta\beta$ decay. Under the assumption of light Majorana neutrino exchange as the dominant method for this decay, we can deduce a value for the effective Majorana neutrino mass

$$m_{\beta\beta} = \left| \sum_i U_{ei}^2 m_i \right|. \quad (2.11)$$

The value of $m_{\beta\beta}$ is highly dependent on the nuclear matrix elements of the decay, on

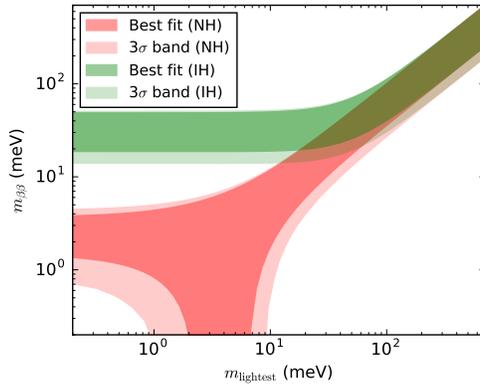
which different models disagree (see [Section 2.2.2](#)). And because the matrix elements U_{ei} are complex, cancellation can occur between the terms of the sum. Oscillation experiments cannot measure α_1 and α_2 , the two Majorana phases in the PMNS matrix ([Equation 2.2](#)), and with all of the phase space for these phases available, it is possible for the value of $m_{\beta\beta}$ to go to zero in the normal hierarchy in the standard 3-neutrino model. In the inverted hierarchy, $m_{\beta\beta}$ is bounded from below.

Comparisons between $m_{\beta\beta}$ and other relevant quantities are shown in [Figure 2.3](#). It is worth noting that the existing of sterile neutrinos of various types can change the shape of these plots dramatically. In particular, the phase space available for $m_{\beta\beta}$ increases, and interference enables $m_{\beta\beta}$ to go to zero even in the inverted hierarchy, depending on the various unmeasured phases. Current and future experiments searching for sterile neutrinos, seeking to determine the mass hierarchy, and measuring neutrino masses should lend some clarity to the situation and limit the available phase space. And in any situation, a measurement of $m_{\beta\beta}$ from $0\nu\beta\beta$ decay outside of the allowed phase space would have profound implications for our understanding of neutrino physics.

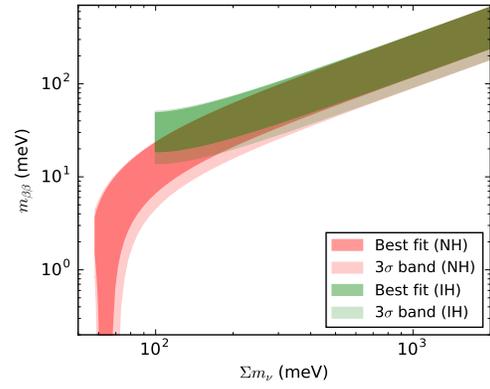
2.2 Neutrinoless double-beta-decay theory

Neutrinoless double-beta ($0\nu\beta\beta$) decay is the hypothetical lepton number-violating decay $(Z, A) \rightarrow (Z + 2, A) + 2e^-$, or double-beta decay without the emission of any neutrinos⁴. Although there are many possible theoretical mechanisms for the decay, the existence of $0\nu\beta\beta$ decay requires that neutrinos are Majorana fermions. This was first argued by Schechter and Valle [[67](#)], who drew a “black box” diagram to demonstrate how $0\nu\beta\beta$ decay by any mechanism couples neutrinos to antineutrinos and thus generates a small but non-zero Majorana neutrino mass (see [Figure 2.4](#)). It

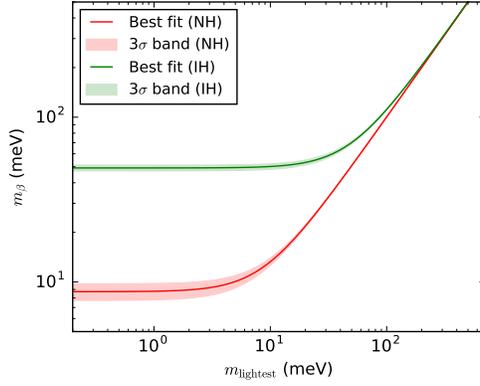
4. The family of double-beta decays also includes the analogous $\beta^+\beta^+$ decay and double electron capture, but here we discuss only $\beta^-\beta^-$ decay.



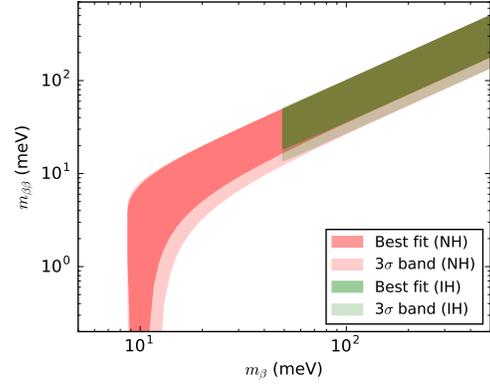
(a) $m_{\beta\beta}$ vs. the lightest neutrino mass



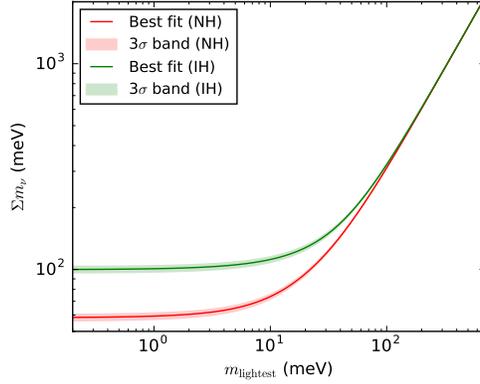
(b) $m_{\beta\beta}$ vs. the sum of the neutrino masses



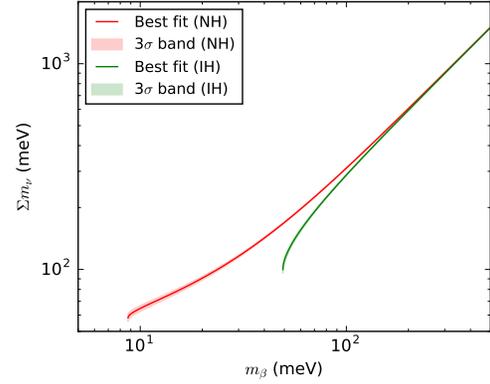
(c) m_β vs. the lightest neutrino mass



(d) $m_{\beta\beta}$ vs. m_β



(e) Sum of the the neutrino masses vs. the lightest neutrino mass



(f) m_β vs. the sum of the the neutrino masses

Figure 2.3: Comparisons of various measurements of the neutrino masses in the normal hierarchy (NH) and inverted hierarchy (IH), using the standard 3-neutrino model. The width of the best fit band for $m_{\beta\beta}$ comes from the unknown Dirac and Majorana phases, and the 3σ uncertainty bands come from uncertainties in the known oscillation parameters. Oscillation parameters are from the PDG [66], and the calculations are based on code by K. Han.

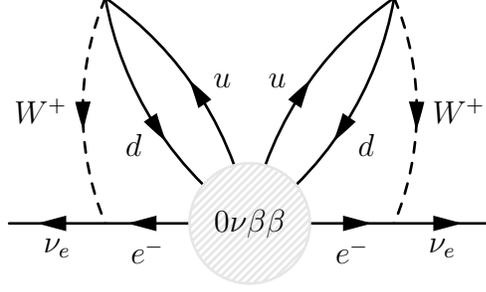


Figure 2.4: Black box diagram for the generation of a Majorana neutrino mass, based on a proposal by Schechter and Valle [67].

is important to note that the neutrino mass not need be principally generated through this mechanism, just that the mass generated by this mechanism is non-zero.

Experiments measure only the half-life of the decay, but to interpret this result, we break the half-life $T_{1/2}^{0\nu}$ down into three components,

$$(T_{1/2}^{0\nu})^{-1} = G_{0\nu} |M_{0\nu}|^2 |f(m_i, U_{ei})|^2, \quad (2.12)$$

where $G_{0\nu}$ is the phase space factor, $M_{0\nu}$ is the nuclear matrix element, and $f(m_i, U_{ei})$ is a function of the neutrino masses and mixing matrix that represents the physical process responsible for the decay. We discuss each one of these factors in turn.

2.2.1 Phase space

The phase space factor represents the available phase space for the decay and encodes the kinematics of the decay. To express the phase space factor for $0\nu\beta\beta$ decay analytically, we introduce the terms

$$w_{0\nu} = \frac{g_A^4 (G \cos \theta_C)^4}{16\pi^5} (m_e c^2)^2 (\hbar c^2) (p_1 c) (p_2 c) \epsilon_1 \epsilon_2, \quad (2.13)$$

where g_A is the axial-vector coupling constant, G is the Fermi coupling constant, θ_C is the Cabibbo angle, m_e is the electron mass, and p_1 , p_2 , ϵ_1 and ϵ_2 are the energy

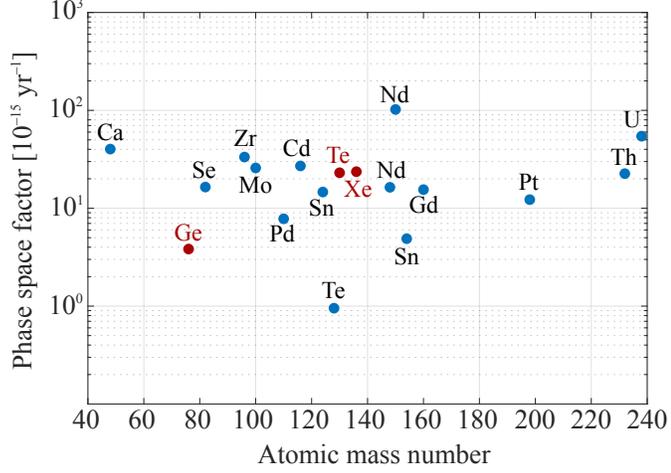


Figure 2.5: Phase space factors for a variety of $0\nu\beta\beta$ decay candidates. In red are ^{76}Ge , ^{130}Te , and ^{136}Xe , isotopes under investigation by some leading $0\nu\beta\beta$ decay searches. Data from calculations by Kotila and Iachello [68].

and momenta of the two emitted electrons; and

$$f_{11}^{(0)} = |g_{-1}(\epsilon_1)g_{-1}(\epsilon_2)|^2 + |f_1(\epsilon_1)f_1(\epsilon_2)|^2 + |g_{-1}(\epsilon_1)f_1(\epsilon_2)|^2 + |f_1(\epsilon_1)g_{-1}(\epsilon_2)|^2, \quad (2.14)$$

where g_{-1} and f_1 are solutions to the radial Dirac equations with energy ϵ and relativistic quantum number ± 1 evaluated at the nuclear radius $R = (1.2 \text{ fm})A^{1/3}$. The phase space factor $G_{0\nu}$ is then represented as

$$G_{0\nu} = \frac{1}{(2 \ln 2)g_A^4 R^2} \int_{m_e c^2}^{Q_{\beta\beta} + m_e c^2} f_{11}^{(0)} w_{0\nu} d\epsilon_1, \quad (2.15)$$

where $Q_{\beta\beta}$ is the Q -value of the decay [68]. As is clear from the integral, the phase space increases with $Q_{\beta\beta}$, so experimental searches for $0\nu\beta\beta$ decay become increasingly difficult for decays with lower Q -values.

A plot of the phase space factors for various $0\nu\beta\beta$ decay candidates is shown in [Figure 2.5](#).

2.2.2 Nuclear matrix elements

The nuclear matrix elements (NMEs) are significantly harder to calculate than the phase space factors and are highly model-dependent. The NME for the decay ${}^A_Z X_N \rightarrow {}^A_{Z+2} Y_{N-2} + 2e^-$ is

$$M_{0\nu} = \langle {}^A X; 0_1^+ | H | {}^A Y; J_F \rangle, \quad (2.16)$$

where H is the weak interaction Hamiltonian and J_F is the angular momentum of the final state. This can be broken up into three components,

$$M_{0\nu} = g_A^2 \left[M_{GT}^{(0\nu)} - \left(\frac{g_V}{g_A} \right)^2 M_F^{(0\nu)} + M_T^{(0\nu)} \right], \quad (2.17)$$

which represent the Gamow-Teller, Fermi, and tensor contributions. For $0\nu\beta\beta$ decay, the ratio of the first to second terms varies by element and by nuclear model, but the Gamow-Teller term is generally larger by a factor of 2–5. There is disagreement between models on the size of the tensor contribution, but it is generally evaluated to be similar than the Fermi term.

A variety of nuclear structure models are used to estimate the NMEs. One model for evaluating the NMEs is the Interacting Shell Model (ISM) [69]. In the ISM, the nucleus is constructed as a collection of fermions that obey the Pauli Exclusion Principle. The basis states are those of the harmonic oscillator with perturbations, and the many-body wave function is composed of a large number of Slater Determinants. This is in contrast to the Interacting Boson Model (IBM) [70]. The IBM, which is useful for even-even nuclei like the $0\nu\beta\beta$ decay candidates, considers pairs of protons or neutrons as bosons instead of as individual fermions. A different approach for evaluating NMEs is the Quasi-Random Phase Approximation (QRPA) [71]. Although the ISM includes all possible shell configurations, it sums over a small number of state energies, due the intense computational requirements. The QRPA, however, uses particle-hole pair and quasiparticle dynamics to include a much larger number

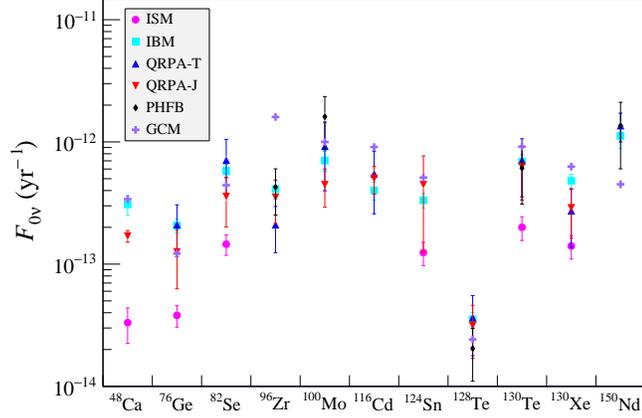


Figure 2.6: Figure of merit for various $0\nu\beta\beta$ decay candidate isotopes, as defined in Equation 2.18. Although ^{130}Te and ^{128}Te appear with roughly equal abundance in natural tellurium, ^{130}Te is predicted to decay an order of magnitude faster. ^{130}Te also has a higher figure of merit than ^{76}Ge and ^{130}Xe , isotopes used in other leading $0\nu\beta\beta$ decay experiments. All calculations are under the assumption of the light Majorana neutrino exchange mechanism. Phase space factors are from Kotila and Iachello [68]. NMEs are presented in a variety of models, labelled ISM [69, 73], IBM [74], QRPA-T [75], QRPA-J [76], PHFB [77], and GCM [78]. Figure from J. Ouellet [79].

of energy states, but in fewer shell configurations.

An outstanding issue is the value of g_A . Most calculations (and all those presented here) use the free values $g_V = 1$ and $g_A = 1.269$ (or similar). Yet comparisons between measured half-lives for β decay and half-lives predicted by the nuclear models indicate that g_A may be significantly quenched in heavy nuclei [72]. Because the decay rate is proportional to g_A^4 , quenching could significantly decrease the decay rate for a given effective Majorana neutrino mass, making it more difficult to observe $0\nu\beta\beta$ decay.

To compare isotopes for $0\nu\beta\beta$ decay searches, it is useful to combine the phase space factors and nuclear matrix elements into an isotope figure of merit

$$F_{0\nu} = G_{0\nu} |M_{0\nu}|^2, \quad (2.18)$$

which is directly proportional to the decay rate (Equation 2.12). This figure of merit for a variety of isotopes in various nuclear models is presented in Figure 2.6.

2.2.3 Light Majorana neutrino exchange

The simplest mechanism for $0\nu\beta\beta$ decay is light Majorana neutrino exchange (Figure 1.2). In this process, a light Majorana neutrino is effectively emitted by one neutron and absorbed by the other. The factor $f(m_i, U_{ei})$ from Equation 2.12 in this scenario is given as

$$f(m_i, U_{ei}) = \frac{m_{\beta\beta}}{m_e} = \frac{1}{m_e} \left| \sum_i U_{ei} m_i \right|, \quad (2.19)$$

where $m_{\beta\beta}$ is the effective Majorana neutrino mass and m_e is included as a normalization constant so that $f(m_i, U_{ei})$ is unitless.

The form of $f(m_i, U_{ei})$ tells us that the decay rate is directly proportional to the square of the effective Majorana neutrino mass. To understand why this is the case, we must examine the chirality and helicity of the exchanged neutrino. When the neutrino is emitted, it must have left-handed chirality; when it is absorbed, it has right-handed chirality. At the same time, helicity is conserved in the process. If neutrinos were massless, this would be impossible. Because they have mass, however, a neutrino emitted with left-handed chirality has a small right-handed helicity component. Thus, the relevant quantity is the probability that the emitted left-handed neutrino has right-handed helicity.

To calculate this probability, we first look at the helicity solutions to the Dirac equation. We use the Dirac basis, with

$$\gamma^0 = \begin{pmatrix} I_2 & 0 \\ 0 & -I_2 \end{pmatrix}, \quad \gamma^k = \begin{pmatrix} 0 & \sigma^k \\ -\sigma^k & 0 \end{pmatrix}, \quad \text{and} \quad \gamma^5 = \begin{pmatrix} 0 & I_2 \\ I_2 & 0 \end{pmatrix}. \quad (2.20)$$

Then we have the right-handed helicity solution u_+ and left-handed helicity solution

u_- , with

$$u_+ = N \begin{pmatrix} \cos \frac{\theta}{2} \\ e^{i\phi} \sin \frac{\theta}{2} \\ \frac{|\vec{p}|}{E+m} \cos \frac{\theta}{2} \\ \frac{|\vec{p}|}{E+m} e^{i\phi} \sin \frac{\theta}{2} \end{pmatrix} \quad \text{and} \quad u_- = N \begin{pmatrix} -\sin \frac{\theta}{2} \\ e^{i\phi} \cos \frac{\theta}{2} \\ \frac{|\vec{p}|}{E+m} \sin \frac{\theta}{2} \\ -\frac{|\vec{p}|}{E+m} e^{i\phi} \cos \frac{\theta}{2} \end{pmatrix}, \quad (2.21)$$

where N is a normalization constant. We can then use the chiral projections $P_R = \frac{1}{2}(1 + \gamma^5)$ and $P_L = \frac{1}{2}(1 - \gamma^5)$ to obtain

$$P_R u_+ = \frac{N}{2} \left(1 + \frac{|\vec{p}|}{E+m} \right) \begin{pmatrix} \cos \frac{\theta}{2} \\ e^{i\phi} \sin \frac{\theta}{2} \\ \cos \frac{\theta}{2} \\ e^{i\phi} \sin \frac{\theta}{2} \end{pmatrix} = \frac{N}{2} \left(1 + \frac{|\vec{p}|}{E+m} \right) u_R \quad (2.22)$$

and

$$P_L u_+ = \frac{N}{2} \left(1 - \frac{|\vec{p}|}{E+m} \right) \begin{pmatrix} \cos \frac{\theta}{2} \\ e^{i\phi} \sin \frac{\theta}{2} \\ -\cos \frac{\theta}{2} \\ -e^{i\phi} \sin \frac{\theta}{2} \end{pmatrix} = \frac{N}{2} \left(1 - \frac{|\vec{p}|}{E+m} \right) u_L. \quad (2.23)$$

where we substitute in the chiral states u_R and u_L , and analogously,

$$P_R u_- = \frac{N}{2} \left(1 - \frac{|\vec{p}|}{E+m} \right) u_R \quad (2.24)$$

and

$$P_L u_- = \frac{N}{2} \left(1 + \frac{|\vec{p}|}{E+m} \right) u_L. \quad (2.25)$$

Then can write

$$u_{\pm} = \frac{N}{2} \left(1 \pm \frac{|\vec{p}|}{E+m} \right) u_R + \frac{N}{2} \left(1 \mp \frac{|\vec{p}|}{E+m} \right) u_L, \quad (2.26)$$

which leads to

$$u_L = \frac{1}{2N|\vec{p}|} [(E+m+|\vec{p}|)u_- - (E+m-|\vec{p}|)u_+]. \quad (2.27)$$

As expected, if the particle is relativistic ($E \gg m$), then $E \approx |\vec{p}|$ and $N \approx 1$, and our left-handed chirality particle is almost entirely in a left-handed helicity state. More precisely, we can expand the u_+ component for $E \gg m$ and substitute $p = \sqrt{E^2 - m^2} \approx E - m^2/2E$, giving

$$\frac{E+m-|\vec{p}|}{2N|\vec{p}|} \rightarrow \frac{m+m^2/2E}{2E-m^2/2E} \rightarrow \frac{m}{2E}. \quad (2.28)$$

Thus, the probability of the neutrino with left-handed chirality having right-handed helicity is proportional to m^2/E^2 . Because of neutrino mixing, m is really an effective mass ($m_{\beta\beta}$), and indeed the decay rate (Equation 2.12) is proportional to $m_{\beta\beta}^2$ in this framework.

2.3 Searching for neutrinoless double-beta decay

Because $0\nu\beta\beta$ decay is such a rare process, all modern searches for $0\nu\beta\beta$ decay have similar themes. In essence, these experiments look for an excess of events with energies right at the Q -value of the decaying isotope, in a region of interest determined by the energy resolution of the detectors. It is common to speak of a sensitivity figure of merit for $0\nu\beta\beta$ decay searches, which is proportional to the true half-life sensitivity

under various assumptions:

$$\text{Sensitivity} \propto a \sqrt{\frac{Mt}{b\Delta E}}, \quad (2.29)$$

where a is the isotopic abundance of the double-beta emitter, M is the mass of the detector, t is the exposure time, b is the background index (the number of background events in a region around the Q -value divided by the size of the energy region, exposure time and detector mass), and ΔE is the energy resolution. This model is for experiments in which the source is embedded in the detector, a common setup for many $0\nu\beta\beta$ decay experiments; the total background is then taken to be proportional to M and the total $0\nu\beta\beta$ -decaying isotope mass is taken to be proportional to aM .

The first term in the sensitivity figure of merit is the isotopic abundance a of the $0\nu\beta\beta$ decay candidate isotope. In addition to the isotope figure of merit (Figure 2.6), the natural isotopic abundance can impact the choice of isotope, as isotopes with low natural abundance require costly and difficult enrichment. The natural isotopic abundance for a variety of $0\nu\beta\beta$ decay candidate isotopes is shown in Figure 2.7. There are two clear outliers on the plot, ^{128}Te and ^{130}Te ; all other isotopes essentially require isotopic enrichment to achieve a large source mass.

Another term in the sensitivity figure of merit is the background index b , which is measured in units of counts/(keV·kg·yr); that is, how many background events are expected in a 1-keV range in one year of data-taking per kilogram of detector mass. The background index relevant to the sensitivity is the background index at the Q -value. All experiments follow similar approaches to reduce the background levels, including using ultra-radiopure materials in the detector, using some sort of active or passive detector shielding, and going deep underground to avoid cosmic rays. But the $0\nu\beta\beta$ decay isotope choice also influences the background index through its Q -value, since natural radioactive backgrounds decrease significantly at higher energies. This

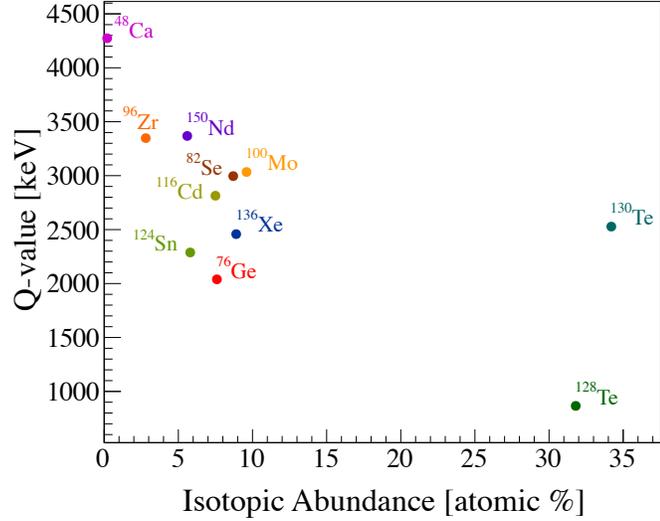


Figure 2.7: Q -values and natural isotopic abundance for a variety of $0\nu\beta\beta$ decay candidate isotopes. Figure courtesy of the CUORE Collaboration.

is another advantage of ^{130}Te over ^{76}Ge and ^{136}Xe , the other two isotopes used most in present-day $0\nu\beta\beta$ decay searches. While ^{48}Ca has the highest Q -value, its natural isotopic abundance is so low that a large-scale experiment with ^{48}Ca is infeasible due to the costs of the isotopic enrichment.

The other terms in the sensitivity figure of merit are the detector mass M , exposure time t , and energy resolution ΔE . The ability to scale up an experiment to large masses, operate it stably for long periods of time (generally several years), and achieve a good energy resolution is highly linked with the particular detector technology used. As a general principle, the experiments that are easiest to scale to large masses (liquid and gas detectors) tend to have poorer energy resolution, while solid crystal- or bolometer-based experiments tend to have superior energy resolution but require many individual segmented detectors in order to achieve a large mass.

A more complete discussion of the experimental sensitivity is in [Section 3.5](#) below.

Experiment	Isotope	Detector technology	Half-life limit	Iso. exposure
GERDA	^{76}Ge	Ionization	$> 5.3 \times 10^{25}$ yr [80]	34 kg·yr
NEMO-3	^{100}Mo	Tracker, calorimeter	$> 1.1 \times 10^{24}$ yr [81]	35 kg·yr
CUORE-0	^{130}Te	Bolometers	$> 4.0 \times 10^{24}$ yr [82]	30 kg·yr
EXO-200	^{136}Xe	Liquid TPC	$> 1.1 \times 10^{25}$ yr [83]	100 kg·yr
KamLAND-Zen	^{136}Xe	Scintillation	$> 1.1 \times 10^{26}$ yr [84]	504 kg·yr

Table 2.1: Recent results from selected $0\nu\beta\beta$ decay experiments. Half-life limits are shown at 90% C.L. Exposure shown is the isotopic mass exposure.

Experiment	Isotope	Detector technology	Sensitivity	Iso. mass	Start
GERDA (Phase II)	^{76}Ge	Ionization	1×10^{26} yr [80]	30 kg	2016
MAJORANA DEMO.	^{76}Ge	Ionization	2×10^{26} yr [85]	26 kg	2016
SuperNEMO	^{82}Se	Tracker, calorimeter	1×10^{26} yr [86]	100 kg	2020?
CUORE	^{130}Te	Bolometers	9×10^{25} yr [87]	206 kg	2017
SNO+	^{130}Te	Scintillation	9×10^{25} yr [88]	800 kg	2018?
EXO-200 (Phase II)	^{136}Xe	Liquid TPC	6×10^{25} yr [89]	76 kg	2016
NEXT-100	^{136}Xe	Gas TPC	6×10^{25} yr [90]	90 kg	2018?
PandaX-III	^{136}Xe	Gas TPC	1×10^{26} yr [91]	180 kg	2019?
KamLAND-Zen	^{136}Xe	Scintillation	2×10^{26} yr [92]	600 kg	2016

Table 2.2: Final projected (3-, 4-, or 5-year) sensitivity at 90% C.L. of current and upcoming $0\nu\beta\beta$ decay experiments, as reported by each experiment. Mass shown is the mass of the $0\nu\beta\beta$ decay candidate isotope. Start dates in the future are necessarily estimates and are indicated by a question mark.

2.3.1 Current efforts

There are several active experiments in the field investigating a variety of isotopes with several different detection techniques. Table 2.1 contains a summary of recent $0\nu\beta\beta$ decay search results, and Table 2.2 shows the projected sensitivity of current and upcoming $0\nu\beta\beta$ decay searches.

One area of active research is the use of germanium semiconductor detectors to measure the energy of electrons released inside the detectors by the $0\nu\beta\beta$ decay of ^{76}Ge . The GERDA collaboration uses high-purity germanium detectors enriched to 86% in ^{76}Ge and immersed in a large volume of liquid argon, which is used for shielding and cooling. GERDA has the current best limit on the ^{76}Ge half-life, $T_{1/2}^{0\nu\beta\beta} > 5.3 \times 10^{25}$ yr (90% C.L.), with an exposure of 34 kg·yr [80]. Another

experiment with similar sensitivity, the MAJORANA DEMONSTRATOR, uses p-type point-contact detectors to search for $0\nu\beta\beta$ decay due to their ability to discriminate between single-site events and multi-site backgrounds. A new collaboration, LEGEND, has been formed with members of both GERDA and Majorana to produce a ton-scale experiment using a phased approach in the future [93].

Two collaborations currently have comparable sensitivities on the ^{136}Xe $0\nu\beta\beta$ decay half-life: EXO-200 and KamLAND-Zen. EXO-200 is a cylindrical time projection chamber filled with liquid xenon enriched to 80.6% in ^{136}Xe . The xenon is also surrounded by avalanche photodiodes that measure scintillation light; this combination allows for event discrimination between single-site decays and broader Compton scattering of gamma ray backgrounds. The KamLAND-Zen experiment suspends a balloon of xenon-doped liquid scintillator into the existing large liquid scintillator detector built for the KamLAND reactor neutrino experiment. EXO-200 has reported a ^{136}Xe half-life limit of $T_{1/2}^{0\nu\beta\beta} > 1.1 \times 10^{25}$ yr (90% C.L.) with an exposure of 100 kg·yr [83], while KamLAND-Zen has reported $T_{1/2}^{0\nu\beta\beta} > 1.1 \times 10^{26}$ yr (90% C.L.) with an exposure of 504 kg·yr [84]. Both KamLAND-Zen and EXO-200 have been upgraded and continue to take data. The EXO-200 collaboration is also planning a new experiment, nEXO, with significantly more ^{136}Xe mass. Two other collaborations, NEXT and PandaX-III, are planning experiments with high-pressure xenon gas TPCs to begin operations in 2018 or 2019.

The two collaborations dedicated to studying ^{130}Te decay are CUORE and SNO+. The SNO+ experiment will use tellurium-loaded liquid scintillator, similar to KamLAND-Zen, to search for $0\nu\beta\beta$ decay of ^{130}Te and is currently in commissioning without tellurium [88]. It will take advantage of the hardware used in the SNO experiment, which was filled with heavy water to detect solar neutrinos. Although SNO+ can more easily use very large masses (the first generation of SNO+ will use approximately 800 kg of ^{130}Te), the limited energy resolution of

liquid scintillators likely limits its first-run half-life sensitivity to be below that of CUORE, the experiment that is the subject of this dissertation.

A qualitatively different experiment, NEMO-3, is the only major $0\nu\beta\beta$ decay experiment where the source isotope is not embedded in the detector. The best half-life limits for many $0\nu\beta\beta$ decay candidate isotopes not addressed by dedicated experiments have come from the NEMO-3 detector, due to its ability to study any isotope that can be added to a thin source film. These source films are placed inside a gas tracking chamber that is surrounded by calorimeters, which allows for both electron tracking and energy reconstruction and thus extremely low backgrounds. The experiment is, however, limited by the relatively poor energy resolution of the calorimeters ($\sim 10\%$) and the low source masses that can be used. The NEMO collaboration is currently scaling up to produce SuperNEMO, a modular detector aimed at improving energy resolution and signal efficiency [86].

Several factors distinguish CUORE from other $0\nu\beta\beta$ decay searches. The relatively high natural abundance of ^{130}Te (significantly greater than that of ^{76}Ge and ^{136}Xe) makes CUORE competitive without the need to enrich the TeO_2 crystals in the candidate isotope. The Q -value of ^{130}Te decay, 2528 keV [94–96], is above the Compton edge of the dominant 2615 keV background line from ^{208}Tl , resulting in low backgrounds in the $0\nu\beta\beta$ decay region of interest. The high Q -value also results in a favorable phase space factor for the decay, reducing the expected half-life compared to other isotopes. Finally, the high resolution of TeO_2 bolometers — approximately 5 keV at the $0\nu\beta\beta$ decay Q -value and the second-best resolution of any $0\nu\beta\beta$ decay technology after germanium detectors — provides for excellent separation of nearby background lines from the region of interest and greatly reduces the unavoidable background from ordinary (two-neutrino) double-beta decay. All of these factors make CUORE one of the most promising of today's $0\nu\beta\beta$ decay experiments.

At the same time, the wide variety of experiments with different isotopes and de-

tection techniques is extremely important to the field. At the very least, we will need to verify any claimed discovery of $0\nu\beta\beta$ decay with other experiments. Verifications with the same isotope and different techniques can ensure that we fully understand our detectors, in particular their energy reconstruction and detection efficiency. And the discovery of $0\nu\beta\beta$ decay with different isotopes can provide valuable information to nuclear physics theorists, helping us reduce the uncertainties and disagreements in the nuclear matrix elements between various nuclear models.

All of these experiments are driven to search for $0\nu\beta\beta$ decay because of the profound implications of such a discovery for our understanding of physics. The conclusive proof that neutrinos are Majorana particles and the direct observation lepton number violation would alter the Standard Model of particle physics. And using the nuclear matrix elements and phase space factors calculated for various isotopes, the half-life measured in $0\nu\beta\beta$ decay can give us direct information about the absolute mass scale of the lightest neutrino, which is as-of-yet unknown.

Chapter 3

CUORE

“Science is curiosity. We all have natural curiosity. Science is a process of investigating. It’s posing questions and coming up with a method. It’s delving in.”

— Sally Ride

CUORE, the Cryogenic Underground Observatory for Rare Events, is an experiment designed to search for neutrinoless double-beta ($0\nu\beta\beta$) decay of ^{130}Te and other rare events. CUORE uses TeO_2 crystals both as the source of ^{130}Te and as bolometers to detect its decay. The crystals are operated in a dilution refrigerator deep underground, at a depth of 3600 m.w.e.¹ at the Laboratori Nazionali del Gran Sasso (LNGS) in Assergi, Italy. The underground location limits the experiment’s exposure to cosmic rays, which can be a significant source of background events in a rare events search.

We discuss below the general principles behind CUORE, before reviewing the history of the CUORE family of detectors and the design of CUORE itself.

1. The unit m.w.e. (meters of water equivalent) is used to measure the natural shielding of underground laboratories. In this case, the mountain above the laboratory provides shielding equivalent to that provided by 3600 m of water.

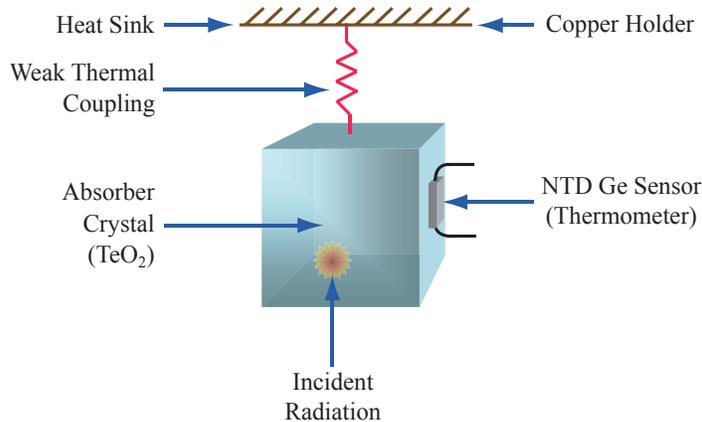


Figure 3.1: Schematic of a CUORE bolometer. Figure courtesy of the CUORE collaboration.

3.1 Detectors

A bolometer is a device that is weakly coupled to a thermal reservoir and that heats up temporarily as it absorbs energy from an incident particle. In general, bolometers consist of an energy-absorbing material strongly coupled to a temperature-dependent resistor. The voltage across this resistor, a proxy for the temperature of the absorber, is then used to determine the energy of the incident particle.

In CUORE, the bolometers are TeO_2 crystals coupled to neutron-transmutation-doped (NTD) germanium temperature-dependent resistors (thermistors) [97]. Neutron transmutation doping allows us to achieve a very uniform dopant concentration and is critical to the high resolution of these thermistors. The $5 \times 5 \times 5 \text{ cm}^3$ crystals are held at the corners by polytetrafluoroethylene (PTFE) supports inside copper frames. The copper frames are thermally linked to the mixing chamber of a dilution refrigerator and held at 15 mK (see Section 3.4). A schematic of the CUORE bolometer setup is in Figure 3.1.

When a particle deposits energy into a TeO_2 crystal, the crystal heats up accordingly. Over the course of a few seconds, the temperature relaxes back to its baseline

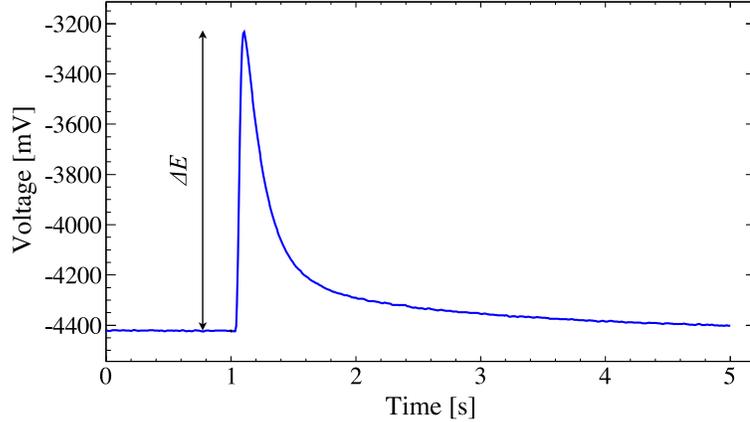


Figure 3.2: A representative NTD thermistor pulse. We use the pulse height to determine the amount of energy deposited in the crystal, while the baseline voltage before the pulse tells us the steady-state temperature of the detector.

value as the energy passes through the PTFE supports² and into the thermal bath (the copper frame and the mixing chamber of the dilution refrigerator). The change in temperature of the crystal is measured by a change in voltage across the thermistor. The rise time of the voltage pulse from the thermistor (defined as the time to rise from 10% to 90% of the maximum voltage height) is approximately 100 ms. The fall time is much longer; it is about 400 ms from 90% to 30% of the maximum, and the pulse takes several seconds to decay completely. Each of the bolometers is unique, and the pulse rise and decay times vary by more than 50% between bolometers. In any case, the NTD thermistor voltage is recorded continuously at 1000 samples/second, so the full pulse shape is acquired accurately (see Figure 3.2).

The total temperature rise of the TeO_2 crystal is determined by its heat capacity. These crystals follow the Debye law very well at temperatures up to ~ 250 mK, with a heat capacity C given by

$$C = amT^3. \quad (3.1)$$

In dedicated tests, a was measured to be $(2.93 \pm 0.04) \times 10^{-3}$ J/(kg K⁴) in TeO_2

2. Some heat is also transferred through the gold wire leads for the thermistor; see Section 3.1.1.

crystals [98]. For the CUORE crystals ($m = 750$ g), this gives a heat capacity of $C = 7.4 \times 10^{-9}$ J/K at 15 mK. This corresponds to a rise of 0.02 mK (0.1% of the absolute temperature) per MeV of energy deposited.

The temperature rise in the crystal is measured by the NTD Ge thermistor glued to its surface. To produce these thermistors, neutrons are fired at germanium wafers, creating acceptor (Ga) and donor (As and Se) impurities. At sufficiently cold temperatures (below approximately 1 K), conduction is dominated by the hopping of electrons from one donor site to another, and the resistance R is expressed as

$$R = R_0 \exp(\sqrt{T_0/T}), \quad (3.2)$$

where T is temperature and R_0 and T_0 are empirically-determined constants that depend on the doping [99]. Typical values for the CUORE bolometers are $R_0 \sim 1 \Omega$ and $T_0 \sim 4$ K.

Because the heat capacity of the crystals and the resistance of the thermistors are extremely temperature dependent, the pulse height that we record depends on both the energy deposited in the crystal and the temperature-dependent gain of the detector. In order to remove the dependence on the baseline temperature when calculating the pulse heights, and thus determine the particle energies, we need a mapping from baseline temperature to gain for each bolometer. To this end, each bolometer also contains a resistive heater, to which we periodically send a fixed amount of energy³.

The time it takes for the temperature of the crystal to return to its baseline value, after it absorbs the energy of an incident particle, is determined by the heat capacity C of the crystal and the thermal conductance G of the links between the crystal and the thermal bath. Their ratio is the time constant τ for the return of the detector

3. Some channels don't have functioning heaters. For those, we can produce a mapping from baseline temperature to gain by using known-energy calibration lines. We will return to this discussion in [Section 5.2.3](#).

temperature to its baseline, or

$$\tau = C/G. \tag{3.3}$$

There are two paths for the heat to dissipate through: the PTFE supports on the corners of the crystal and the (very thin) gold wires connecting the NTD thermistor and heater to the wiring trays. Based on our measurements of the detector pulse shape and the observed time constant $\tau \sim 1$ s, we empirically determine $G \sim 2 \times 10^{-9}$ W/K, or in different units, $G \sim 10$ MeV/(mK s).

3.1.1 Detector response function

In order to determine the detector response function (i.e., the signal seen when energy is deposited in the detector), we model a CUORE-style bolometer as a thermal circuit, following the approach of M. Vignati [100]. The circuit is broken into three parts: the TeO₂ crystal, PTFE supports, and Ge thermistor; these are shown in [Figure 3.3](#). The crystal has a heat capacity C_c ; it is connected to the supports with a thermal contact resistance R_{cs} and to the thermistor with a thermal resistance R_g (incorporating contact resistance and the resistance of the glue itself). The supports have a heat capacity C_s and a thermal resistance R_s to the mixing chamber of the dilution refrigerator, which we consider to be a constant-temperature heat bath. The thermistor has a phonon thermal resistance R_p and electron gas thermal resistance R_e to the heat bath. It also has a heat capacity C_e from the electron gas⁴. Because the heat capacity from the phonons is proportional to T^3 (Debye model) while the heat capacity from the electrons is proportional to T (Fermi-Dirac statistics), the heat capacity from phonons is comparatively negligible at extremely low temperatures.

Using power conservation at the four nodes indicated in [Figure 3.3](#), we can write

4. In reality, there is another heat load on the thermistors due to the steady-state current used to measure their voltage. This has the effect of slightly modifying the effective values of R_e and C_e from their nominal values [100].

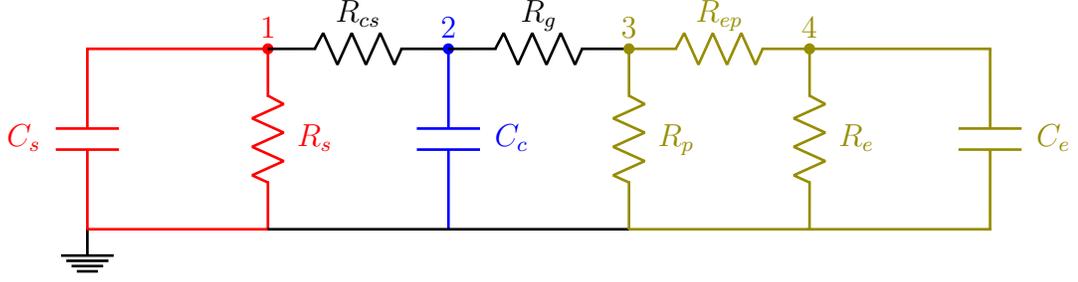


Figure 3.3: Thermal circuit of a CUORE bolometer. In red are the crystal supports (PTFE), in blue is the crystal itself (TeO_2), and in olive is the thermistor (Ge). The thermistor thermal capacity and resistance is broken down into conductance through the gold wires with phonon thermal resistance R_p and electron gas thermal resistance R_e , and the capacitance of the electron gas C_e .

the differential equations

$$\begin{aligned}
 C_s \frac{dT_1}{dt} + \frac{T_1}{R_s} + \frac{T_1 - T_2}{R_{cs}} &= 0 & C_c \frac{dT_2}{dt} + \frac{T_2 - T_1}{R_{cs}} + \frac{T_2 - T_3}{R_g} &= 0 \\
 \frac{T_3}{R_p} + \frac{T_3 - T_2}{R_g} + \frac{T_3 - T_4}{R_{ep}} &= 0 & C_e \frac{dT_4}{dt} + \frac{T_4}{R_e} + \frac{T_4 - T_3}{R_{ep}} &= 0.
 \end{aligned} \quad (3.4)$$

Note that all of these temperatures are relative temperatures above a base temperature T_0 . Using Laplace transforms, we write these equations as

$$\begin{aligned}
 C_s [sT_1 - T_1(0)] + \frac{T_1}{R_s} + \frac{T_1 - T_2}{R_{cs}} &= 0 & C_c [sT_2 - T_2(0)] + \frac{T_2 - T_1}{R_{cs}} + \frac{T_2 - T_3}{R_g} &= 0 \\
 \frac{T_3}{R_p} + \frac{T_3 - T_2}{R_g} + \frac{T_3 - T_4}{R_{ep}} &= 0 & C_e [sT_4 - T_4(0)] + \frac{T_4}{R_e} + \frac{T_4 - T_3}{R_{ep}} &= 0.
 \end{aligned} \quad (3.5)$$

We recognize $C_s T_1(0)$, $C_c T_2(0)$, and $C_e T_4(0)$ as the energy in the three capacitors at time $t = 0$. To model an energy deposition E in the TeO_2 crystal, we set $C_c T_2(0) = E$ and $C_s T_1(0) = C_e T_4(0) = 0$. Simplifying the equations above and using thermal

conductivity $K = 1/R$ for each element, this leaves

$$(K_{cs} + K_s + C_s s)T_1 - K_{cs}T_2 = 0 \quad (3.6a)$$

$$-K_{cs}T_1 + (K_{cs} + K_g + C_c s)T_2 - K_g T_3 = E \quad (3.6b)$$

$$-K_g T_2 + (K_{ep} + K_g + K_p)T_3 - K_{ep}T_4 = 0 \quad (3.6c)$$

$$-K_{ep}T_3 + (K_e + K_{ep} + C_e s)T_4 = 0. \quad (3.6d)$$

The full solution to this system of equations is quite complex. We are, however, only really interested in T_4 , the signal in the thermistor. Eliminating T_1 , T_2 , and T_3 from this system, we can then write T_4 as

$$\begin{aligned} & [(K_{cs} + K_g + C_c s)(K_{cs} + K_s + C_s s) - K_{cs}^2] \times \\ & [(K_{ep} + K_g + K_p)(K_e + K_{ep} + C_e s) - K_{ep}^2] \\ & - K_g^2(K_e + K_{ep} + C_e s)(K_{cs} + K_s + C_s s) = \frac{EK_g K_{ep}(K_{cs} + K_s + C_s s)}{T_4} \end{aligned} \quad (3.7)$$

The left side of this equation is a third-order polynomial in s . Thus, we can write T_4 as

$$T_4(s) = \frac{Ea(s - b)}{(s - s_1)(s - s_2)(s - s_3)}, \quad (3.8)$$

where a , b , s_1 , s_2 , and s_3 are constants determined by the various thermal resistances and heat capacities that we do not reproduce here. Importantly, s_1 , s_2 , and s_3 are all negative, because the (expanded) polynomial on the left-hand side of [Equation 3.7](#) has only positive terms as all of the physical quantities (thermal resistance and conductivity) are positive.

Transforming back into time space, we have

$$T_4(t) = Ea \left[\frac{s_1 - b}{(s_1 - s_2)(s_1 - s_3)} e^{s_1 t} + \frac{s_2 - b}{(s_2 - s_1)(s_2 - s_3)} e^{s_2 t} - \frac{s_3 - b}{(s_3 - s_1)(s_3 - s_2)} e^{s_3 t} \right]. \quad (3.9)$$

The coefficients of the second and third terms sum to the negative of the coefficient of the first term. Thus, we can define a new constant α and a new scaling factor γ to rewrite this as

$$T_4(t) = \gamma E \left[-e^{-t/\tau_1} + \alpha e^{-t/\tau_2} + (1 - \alpha)e^{-t/\tau_3} \right], \quad (3.10)$$

where we have used $s_i = -1/\tau_i$ to put this equation in a more familiar form: a pulse with one rise time constant and two decay time constants.

Despite being an approximation — we assumed above that the thermal conductivities and resistances do not change with temperature — this function describes our pulses quite well. CUORE pulses do appear to have one rise time constant and two fall time constants, and the temperature does increase approximately linearly with the energy deposited. We cannot, however, use this function to determine our expected pulse shape, as the various physical parameters that determine the time constants are themselves determined by fitting the pulse shape to this function. These parameters also vary between bolometers in the towers.

3.1.2 Crystal production

The performance of the bolometers depends critically on the quality of the TeO_2 crystals used as the heat absorbers. The crystals must be extremely radiopure, have very low impurity levels, have very tight dimensional tolerances, and have nearly perfect crystal structure. And because CUORE is composed of 988 individual bolometers, the process for producing crystals must be reliable, accurate, and highly reproducible.

The crystals were grown by the Shanghai Institute of Ceramics at the Chinese Academy of Sciences. The production process can be divided into two phases: crystal synthesis and crystal polishing. The primary concern of the first phase is the production of near-perfect crystals with extremely low bulk contamination, and the

focus of the second is to process the crystals into the correct shape while avoiding surface contamination.

The crystal production begins with pure metallic tellurium. The metallic tellurium is dissolved with aqua regia (a solution of HNO_3 and HCl) and precipitated with ammonia. This leaves behind TeO_2 powder, which is washed repeatedly and dried. This powder is then dissolved in hydrochloric acid, forming TeCl_4 , which is filtered and precipitated with NH_4OH . After washing the precipitate and evaporating away the water, the precipitate is heated to $680\text{ }^\circ\text{C}$ for 24 hours to produce a 99.99%-pure TeO_2 powder for crystal growth. The crystals are then grown with the vertical Bridgman technique [101] in platinum crucibles. In this process, the crucible is heated to approximately $830\text{ }^\circ\text{C}$, above the melting point of the TeO_2 powder, and slowly lowered (0.6 mm/hour) inside a furnace with a temperature gradient while the furnace temperature is slowly raised ($3\text{ }^\circ\text{C/hour}$). Thus, the solid-liquid interface position is steady as the crystal grows and is lowered [102].

After growing, the crystal is inspected and the best portions are chosen to be ground down into powder; the rest is not used. The powder is then subject to the same cleaning process as before: dissolution with TeCl_4 , precipitation with NH_4OH , filtering, washing, and heating. TeO_2 crystals are grown once again from this now-extremely-clean TeO_2 powder [103]. This double-growth process is the key to the cleanliness of the CUORE crystals, allowing us to reach bulk contamination levels of $< 2 \times 10^{-14}\text{ g/g}^5$ for ^{238}U and $< 3 \times 10^{-14}\text{ g/g}$ for ^{232}Th (90% C.L.) [104].

Following the double-growth procedure, the crystals are taken to be cut to size and polished. Each crystal is first subjected to rough mechanical processing, in which it is X-ray oriented, cut almost down to size, and cleaned. Then, inside a clean room, each face of the crystal is chemically etched and mechanically polished. Both the etching

5. The units g/g are used to signify that the value provided is a mass fraction as opposed to a volume fraction.

and polishing remove approximately 10^5 atomic layers from the surface of the crystals, removing the surface impurities introduced by the rough mechanical processing and producing flat, high-quality surfaces [103]. The cleaning leaves the crystals with surface contamination levels of $< 4 \times 10^{-9}$ Bq/cm² for ²³⁸U and $< 2 \times 10^{-9}$ Bq/cm² for ²³²Th (90% C.L.) [104].

Finally, the crystals are triple-vacuum-packed in polyethylene bags to avoid radon exposure, individually barcoded, and placed in dedicated crates to be sent to LNGS. The crystals are shipped by sea to minimize cosmogenic activation and placed underground immediately on arrival. Even during the approximately 1-month above-ground exposure after crystal growth and before arriving underground, cosmogenic neutrons and protons result in non-negligible production of ^{110m}Ag and ⁶⁰Co, two long-lived isotopes that decay with energy greater than the Q -value of $0\nu\beta\beta$ decay of ¹³⁰Te. Analyses of the cosmogenic neutron and proton activation estimate that these isotopes contribute under 1% of the background at the $0\nu\beta\beta$ decay Q -value [105, 106].

3.2 Before CUORE

CUORE is the latest in a series of progressively more sensitive $0\nu\beta\beta$ decay searches with TeO₂ bolometers (see Figure 3.4). The technology, mass, and sensitivity of these experiments have improved steadily over the past 25 years, following the proposal of the technique by Fiorini and Niinikoski in 1984 [107]. Starting with a single 6-g crystal [108], experiments grew to become 21-g, 34-g, 73-g and 334-g crystals [109–111], a 4-crystal array [112], and a 20-crystal tower [113]. The success of these early experiments lead to the idea of CUORE, a massive scale-up to a ton-scale detector [114]. But first, we have Cuoricino and CUORE-0, which paved the way for CUORE’s successful operation.

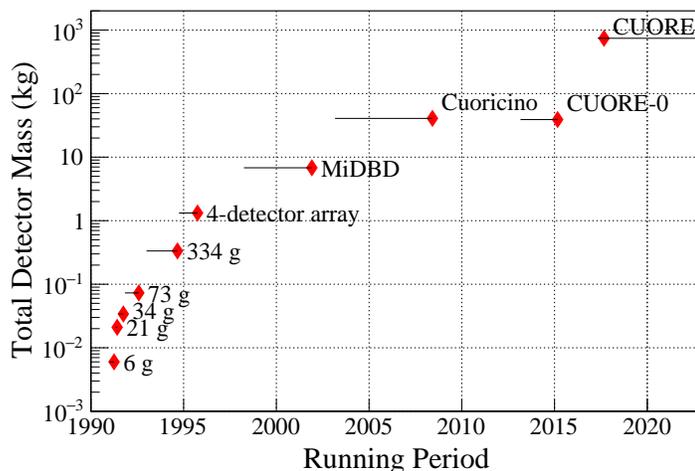


Figure 3.4: History of $0\nu\beta\beta$ decay experiments with TeO_2 bolometers. All of the experiments prior to the “4-detector array” are single-crystal experiments, identified by the crystal mass.

3.2.1 Cuoricino

Cuoricino was the first experiment to utilize a detector tower similar to that used in CUORE. The 62-crystal tower contained two different crystal sizes — $3 \times 3 \times 6 \text{ cm}^3$ and $5 \times 5 \times 5 \text{ cm}^3$ — and the smaller crystals were a mix of natural-abundance tellurium crystals and crystals enriched in ^{130}Te and ^{128}Te . The overall detector mass was 40.7 kg, with 11.3 kg of ^{130}Te , and it was operated in a cryostat underground at LNGS. The Cuoricino tower is shown in [Figure 3.5a](#).

Cuoricino took data from 2003 to 2008, accumulating 19.75 kg·y of ^{130}Te exposure. With this exposure, it set the best experimental limit to date on the $0\nu\beta\beta$ decay half-life of ^{130}Te : $T_{1/2}^{0\nu\beta\beta} > 2.8 \times 10^{24} \text{ y}$ (90% C.L.) [[115](#)]. The final Cuoricino spectrum in the $0\nu\beta\beta$ decay region of interest (ROI) is in [Figure 3.5b](#).

In the Cuoricino data, there are two background sources in the ROI: α particles and γ rays. The γ background in the ROI comes almost entirely from 2615 keV γ rays from ^{208}Tl decay, the result of ^{232}Th bulk contamination in the detector towers and in the cryostat. Many of the γ rays emitted in ^{208}Tl decay Compton scatter inside the

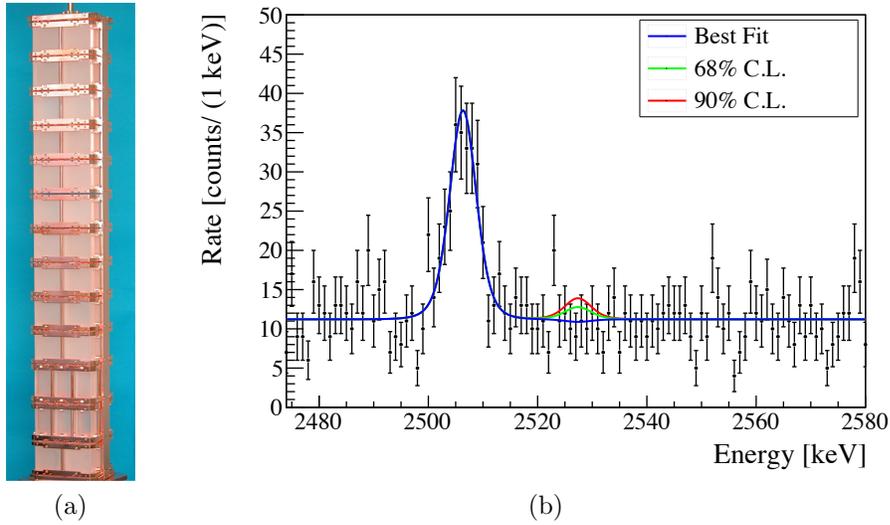


Figure 3.5: (a) The Cuoricino detector tower, showing $3 \times 3 \times 6 \text{ cm}^3$ and $5 \times 5 \times 5 \text{ cm}^3$ crystals. (b) The final Cuoricino spectrum in the $0\nu\beta\beta$ decay region of interest, summed over all channels, showing the best-fit line and the 68% and 90% C.L. upper limits. The peak around 2505 keV is from ^{60}Co decay. Figure from [115].

cryostat shielding or other passive material before reaching the detectors, and thus deposit less than their full energy in the crystals, contributing to the background in the ROI. The α background, on the other hand, which is dominant in Cuoricino, comes only from surface contamination, and is from a variety of different α lines (see Figure 3.6). Any α contamination not on the crystal surface would deposit its full energy in the crystal, and thus would not degrade and contribute to the background in the ROI. This bulk contamination of the crystals is seen in the 3.3 MeV line resulting from ^{190}Pt α decay, for example, which is from the crucibles in which the crystals are grown, and which does not affect the background in the ROI.

The results from Cuoricino also showed that the energy resolution of the larger crystals ($\text{FWHM} = 6.3 \pm 2.5 \text{ keV}$) was superior to that of the smaller crystals ($\text{FWHM} = 9.9 \pm 4.2 \text{ keV}$) and enriched crystals ($\text{FWHM} = 13.9 \pm 5.3 \text{ keV}$). Due to this, for CUORE-0 and CUORE, we use crystals of the larger size made from natural-abundance tellurium.

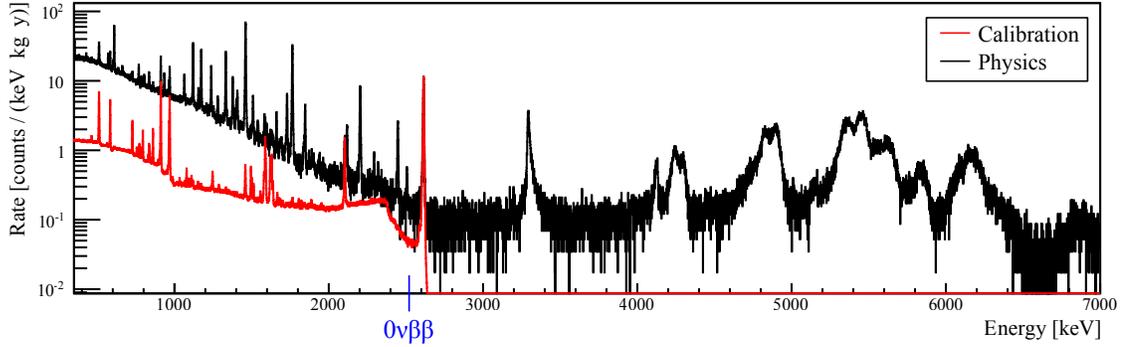


Figure 3.6: The final Cuoricino physics (background) energy spectrum, summed over all channels, with the calibration spectrum overlaid. The calibration spectrum is normalized to the height of the 2615 keV peak from ^{208}Tl . The Q -value of ^{130}Te $0\nu\beta\beta$ decay is indicated.

3.2.2 CUORE-0

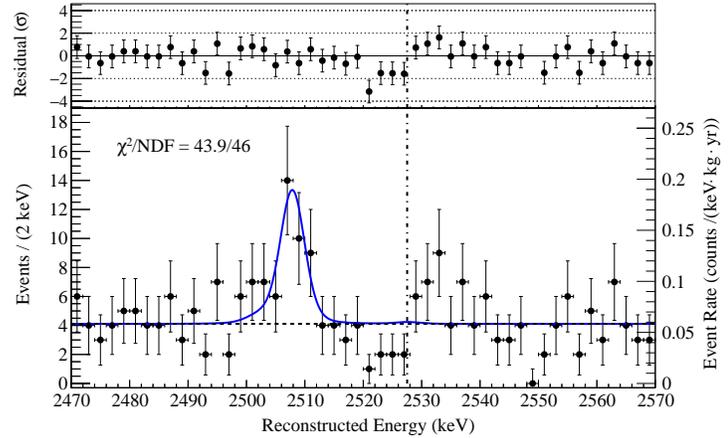
The design of the CUORE detector towers was finalized following the results of Cuoricino. CUORE would be built with a reduced amount of copper in the tower frame, and the detector parts would be subject to significantly improved cleaning and surface treatments. Because CUORE is composed of 19 towers, a semi-robotic tower assembly line was created in order to optimize the tower construction process (see [Section 3.3](#)). We operated the first tower created with the new design from the assembly line in the Cuoricino cryostat as the CUORE-0 experiment.

CUORE-0 served both as a validation of our improved cleaning and detector construction procedures and as a sensitive $0\nu\beta\beta$ decay search in its own right. It operated from 2013–2015, accumulating 9.8 kg·y of ^{130}Te exposure, approximately half the total exposure of Cuoricino. With CUORE-0, we set a limit on the $0\nu\beta\beta$ decay half-life of ^{130}Te of $T_{1/2}^{0\nu\beta\beta} > 2.7 \times 10^{24}$ y (90% C.L.), or $T_{1/2}^{0\nu\beta\beta} > 4.0 \times 10^{24}$ y (90% C.L.) when combined with the Cuoricino data [82]. The final CUORE-0 spectrum in the $0\nu\beta\beta$ decay ROI is in [Figure 3.7](#).

The stringent materials selection and cleaning, as well as a reduced amount of copper in the tower frame, helped to greatly reduce the backgrounds in CUORE-



(a)



(b)

Figure 3.7: (a) The fully assembled CUORE-0 tower inside a glovebox. (b) The final CUORE-0 spectrum in the $0\nu\beta\beta$ decay region of interest, summed over all channels, showing the best-fit line. The Q -value of ^{130}Te $0\nu\beta\beta$ decay is indicated by the vertical dashed line. The peak around 2505 keV is from ^{60}Co decay. Figure from [82].

0 with respect to Cuoricino (see Figure 3.8). The background in the α -dominated region above 2.7 MeV is particularly improved over Cuoricino, decreasing by a factor of 7. In the $0\nu\beta\beta$ decay ROI, the background was improved by a factor of 3 to 0.058 ± 0.004 (stat) ± 0.002 (syst) counts/(keV·kg·yr). The background improvement at and below the ROI, where the background is dominated by γ rays, is limited by the fact that CUORE-0 was operated in the Cuoricino cryostat, which is the main source of the γ background.

3.3 Creating the detector towers

The CUORE detectors towers are built from as few components as possible. The minimization of components helps in constructing the towers, as there are 988 individual detectors in 19 towers, but is primarily necessary because each component increases the overall level of radioactive background in the experiment. All detector components are subject to strict protocols to ensure their cleanliness and radiopurity, and

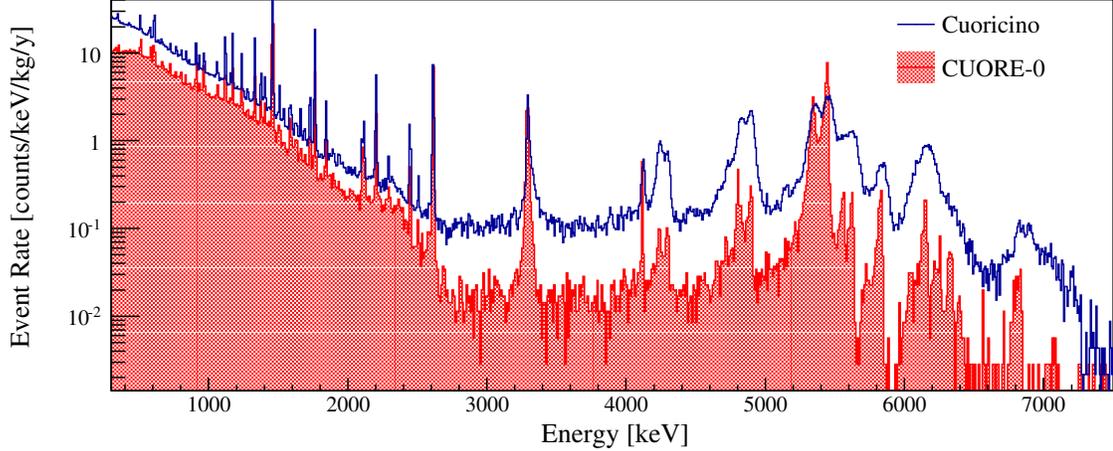


Figure 3.8: Backgrounds in the Cuoricino and CUORE-0 experiments. The backgrounds are reduced throughout moving from Cuoricino to CUORE-0, but are particularly reduced in the region above 2.7 MeV, which is dominated by α particles. Figure from [82].

samples of each are assayed to determine their final bulk and surface contamination levels.

3.3.1 Parts selection and cleaning

The components of a CUORE tower are the 52 TeO_2 crystals; a copper tower frame and wire trays; PTFE crystal holders; silicon heaters and NTD germanium thermistors, along with gold signal wire and glue that holds them to the crystals; and copper-on-polyethylene-naphthalate signal tapes (see Figure 3.9). The masses and bulk contamination levels of each of these components is listed in Table 3.1.

Copper constitutes the largest portion of the detector tower mass, besides from the crystals, and the cleanliness of the copper tower support frame is crucial to CUORE's low background. The 3.7 kg of copper parts in each CUORE tower is made from an Aurubis high-purity Electrolytic Tough Pitch (ETP1) copper alloy, NOSV. This copper has extremely low bulk contamination levels of $< 5 \times 10^{-12}$ g/g for ^{238}U and $< 5 \times 10^{-13}$ g/g for ^{232}Th (90% C.L.) [117]; however, the casting and machining

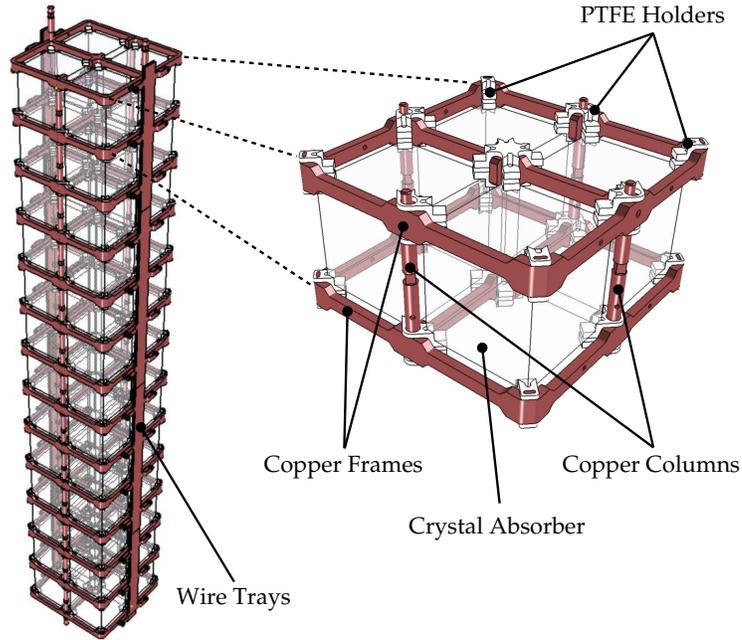


Figure 3.9: Components of a CUORE detector tower. Figure from [116].

Component	Mass [g]	^{238}U contamination limit		^{232}Th contamination limit	
		[g/g]	[bolo $^{-1}$ yr $^{-1}$]	[g/g]	[bolo $^{-1}$ yr $^{-1}$]
TeO $_2$ crystals	742 000	5×10^{-14}	10	2×10^{-13}	2000
Cu structure	70 000	5×10^{-12}	100	5×10^{-13}	500
PTFE holders	5500	2×10^{-12}	4	2×10^{-12}	100
Cu-PEN tape	3000	1×10^{-10}	100	4×10^{-10}	20 000
Ge thermistors	42	1×10^{-9}	20	1×10^{-9}	500
Si heaters	6.8	2×10^{-10}	0.5	8×10^{-11}	7
Au bonding wires	1.5	1×10^{-9}	0.6	1×10^{-8}	200
Glue	0.44	8×10^{-10}	0.1	2×10^{-10}	1

Table 3.1: Masses and upper limits on the bulk contamination levels of the CUORE detector tower components. All contamination data are limits at 90% C.L., expressed in grams of the contaminant per gram of the component and in decays per year per bolometer. Note that this is not a measure of the number of events actually recorded in the bolometers; it is just a metric to compare the magnitude of the possible contaminations, taking into account the different masses of the tower components and two different isotopes.

of the copper components introduces significant surface contamination that must be mitigated.

To minimize surface contamination levels, all copper parts are subject to an aggressive cleaning procedure, consisting of precleaning, mechanical abrasion (tumbling), electropolishing, chemical etching, and plasma etching. The precleaning uses tetrachloroethylene, acetone, ethanol, and alkaline soap to remove large contaminants, such as grease and oil, introduced by machining. Following precleaning is tumbling, where abrasive alumina powder is used in a wet environment to erode away approximately 1 μm of the copper surface. Next is electropolishing, where the copper is oxidized and the oxide is dissolved in phosphoric acid and butanol, removing approximately 100 μm of material and resulting in a smooth, mirror-like surface. Finally the copper undergoes chemical etching, where the surface is chemically eroded and passivated with sulfamic acid, and plasma etching, which further erodes the surface with a plasma in vacuum. The combined cleaning procedure leads to surface contamination levels of $< 7 \times 10^{-8}$ Bq/cm² for both ²³⁸U and ²³²Th (90% C.L.) [117].

Holding the crystals in place are a variety of PTFE blocks, together weighing 290 g per tower. We designed the PTFE holders for strong mechanical and thermal coupling to the crystals at cryogenic temperatures. They are machined from virgin-grade PTFE with extremely low bulk contamination levels of $< 2 \times 10^{-12}$ g/g for both ²³⁸U and ²³²Th (90% C.L.) [116]. To minimize surface contamination levels, the holders are washed with soap and ultra-pure nitric acid before they are used in the tower assembly.

Glued to the crystals are thermistors made from neutron-irradiated germanium wafers. After irradiation, the wafers are polished, etched with nitric and hydrofluoric acid, and doped with boron, and gold pads are added to the sides and top. The small gold pads on the top allow for wire bonding on the front of the thermistors. The thermistors are $3.0 \times 2.9 \times 0.9$ mm³ (see Figure 3.10) and have bulk contamination

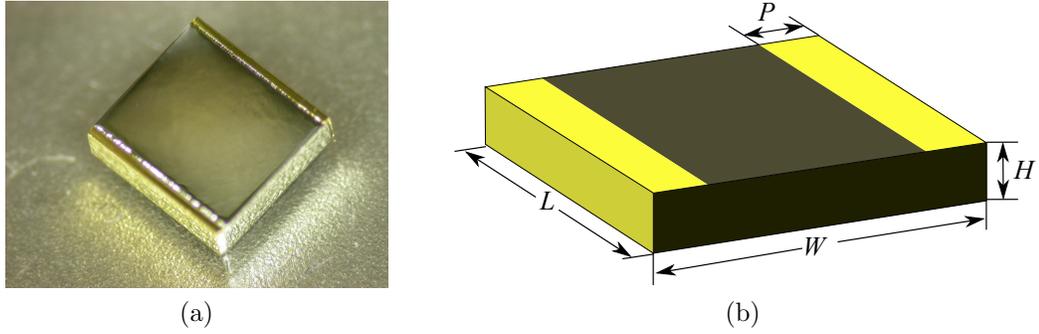


Figure 3.10: (a) A photograph of a CUORE-style thermistor, with gold pads on two sides and the top. (b) A sketch of the thermistor geometry. Nominal dimensions are $L = 3.0$ mm, $W = 2.9$ mm, $H = 0.9$ mm and $P = 0.2$ mm. Figure from [116].

levels of $< 1 \times 10^{-9}$ g/g for both ^{238}U and ^{232}Th (90% C.L.) [116].

Also glued to the crystals are heaters produced from silicon wafers. We add a doped meander to the $2.33 \times 2.40 \times 0.52$ mm³ silicon chips with the planar process, and we add aluminum pads for bonding. The meander has a resistance of 300 Ω . The heaters have bulk contamination levels of $< 2 \times 10^{-10}$ g/g for ^{238}U and $< 8 \times 10^{-11}$ g/g for ^{232}Th (90% C.L.) [116]. They are bonded with gold wire in the same way as the thermistors.

The gold wire used to bond the thermistors and heaters has a diameter of 25 μm and contamination levels of $< 1 \times 10^{-9}$ g/g for ^{238}U and $< 1 \times 10^{-8}$ g/g for ^{232}Th (90% C.L.). They are glued to the crystals with Araldite⁶ Rapid, a strong, fast-setting, 2-component epoxy with good thermal conductivity. The Araldite Rapid also has low bulk contamination levels of $< 8 \times 10^{-10}$ g/g for ^{238}U and $< 2 \times 10^{-10}$ g/g for ^{232}Th (90% C.L.) [116].

The last component of the CUORE towers is the signal paths. Both the thermistor and heater signals travel up the tower on copper-insulator tapes with a polyethylene naphthalate substrate (“Cu-PEN tapes”). Copper traces, 17 μm thick, are etched on the tape (see Figure 3.11). Cu-PEN tapes are located on both sides of each tower,

6. Araldite is a registered trademark of Huntsman Advanced Materials.

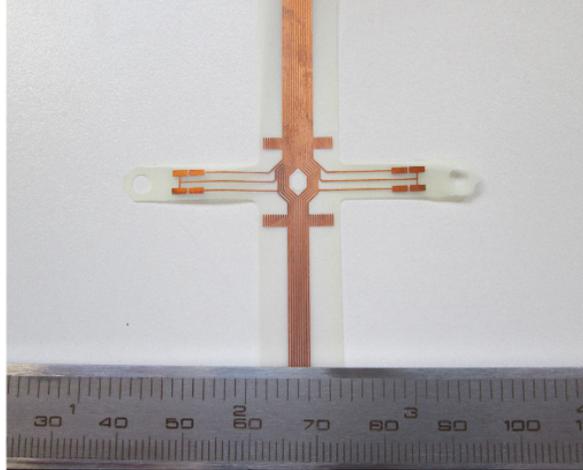


Figure 3.11: Copper pads and traces on Cu-PEN tapes. Photograph courtesy of the CUORE Collaboration.

each carrying the signals for half of the tower’s bolometers. The tapes have bulk contamination levels of $< 1 \times 10^{-10}$ g/g for ^{238}U and $< 1 \times 10^{-10}$ g/g for ^{232}Th (90% C.L.) [116].

3.3.2 Tower construction assembly line

With all of the components procured, processed, and cleaned, we can build the CUORE towers. The 19 towers and 988 crystals are constructed and instrumented on an assembly line, following strict procedures and protocols. We do not expose the materials in the tower to air at any point after they are cleaned⁷, with all work taking place in nitrogen-flushed gloveboxes, and we minimize contact between the tower components and other materials as much as possible. To avoid cosmogenic activation, we keep all parts underground and perform the tower assembly underground as well.

The first step in creating the CUORE towers is gluing a thermistor and heater to each crystal. One of the lessons learned from Cuoricino is that the coupling of the thermistors to the crystals greatly affects the bolometer signals. To this end,

⁷. As we will discuss, the completed towers are exposed to radon-free air for a short period while they are installed in the cryostat.

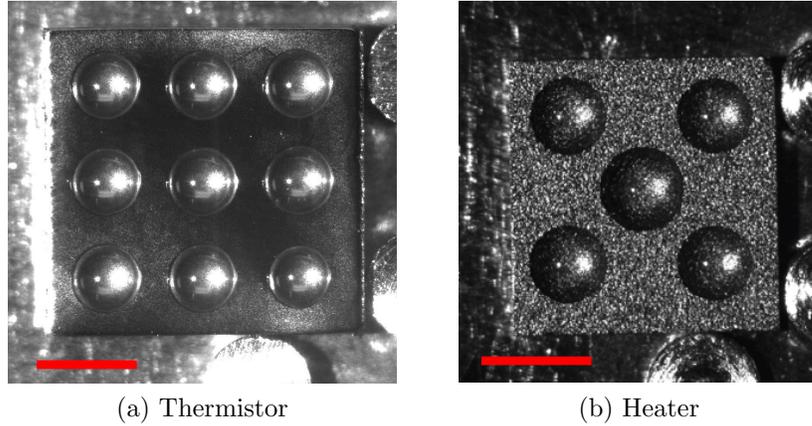


Figure 3.12: Glue spots as deposited by a robotic arm. Red scale bar represents 1 mm. Photographs courtesy of the CUORE Collaboration.

we use a semi-robotic system to make the gluing process precise and reliable. First, we place a thermistor and heater on a positioning platform. A dual-cartridge epoxy dispenser forces the epoxy through a static mixer, the mixed epoxy is injected into a syringe, and a robotic arm deposits a series of glue dots on the back side of the thermistor and heater. This glue dot matrix, consisting of 9 dots for the thermistors and 5 dots for the heaters (see [Figure 3.12](#)), provides good thermal and mechanical coupling while allowing for the relative thermal contraction of the TeO_2 crystals and the silicon and germanium chips without stress fractures developing. Once the glue is in place, another robotic arm picks up a TeO_2 crystal and places it on a cradle that is lowered onto the chips. Vacuum suction keeps the crystal 50 μm above the chips while the epoxy dries. All of this occurs inside a single glovebox, shown in [Figure 3.13](#).

All work following the gluing takes place in glove boxes mounted above a single working plane. In the center of this plane is a moving platform that can lower, raise, and rotate a tower. Below the working plane is a “garage” that can be closed off from the glove box; the moving platform can lower a tower into and out of this garage. The glove boxes are continually flushed with nitrogen, and when the glove boxes are changed and the working area is exposed to air, the tower is isolated in the

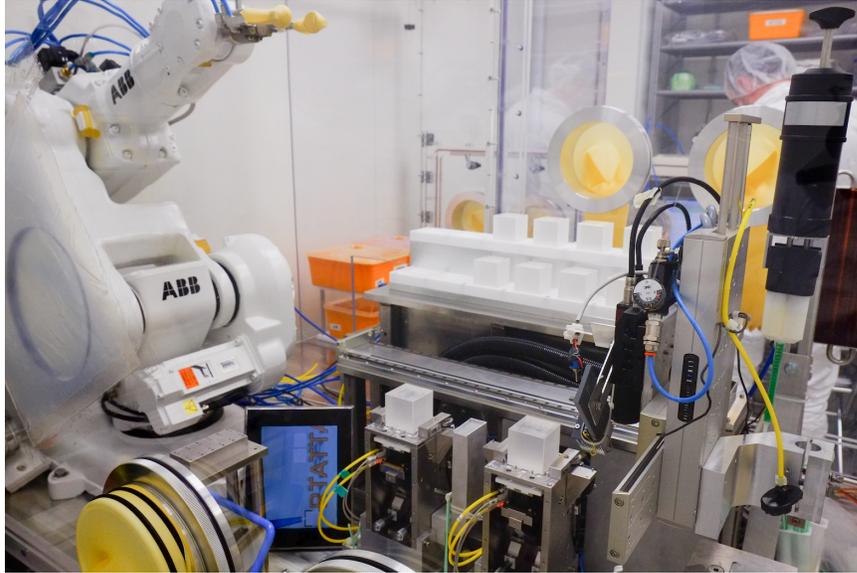


Figure 3.13: The glovebox used for gluing thermistors and heaters to the crystals, with the robotic arm on the left, two crystal gluing platforms in the center foreground, and two rows of glued crystals in the background. Photograph courtesy of the CUORE Collaboration.

nitrogen-flushed garage.

Once we have a heater and thermistor attached to each crystal that will belong to a single tower, we begin the mechanical assembly of the tower. We build each tower floor by floor from the copper frames and columns, PTFE holders, and TeO_2 crystals, and we lower it into the garage as each floor is completed. This assembly takes place in a glovebox with two sets of gloves on either side of the tower and is done by hand (see [Figure 3.14](#)). Following this, we place a taller glovebox on the working plane and attach the Cu-PEN tapes to the sides of the towers with Araldite Standard 2-component epoxy.

Next, the heater and thermistor on each crystal are electrically connected to copper pads on the Cu-PEN tape. We use a manual wire bonder for this task. For each connection, two gold wires are ball-bonded to the pad on the heater or thermistor and wedge-bonded to corresponding copper pad on the Cu-PEN tape, and then the wedge bonds are reinforced with security ball bonds (see [Figure 3.15](#)). When the

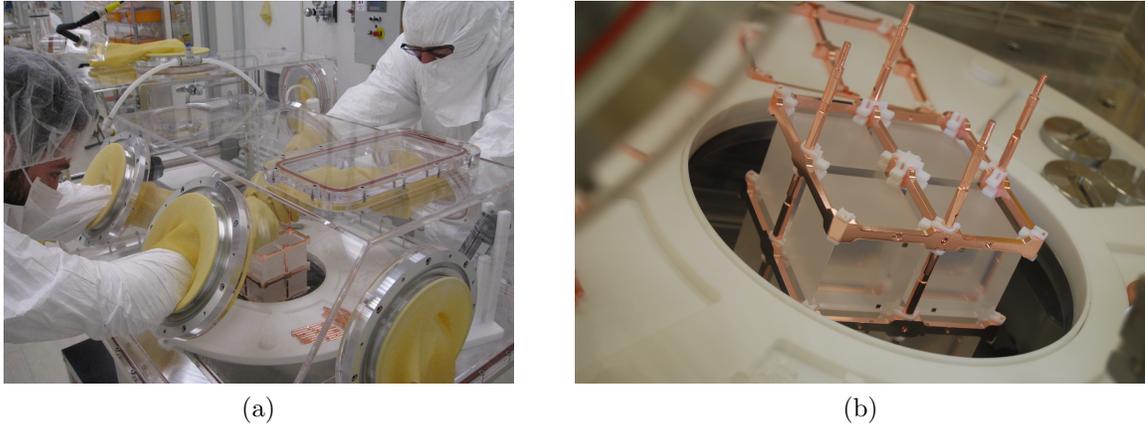


Figure 3.14: (a) The tower assembly glovebox, with a CUORE detector tower being assembled. (b) A partially assembled CUORE tower. Photographs courtesy of the CUORE Collaboration.

bonding is finished, a copper cover is added to the outside of the Cu-PEN tapes and the tower is complete.

3.4 Cryostat

In order to reach the domain of very low heat capacity for the TeO_2 crystals and a steep resistance-vs-temperature curve for the thermistors, the CUORE detectors must be operated stably near ~ 10 mK. The cryostat must not only cool the ton of detector towers and several tons of shielding, but also be strongly vibrationally damped and have extremely low radioactivity levels, especially in the colder stages near the detectors. To meet these goals, we use one of the most powerful dilution refrigerators ever built. The dilution refrigerator is assisted by pulse tube cryocoolers, allowing for cryogen-free, stable operation over the expected 5-year lifetime of the experiment.

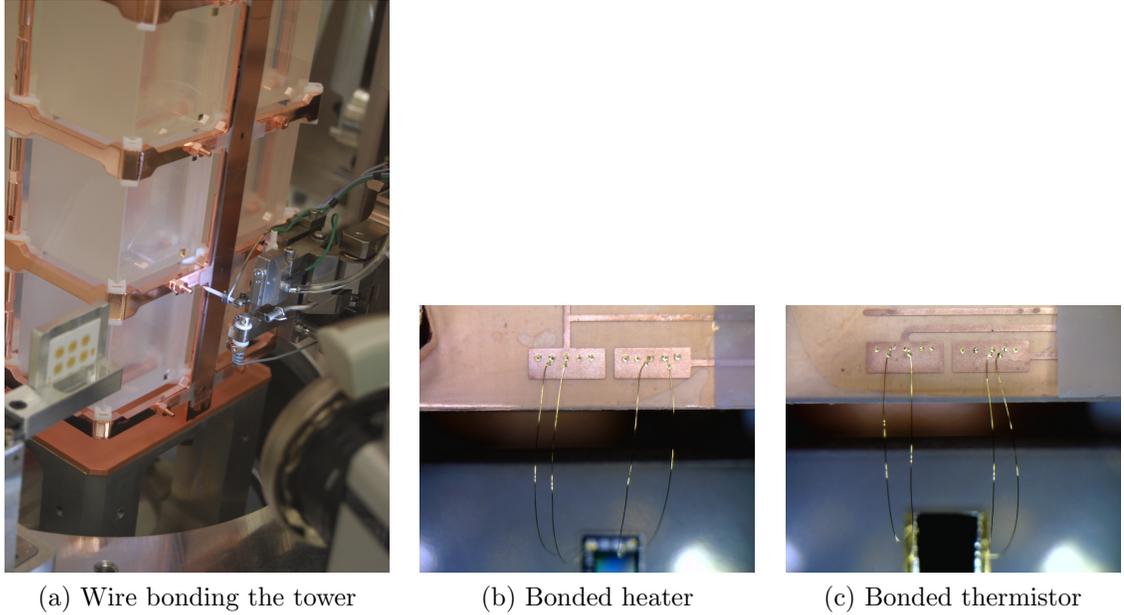


Figure 3.15: Wire bonds between pads on the Cu-PEN tape and a CUORE heater and thermistor. All wire bonds are doubled for redundancy. Photographs courtesy of the CUORE Collaboration.

3.4.1 Composition and shielding

Like all dilution refrigerator systems, the CUORE cryostat is composed of successively colder plates and corresponding heat shields (see [Figure 3.16](#)). The shields are nominally kept at 300 K, 40 K, 4 K, 600 mK (the still), 50 mK (the heat exchanger), and 10 mK (the mixing chamber), with the detector located inside the 10-mK shield. All the shields are copper, and the 300-K and 4-K shields are vacuum tight.

As is customary, the cryostat is supported from above (see [Figure 3.17](#)). The 300-K plate (i.e., the top of the cryostat) is held by a frame attached to four columns. The columns are bolted to a concrete foundation, which is separated from the bedrock by four elastomer platforms. All of the other cryostat plates and thermal shields are supported from the 300-K plate or from each other, with one exception: a thick copper plate dedicated to holding the detector towers is suspended independently from a Y-shaped beam above the cryostat. This beam is placed on three stiff springs

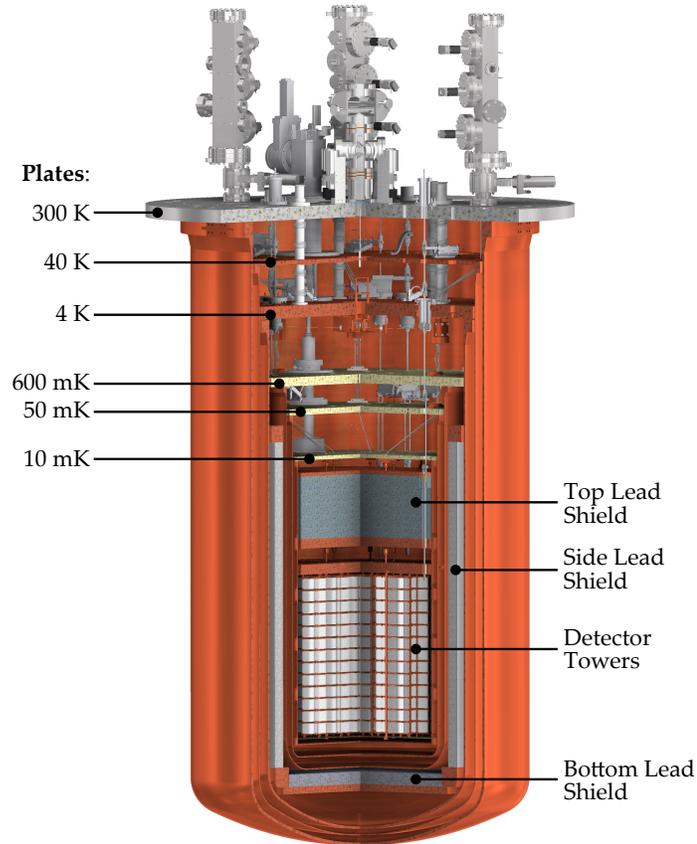


Figure 3.16: A rendering of the CUORE cryostat, showing all copper plates and vessels, lead shielding, and the detector towers.

for additional vibration isolation. This plate suspended from this beam is thermally coupled to the 10-mK shield and mixing chamber, and the detectors are attached to the bottom of the plate (see [Figure 3.18](#)). For a short length above the mixing chamber, Kevlar rope links the stainless steel and copper rods supporting the plate, which minimizes the heat load on the mixing chamber due to the suspension system.

Several layers of cold shielding surround the detector. Above the detector is a 30-cm layer of lead, kept at 50 mK, and around the sides and bottom are 6-cm lead shields at 4 K. The cold lead is supplemented by the copper cryostat vessels, which provide some additional shielding.

The side and bottom lead shields are particularly notable as they are composed of low-radioactivity lead recovered from an ancient Roman shipwreck. All lead contains

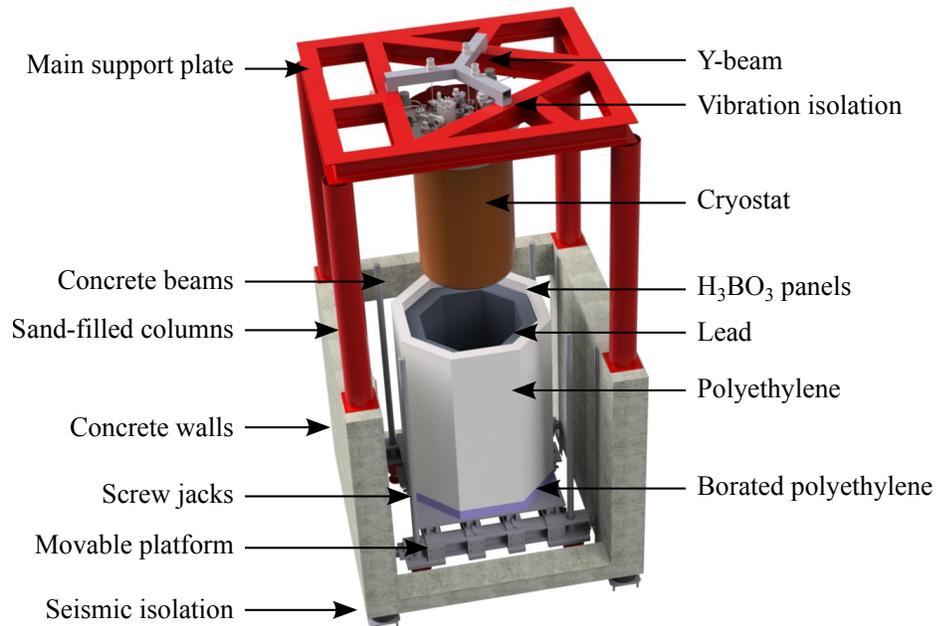


Figure 3.17: A rendering of the CUORE cryostat cryostat support structure and external shielding. The external lead and polyethylene shields are raised around the cryostat when taking data. Figure courtesy of the CUORE Collaboration.

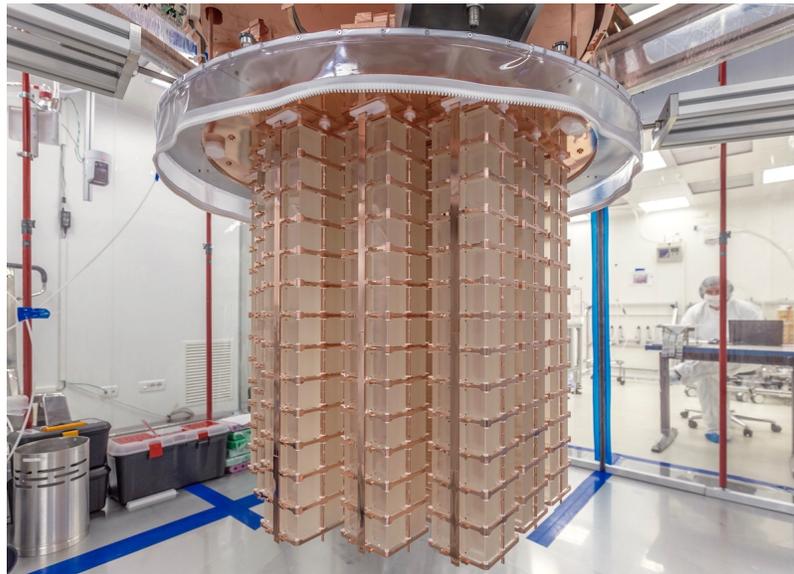


Figure 3.18: A photograph of the detector towers installed in the cryostat. Photograph courtesy of the CUORE Collaboration.

radioactive ^{210}Pb when mined, as it is produced by the ^{238}U decay chain; however, after it is mined, the ^{210}Pb is no longer replenished and begins to decay away with a half-life of 22 years. This makes ancient lead a very attractive source of shielding for the CUORE cryostat, as radioactive shielding is self-defeating. The lead used in CUORE has ^{210}Pb levels below 10 mBq/kg, the lowest levels ever measured in any type of lead [118].

Outside the cryostat, a 18-cm polyethylene shield thermalizes environmental neutrons, and a 2-cm layer of boric acid (H_3BO_3) absorbs them. Inside this is a 25-cm layer of lead bricks for environmental γ ray shielding. These external shields are octagonal and are raised around the cryostat when taking data.

3.4.2 Cooling power

The first line of cooling in the cryostat is provided by 5 two-stage pulse tube cryocoolers, each with a cooling power of 40 W at 45 K and 1.5 W and 4.2 K. These pulse tubes have remote motors to reduce vibrations on the cryostat and are thermalized to the cryostat plates with flexible mechanical couplings. As long as at least 4 of the 5 pulse tubes are functioning, the cryostat can maintain a steady base temperature.

The coldest temperatures in the cryostat are realized by a custom dilution refrigerator made by Leiden Cryogenics. This dilution refrigerator has 3000 μW of cooling power at 120 mK and 5 μW at 10 mK, and it can reach a base temperature below 6 mK circulating 8 mmol/s of ^3He - ^4He mixture.

Each time the cryostat is temperature cycled, it must cool approximately 15 tons of material to 4 K. With the 5 powerful pulse tube coolers alone, this would take approximately half a year. To greatly shorten this time, we inject cold helium gas into the cryostat to cool it down. The helium is circulated through the main cryostat and into a second, small cryostat, where it is cooled again and circulated back into the main cryostat. The second cryostat maintains a temperature gradient of under 40 K

between the outgoing and incoming gas to avoid thermal stress on the main cryostat. With this process, the cooling time to 4 K is lowered to approximately 2 weeks [119].

3.5 Sensitivity

When designing an experiment to search for $0\nu\beta\beta$ decay, and later, when interpreting the results, it is prudent to calculate the sensitivity of the experiment to this decay. Sensitivity calculations can help in the design and commissioning process, informing decisions on isotope choice and abundance, total detector mass and operating time, background goals, and detector resolution requirements. And it is helpful for comparing results to expectations, ensuring that the two are consistent.

If $0\nu\beta\beta$ decay is not discovered, the goal of a $0\nu\beta\beta$ decay experiment is to set an upper limit on the decay rate (or, equivalently, a lower bound on the half-life). There are various approaches to constructing this upper limit, which can be broadly classified as either Bayesian or frequentist. This limit is defined in terms of a confidence interval (frequentist) or credible region (Bayesian); for $0\nu\beta\beta$ decay searches, we conventionally express results at the 90% confidence level. The frequentist interpretation of the decay-rate limit is that the experiment has a 90% chance of measuring a decay rate less than the limit; it is based on no assumptions about the true decay rate. The Bayesian interpretation is that there is a 90% chance that the true decay rate is less than the upper limit; this relies on an assumption about the true value of the decay rate, which is often taken to be equally likely for every physical (greater-than-zero) value. Both interpretations are regularly used, and they generally produce similar results.

The sensitivity of the experiment is, in broad terms, the expected upper limit that the experiment can set on the $0\nu\beta\beta$ decay rate, or the lower limit it can set on the half-life. In practical terms, there are many definitions of the sensitivity. We pursue

two approaches to calculating the CUORE sensitivity, one theoretical and one based on Monte Carlo simulations.

3.5.1 Theoretical approach

In the theoretical approach, we define the experimental sensitivity as the half-life that results in the smallest mean signal that is greater than the background fluctuations at a 90% confidence level, following the approach in [120]. The mean signal S is the product of the number of ^{130}Te atoms, the probability that a given atom decays during the experiment, and the probability of detecting that decay. We write this as

$$S = \epsilon \left(\frac{MN_A a \eta}{W} \right) \left(\frac{t \ln 2}{T_{1/2}^{0\nu}} \right), \quad (3.11)$$

where ϵ is the $0\nu\beta\beta$ decay detection efficiency, M is the mass of the TeO_2 crystals, N_A is Avogadro's number, a is the isotopic abundance of ^{130}Te , η is the stoichiometric coefficient (1 for this decay), W is the molecular mass of TeO_2 , t is the length of data-taking, and $T_{1/2}^{0\nu}$ is the (unknown) half-life of the decay. This expression assumes that $t \ll T_{1/2}^{0\nu}$ in order to make the approximation $1 - (1/2)^{t/T_{1/2}} \approx (\ln 2)(t/T_{1/2})$; experimental limits put the half-life at the level of 10^{25} years, so this assumption should not be a problem for CUORE. For TeO_2 with a natural isotopic abundance of ^{130}Te , this expression evaluates to

$$S = \epsilon (Mt \text{ [kg}\cdot\text{yr]}) \left(\frac{8.9 \times 10^{23}}{T_{1/2}^{0\nu} \text{ [yr]}} \right), \quad (3.12)$$

where the product Mt is referred to as “exposure” and is expressed in units of $\text{kg}\cdot\text{yr}$.

The mean signal must be compared to the background fluctuations in order to determine a confidence level for the observation. The mean background B in an

energy region ΔE is defined as

$$B = bMt \cdot \Delta E, \quad (3.13)$$

where Mt is the exposure and b is the background index in units of counts/(keV·kg·yr). This formulation makes it straightforward to calculate the total number of expected background events for a detector of given mass with a given live time, and the background index is widely used in the field as a measure of overall background levels. The true background observed is Poisson-distributed with a mean of B .

We take a frequentist view and look for the number of signal events S such that, with the expected number of background events B in our energy region ΔE , there is a probability α of observing at least that many signal events by background fluctuations alone. That is, we solve

$$\sum_{k=S+B}^{\infty} p_B(k) = \alpha, \quad (3.14)$$

where $p_B(k)$ is the Poisson probability mass function with mean B evaluated at k . This is a discrete sum, but we can use a continuous form to evaluate the sensitivity by using the cumulative distribution function of the Poisson distribution, which gives

$$\sum_{k=S+B}^{\infty} p_B(k) = 1 - \frac{\Gamma(S+B, B)}{\Gamma(S+B)} = \alpha, \quad (3.15)$$

where $\Gamma(x, y)$ is the incomplete gamma function. We numerically solve this equation for S , given B (evaluated with [Equation 3.13](#)) and α (0.10 for a 90% C.L. limit), and convert it to a half-life sensitivity using [Equation 3.12](#). The CUORE sensitivity calculated with this method is shown in [Figure 3.19](#) for different background levels and energy resolutions. Note that in this formulation, the effective efficiency ϵ for detecting the decay also depends on the energy window used. A smaller energy

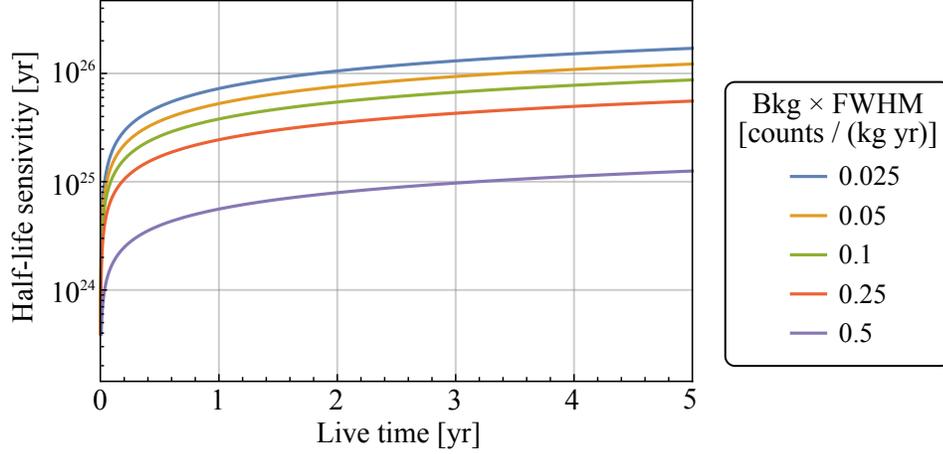


Figure 3.19: The CUORE sensitivity over time, calculated theoretically with frequentist methods. The different curves represent different values of the product of background index (in counts/(keV·kg·yr)) and the FWHM detector energy resolution (in keV). The CUORE goal, a background index of 0.01 counts/(keV·kg·yr) with a 5 keV energy resolution, is represented by the “0.05” line in orange.

window means a smaller number of background counts, but since the $0\nu\beta\beta$ decay signal is Gaussian distribution around the Q -value, then smaller energy windows will reduce the percentage of actual $0\nu\beta\beta$ decays in the examined energy region. The optimal energy window depends on the background level, energy resolution, and exposure, but is approximately 1.2 times the full width half max (FWHM) of the expected signal.

In the case of larger backgrounds, the Poisson distribution can be approximated as a Gaussian distribution. In this case, the Gaussian has mean B and standard deviation \sqrt{B} . Then, as in the Poisson case, we look for the number of signal events S such that there is a probability α of observing at least that many signal events by background fluctuations alone. For a Gaussian distribution, this means that the signal must be a certain number of background standard deviations above zero, or $S = n\sqrt{B}$. Using this in equation [Equation 3.12](#), and assuming that the energy region ΔE is proportional to the resolution of the detector, we recover the sensitivity figure of merit in [Equation 2.29](#).

3.5.2 Monte Carlo

A complementary method of determining the sensitivity is through a Bayesian analysis of Monte Carlo simulations. In this approach, we feed a Monte Carlo simulation with parameters like the exposure, detector resolution, background index, and rate of any nearby background lines, and we generate a large number of spectra assuming no $0\nu\beta\beta$ decay signal. We fit each generated spectrum with a binned maximum likelihood fit and extract an upper limit on the decay rate from each, assuming a flat prior in the physical region, where the decay rate is nonnegative. The median of these decay rates is then the sensitivity of the experiment. This binned maximum likelihood fit similar to the unbinned fit used in the analysis of the experimental data and will be described fully in [Section 6.3](#).

The Monte Carlo approach for calculating the sensitivity is useful for comparing our observed limit to the limit that we expect, since it is determined similarly to how we calculate our experimental limit. It can also take into account channel-dependent resolutions and nearby spectral features. It generally yields results that are compatible with the sensitivity calculated with the frequentist theoretical approach above but is significantly more computationally expensive.

A fully Bayesian sensitivity study was performed just before CUORE began acquiring data. Using the CUORE background goal of 0.01 counts/(keV·kg·yr) and resolution goal of 5 keV (FWHM), the expected sensitivity of CUORE to $0\nu\beta\beta$ decay at 90% C.L. is 2×10^{25} yr with 3 months and 9×10^{25} yr with 5 years of live time [87]. We will return to a discussion of the CUORE sensitivity using the first physics results in [Section 6.3.2](#).

Chapter 4

Detector Calibration System¹

As a particle passes through a CUORE crystal, the energy it deposits is converted into phonons, causing a temperature rise that is measured by a neutron-transmutation-doped germanium thermistor [97]. The relationship between the thermistor voltage reading and the original particle energy is nonlinear and unique to each bolometer–thermistor pair [122]. Because the signature of $0\nu\beta\beta$ decay is a peak in the energy spectrum at the Q -value of the decay, a precise understanding of the bolometer energy scale is critical for detecting this process. In addition, detecting other rare processes, such as two-neutrino double-beta decay, requires an understanding of the spectrum over a wide range of energies. As a result, absolute energy calibration of each bolometer using sources at a variety of energies is required. The response of the bolometers and thermistors is highly temperature-dependent, and as such, the calibration must be performed with the bolometers at their base temperature.

During the projected 5-year operating period of CUORE, we will calibrate the bolometer–thermistor pairs regularly, as detector conditions can change over time. In CUORE-0, a predecessor experiment, calibration was performed monthly [82], and

1. This chapter incorporates content from J. S. Cushman *et al.* [121]. Used with permission from Elsevier. The work described in this chapter was performed by all of the coauthors on this paper: A. Dally, C. J. Davis, L. Ejzak, D. Lenz, K. E. Lim, K. M. Heeger, R. H. Maruyama, A. Nucciotti, S. Sangiorgio, T. Wise, and me. The text is my own.

the calibration frequency to date is similar for CUORE. Each calibration period is kept as short as possible to maximize the live time for physics data taking. At the same time, because of the ~ 4 -second response and recovery time of the bolometers [123], the source activity must be sufficiently low to avoid pile-up. Because of the compact configuration of the CUORE detector towers, the outer bolometer towers in the cryostat partially shield the innermost towers from external radiation. As a result, to achieve calibration periods of one or two days while not saturating the outer detectors, calibration sources must be placed near each bolometer throughout the tower array during calibration. Because these sources are removed from the cold detector region of the cryostat during physics data taking, they must be cooled to the cryostat’s base temperature for each calibration and subsequently warmed up again.

Deploying calibration sources into a cold cryostat without significant disruption to the cryostat operating temperature poses demanding technical challenges. To accomplish this task, we have designed and implemented the CUORE Detector Calibration System (DCS). In [Section 4.1](#), we present an overview of the DCS, including the design and production of the calibration sources, the motion control and monitoring hardware, and the tubes and other hardware that guide the calibration sources through the cryostat. In [Section 4.2](#), we discuss the electronic control system for the DCS and the remote and automatic DCS software controls. Finally, we present and discuss the results of a calibration source deployment down to base temperature in the CUORE cryostat in [Section 4.3](#).

4.1 System overview

The CUORE cryostat contains a large custom-built cryogen-free dilution refrigerator assisted by pulse tube cryocoolers [124]. It comprises six plates and corresponding copper vessels held at successively colder temperatures (see [Figure 3.16](#)). A stainless

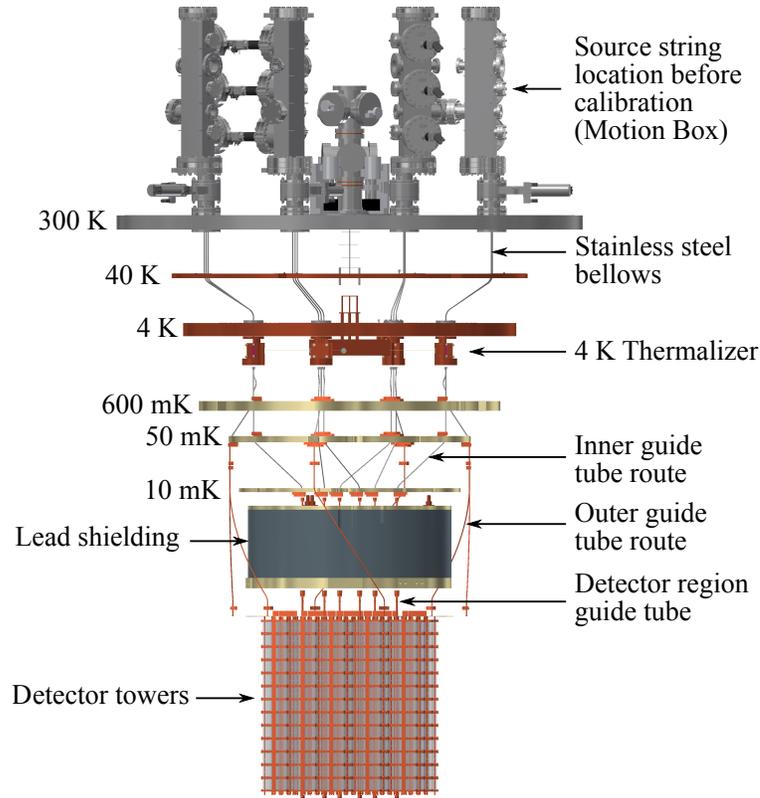


Figure 4.1: Illustration of the DCS in the CUORE cryostat.

steel room-temperature (300 K) plate provides all connections to the outside of the cryostat and supports the calibration system and other hardware. Copper plates at 40 K and 4 K are cooled by pulse tube cryocoolers. A gold-plated copper plate at 600 mK is thermally coupled to the still of the dilution unit, and there are similar plates at 50 mK and at 10 mK, coupled to the heat exchanger and mixing chamber, respectively. The detector towers are located below the mixing chamber, underneath the top lead shielding. The DCS is the motion and thermalization hardware that guides calibration sources into the cryostat and extracts them after each calibration period has concluded.

Each calibration source carrier is a collection of individual source capsules attached to a continuous string. There are 12 calibration source carriers (“source strings”) in

total. During physics data taking, these source strings are wound on spools above the cryostat at room temperature, outside of the internal lead and copper cryostat shielding, to avoid gamma rays from the strings reaching the detectors. The spools are contained inside vacuum-tight stainless steel enclosures that are connected to the inner cryostat vacuum through gate valves. At the beginning of each calibration period, motorized and computer-controlled spools lower the source strings under their own weight into the cryostat.

As the strings are lowered, they are guided through the cryostat by a series of polytetrafluoroethylene (PTFE), stainless steel, and copper tubes (collectively referred to as “guide tubes”). As the strings pass the 4-K stage of the cryostat, a thermalization mechanism consisting of two parallel, spring-loaded copper blocks squeezes the sources to cool them down to 4 K. Below this stage, the guide tubes divide into inner and outer paths. The six inner paths bring strings to the innermost region of the cryostat where the detectors are mounted, whereas the six outer paths bring strings to just outside the 50-mK vessel. The inner strings pass through lead shielding above the detectors and reach their final positions inside copper tubes mounted between the detectors; the outer strings are allowed to hang freely because they are outside the detector region. An illustration of the DCS in the CUORE cryostat is shown in [Figure 4.1](#).

4.1.1 Requirements

The primary function of the DCS is to deploy calibration sources into the cryostat and cool them down without affecting the operating temperature of the detectors. The design and construction of the system is largely driven by the strict thermal and radioactivity requirements of the experiment. The DCS must respect the thermal load requirements of each stage of the dilution refrigerator, both when sources are stationary and when they are moving. The hardware that remains in the cryostat

Thermal stage	Available cooling power
40 K	1 W
4 K	300 mW
600 mK	600 μ W
50 mK	1 μ W
10 mK	1 μ W

Table 4.1: Cooling power budgeted for the calibration system at all thermal stages of the cryostat.

during physics data taking must also make a negligible contribution to the radioactive background in the $0\nu\beta\beta$ -decay energy region of interest around 2528 keV [94–96]. Finally, the system must operate safely and stably over the lifetime of the experiment and must be flexible, allowing us to change or replace any of the calibration sources as necessary.

The source strings and guide tubes must only minimally impact the thermal conductivity between the various thermal stages of the cryostat. In addition, there must be no straight-line access between the very different temperature regions of the cryostat to minimize thermal radiation from warmer to colder stages. During string motion, frictional heating, which can be particularly problematic at the coldest stages of the cryostat, must also be minimized. Moreover, during string lowering, the source strings must be cooled as they are lowered to avoid dissipating large amounts of heat in the colder parts of the cryostat. Specifically, the DCS is designed such that any heat load due to source deployment and extraction can be compensated for by the temperature stabilization system of the cryostat, thereby avoiding any effects on the cryostat base temperature. The design goal for the cooling power available to the DCS in the cryostat is shown in [Table 4.1](#).

Another set of requirements for the DCS arises from the low-background environment that is necessary for a rare-event search with CUORE. Most importantly, the detectors must be shielded from the radioactivity of the calibration sources during

physics data taking, which necessitates removing the sources from the detector region. In addition, the radioactivity of the materials used to construct the DCS must be low, especially for the source string guide tubes in the detector region. To achieve the CUORE background goal of 10^{-2} counts/(keV·kg·yr) in the $0\nu\beta\beta$ decay region of interest, this necessitates constructing all DCS hardware from only ultrapure copper, with ^{232}Th and ^{238}U bulk contamination levels at or below 10^{-12} g/g, in the detector region of the cryostat; in this region, we construct the DCS hardware from the same NOSV copper used in the frames of the CUORE detector towers (see [Section 3.3.1](#)).

The DCS must also be both fail-safe and very unlikely to fail. The system must be constructed such that the risk of active source material escaping into or remaining behind in the cryostat is essentially zero. It must also be designed such that the source capsules do not become stuck in the cryostat. If any source material were to become stuck or remain behind in the cryostat following a calibration period, it would be necessary to warm up the cryostat to room temperature and open it to extract this material, which is a lengthy process. Therefore, we use a variety of sensors to monitor the system to ensure that any abnormal behavior will be caught before any damage occurs to the source strings or the cryostat.

Finally, the calibration sources must be replaceable without warming up the cryostat to allow a variety of different source isotopes to be inserted if desired.

4.1.2 Calibration sources

The Q -value of ^{130}Te $0\nu\beta\beta$ decay is 2528 keV; thus, a variety of gamma decay lines are appropriate for calibrating all energies up to this Q -value. For CUORE, ^{232}Th sources are a natural choice because of the wide variety of lines provided by the daughter isotopes of ^{232}Th , including the strong ^{208}Tl line at 2615 keV, which enables precise energy calibration near the Q -value. Sources containing ^{232}Th were also used in Cuoricino and CUORE-0, predecessor experiments to CUORE, thus providing us

with a vetted analysis framework for performing this calibration [125, 126]. The long half-life of ^{232}Th , 1.4×10^{10} years [127], means that we will not need to replace the ^{232}Th sources over the 5-year lifetime of the experiment.

The source material is commercially available thoriated tungsten wire that fits into small source capsules. The source activity is tuned by choosing the desired thorium composition of the wire (1% or 2% thorium by weight, nominal) and cutting it to a precise length. The source capsules accept wires of up to 5.1 mm in length and 0.50 mm in diameter. The activity of the uncut source wires was measured by the Berkeley Low Background Facility [128] to a precision of approximately 3%.

Wire electrical discharge machining makes it possible to cut the source wires reliably, with clean, hard corners and no fraying. Prior to machining, we place the source wires inside small aluminum holders, which are themselves placed inside a larger aluminum jig (see Figure 4.2). Perpendicular cuts are made through the entire jig, with the sources inside, to produce wires of the desired length. The aluminum holders ensure that the source wires remain in place before and during this cutting. We then dissolve the aluminum holders in a strong sodium hydroxide (NaOH) solution and clean the source wires with citric acid ($\text{C}_6\text{H}_8\text{O}_7$). The cutting procedure has a precision of a few hundredths of a millimeter, corresponding to an additional uncertainty on the string activity of approximately 1%.

After the cut source wires are prepared, we place them inside small copper capsules (8.0 mm in length, 1.6 mm in diameter) that are crimped onto Kevlar² strings. The capsule pitch along the string is 2.9 mm. This design allows the source strings to be flexible as they pass through the bends of the guide tubes in the cryostat, and the use of continuous strings minimizes the risk of capsules detaching from the source strings inside the cryostat. PTFE heat-shrink tubing around the capsules reduces friction and covers any edges on the source capsules that could impede their smooth motion

2. Kevlar is a registered trademark of E.I. du Pont de Nemours and Company.

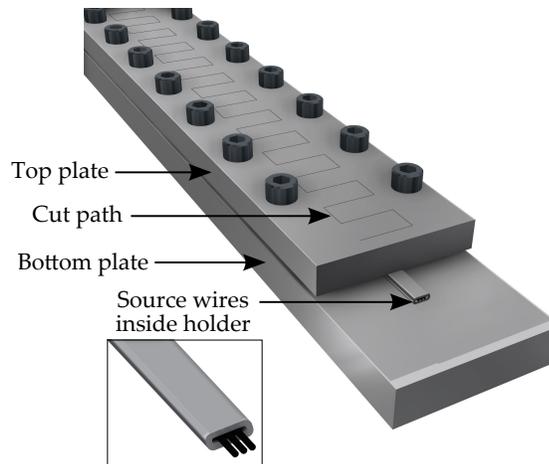


Figure 4.2: Rendering of the jig used for the wire electrical discharge machining of the source wires. The top and bottom plates and the source wire holder are aluminum.

inside the guide tubes. For the strings, we use Kevlar coated in PTFE³ because of its high tensile strength and low coefficient of friction. Kevlar’s low thermal conductivity at cryogenic temperatures [129] limits heat flow along the length of the string, and the PTFE coating reduces the fraying of the Kevlar string over time. The string has a diameter of 0.25 mm and a rated tensile strength of 110 N over the diameter of the string [130]. The capsule crimps are at least as strong; we are unable to slide the capsules along the string with a force less than 110 N, and any greater force breaks the string. A schematic diagram and a photograph of a source capsule are shown in Figure 4.3a.

The source strings are lowered under their own weight as they are deployed from the top of the cryostat. To this end, we add eight larger and heavier copper capsules (6.4 mm in length, 3.2 mm in diameter, 0.4 g each) to the bottom of each string; these capsules are also loaded with source wire. In addition, we attach a small PTFE ball to the bottom of each string to help the strings enter the guide tubes inside the cryostat (see Figure 4.3b).

There are 12 deployment positions for the source strings. The distribution of the

3. W.F. Lake Corporation. PTFE Coated Aramid Thread (R722-70). <http://www.wflake.com>

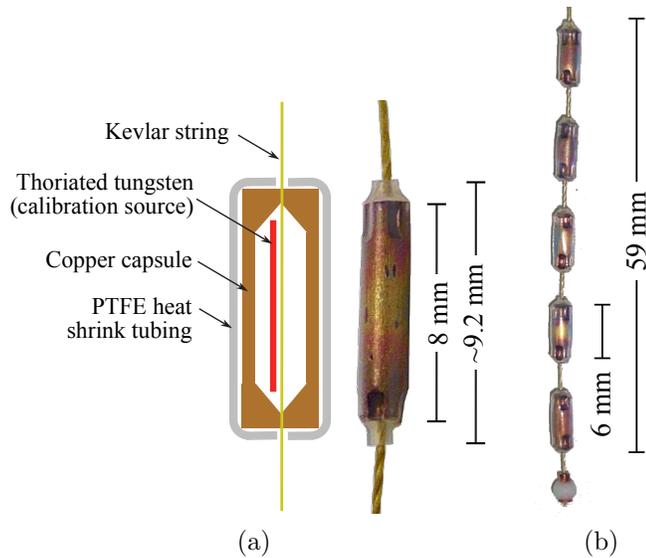


Figure 4.3: (a) Schematic and photograph of an assembled source capsule. (b) Photograph of five heavier bottom capsules and the PTFE guide ball at the bottom of a source string.

source capsules and intensities along each string matches the height of the detector towers and is optimized for event rate uniformity. The six inner strings, with sources distributed over 83 cm, are guided into the detector area and placed among the detector towers to irradiate the innermost towers. The six outer strings, with sources distributed over 80 cm, are allowed to hang outside the detector region, though still inside the cryostat’s lateral lead shielding, to provide additional gamma rays to the outer towers. The locations of the calibration strings in their deployed positions with respect to the detector towers are shown in [Figure 4.4](#).

In high-rate calibration tests with CUORE-like crystals, we determined that the optimal trigger rate for calibration is below 150 mHz per detector [125]. Above this rate, pile-up effects become dominant and the energy resolution is degraded. Thus, guided by Geant4-based Monte Carlo simulations [131] of the calibration sources in the CUORE cryostat, we selected summed source activities of 3.6 Bq of ^{228}Th for each of the six inner strings and 19.4 Bq for each of the outer strings; this configuration

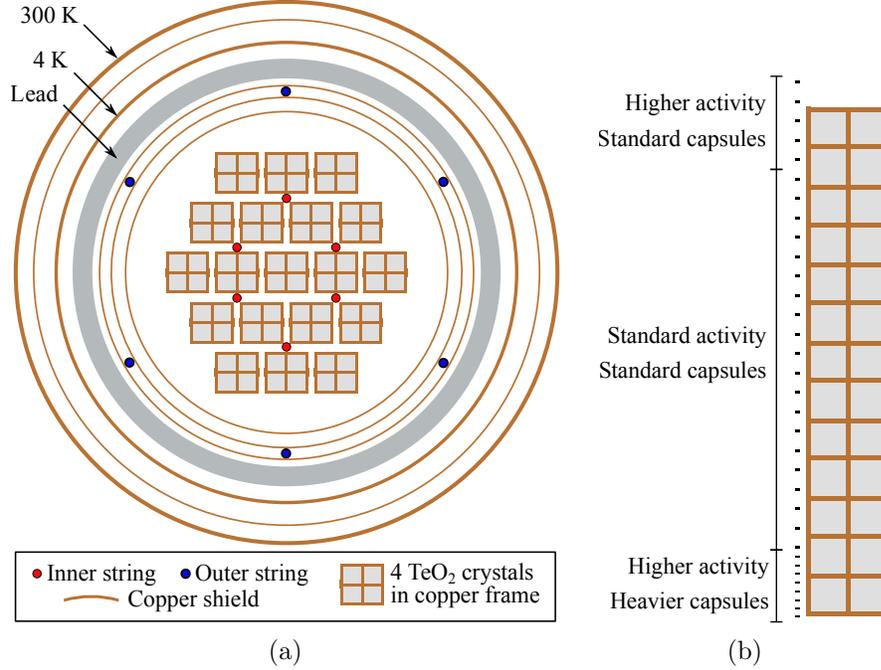


Figure 4.4: (a) Top-down cross-sectional view of the 19 detector towers, showing the locations of the source strings in their calibration positions. The outer strings (19.4 Bq of ^{228}Th each) are shown in blue; the inner strings (3.6 Bq each) are shown in red. The circles represent the cryostat vessels, which are identified as follows, starting from the outside: 300 K, 40 K, 4 K, lateral lead shielding, 600 mK, 50 mK, and 10 mK. (b) The height of the source capsules with respect to the towers, in their deployment position. Capsule size not to scale. The outer strings do not have the upper-most capsule in this figure.

results in an average trigger rate of approximately 100 mHz per bolometer. The simulated average event rates of the bolometers in each column of 13 crystals due to the inner strings alone, the outer strings alone, and all strings together are shown in [Figure 4.5](#). The effective event rate per bolometer from all source strings, after pile-up rejection, is significantly less than the 100 mHz trigger rate; for this simulation, we required that the times since the previous trigger and until the next trigger on the same bolometer must be greater than 3.1 s and 4.0 s, respectively, as in CUORE-0 [82].

The simulated calibration spectra produced by the sources are shown in [Figure 4.6](#). The source activity is divided among the 33 or 34 capsules on each outer or inner string, respectively, with the activity at the top and bottom of the capsule region of

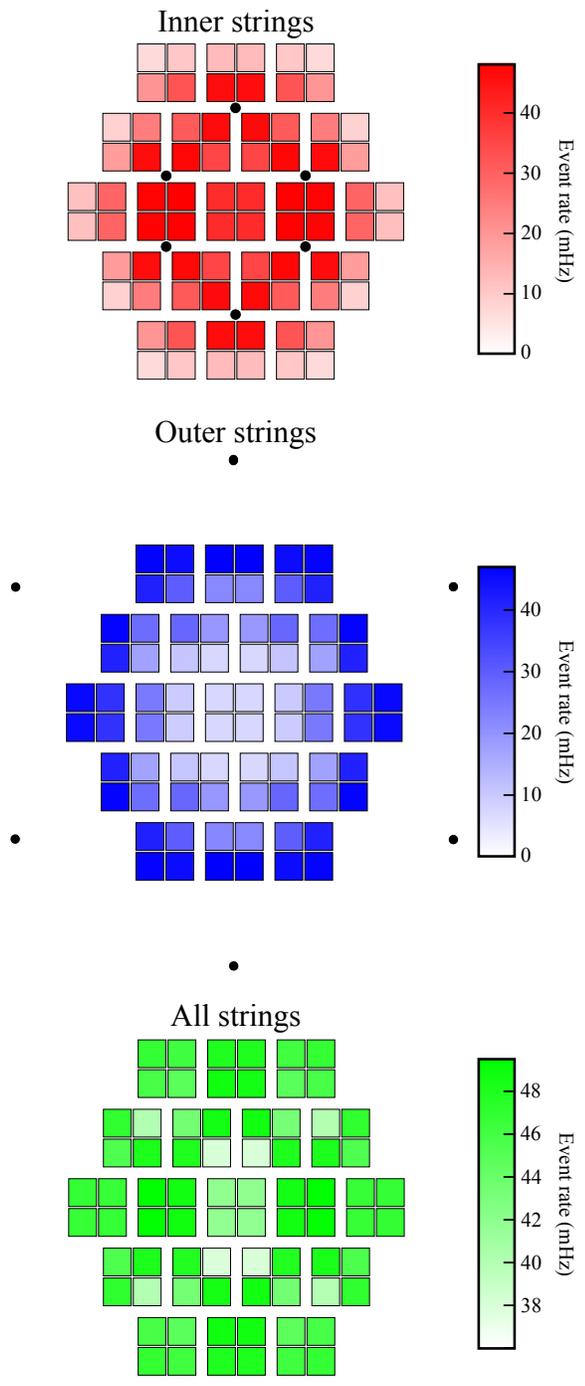


Figure 4.5: The simulated event rates per bolometer due to the inner strings alone (red), the outer strings alone (blue), and all strings (green), averaged over each column of 13 crystals. The total effective event rate is less than the sum of those from the inner strings and the outer strings as a result of pile-up rejection. Simulations performed by C. Davis.

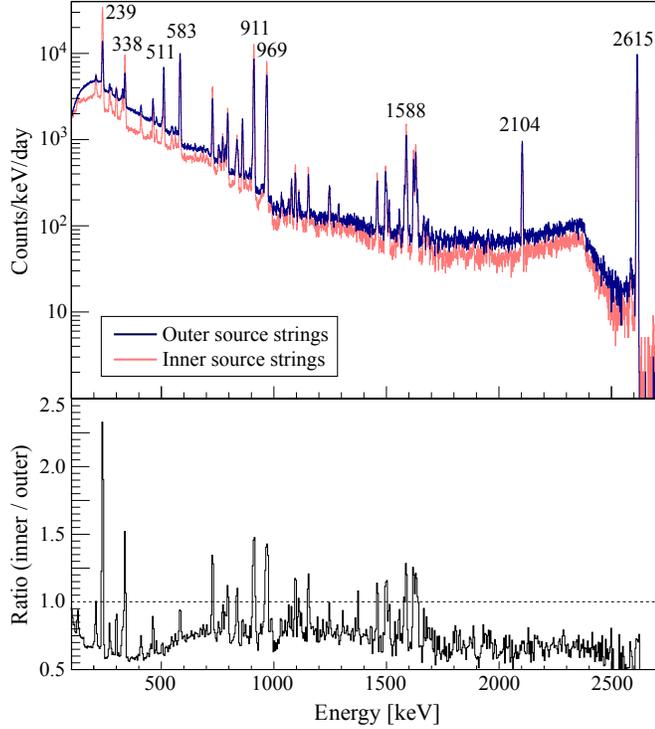


Figure 4.6: Top: A simulated CUORE calibration spectrum, summed over all channels. The spectra produced by the inner strings and outer strings are separated and overlaid, and the energies of important lines for calibration are labeled in units of keV. Bottom: The ratio of the counts due to the inner strings divided by the counts due to the outer strings. The outer strings have a lower peak-height-to-background ratio, particularly at lower energies, because of the presence of the copper vessels between the sources and the detectors. The ratio above 2615 keV is not shown because of low statistics. Simulations performed by C. Davis.

the strings higher than that in the middle to compensate for solid-angle effects. The capsules with increased activity are distributed over $\sim 25\%$ of the string length but account for $\sim 50\%$ of the total activity. This non-uniform activity distribution was adopted because the cryostat geometry prevents the simple extension of the active source length beyond the top and bottom of the towers. Calibration sources cannot be deployed below the bottom of the bolometer towers because of the presence of the 10-mK cryostat vessel surrounding the towers. Above the towers, the lead shielding would block the gamma rays from these low-activity calibration sources (as shown in [Figure 3.16](#)).

4.1.3 Motion control and monitoring hardware

Before and after each calibration period, the source strings are wound on spools above the cryostat, at room temperature. These spools are attached to stepper motors that turn to deploy the strings into or extract them from the cryostat. Four stainless steel enclosures (“motion boxes”) contain the 12 source string spools and the motors that drive them; each motion box is equipped with three motors and thus controls three strings. The motion boxes can be pumped down to vacuum or vented without affecting the cryostat vacuum, allowing the replacement of any source strings as necessary during the operation of CUORE. Each motor and spool in the motion box is instrumented with several sensors to ensure the fail-safe operation of the system inside the cryostat. Many of these sensors provide redundant information to mitigate the risk of the strings being in an unknown position in the cryostat or breaking as a result of excessive tension or other exceptional circumstances.

The motion boxes are stainless steel weldments that are mounted to gate valves on top of the 300-K plate of the cryostat. The main volume of a motion box is a welded box with dimensions of $126 \times 79 \times 607 \text{ mm}^3$. Along one face are three 6.75” CF flanges (125 mm in inner diameter) that protrude from this box. Mated with these are three CF flanges that hold the motors, sensors, and spools containing the source strings. The flanges are aligned such that the source strings on each spool feed directly into vertical guide tubes below. Opposite the motors in the motion box are three glass viewports, which we can uncover to observe the movement of a string while it is under vacuum, if necessary. The sides of the motion box contain two additional CF flanges: one hosts a vacuum pressure gauge, and the other is for the vacuum pumping line. A rendering of a motion box is shown in [Figure 4.7](#).

The motion boxes are sealed to the gate valves below them with fluoroelastomer O-rings. Below the gate valves, polyetheretherketone (PEEK) spacers electrically isolate the motion boxes, and thus the stepper motors, from the cryostat. The gate

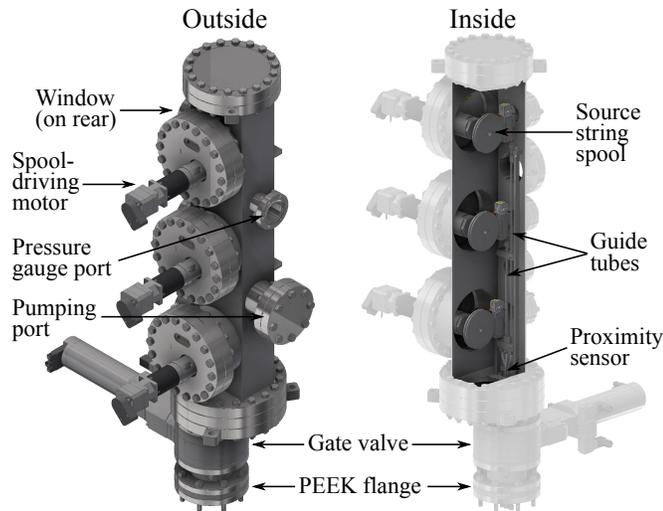


Figure 4.7: Rendering of a single motion box, which controls three strings. Four motion boxes are mounted above the 300-K plate of the cryostat.

valves and PEEK spacers are also sealed to the cryostat and to each other with fluoroelastomer O-rings, under a compression force provided by threaded stainless steel rods and nuts that are electrically insulated from the cryostat plate by nylon sleeves and washers.

Rotational motion is transmitted from the motor shaft outside the motion box vacuum to the source string spool through a rotary motion feedthrough⁴. This direct drive feedthrough includes a fluoroelastomer seal around the shaft; as the shaft rotates, the pressure of the fluoroelastomer against the shaft maintains the vacuum inside the motion box. The feedthrough is attached to the motor shaft outside the vacuum with a bellows coupling. A diagram of the motor, spool, and feedthrough mounted on a CF flange is shown in [Figure 4.8](#).

The stepper motors that raise and lower the strings⁵ have a 1.8° step angle and

4. MDC Vacuum Products. Direct Drive Rotary Motion Feedthrough (652000). <http://www.mdcvacuum.com>

5. Kollmorgen. CT Series Stepper Motor (CTP10ELF10MMA00). <http://www.kollmorgen.com>

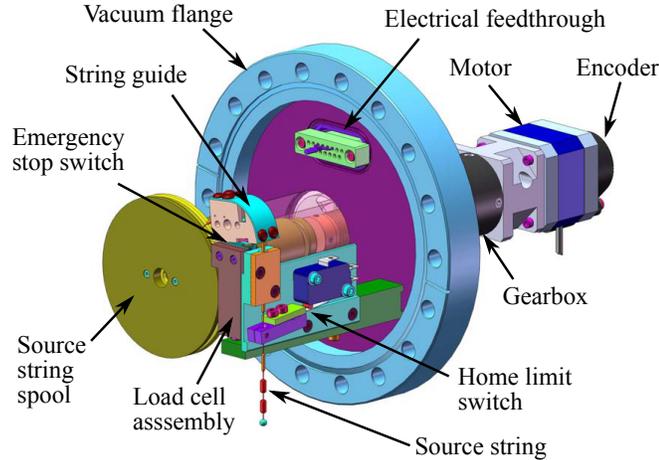


Figure 4.8: False-color diagram of the motor and spool controlling a single source string.

are connected to a 320:1 planetary gearbox (gear reducer)⁶. Thus, 64 000 steps of the stepper motor result in one full rotation of the source string spool. The strings can be moved smoothly at speeds ranging from below 5 mm/minute to above 500 mm/minute.

Information regarding the position and velocity of the strings in the cryostat is provided by a rotary encoder on each motor⁷. This 1000 cycles-per-revolution optical incremental encoder, combined with the gear reducer installed on each motor, yields a position resolution of approximately 1 μm , which is particularly useful for velocity measurements during the slowest parts of the deployment and extraction. By measuring the motion of the motor shaft, the encoders verify that the string is actually moving inside the cryostat and can alert the system operator of non-operational or malfunctioning motors or motor controllers.

As a string is lowered into the cryostat, we continuously monitor the string tension

6. Anaheim Automation. GBPN-064x-FS Series Planetary Gearbox (GBPN-0403-320-AA171-197). <http://www.anaheimautomation.com>

7. US Digital. E5 Optical Kit Encoder (E5-1000-197-IE-D-D-G-B). <http://www.usdigital.com>

using a load cell⁸. The load cell has a 3000- Ω thin-film strain gauge bridge and a range of 0 to 2000 g. A 1- μ F capacitor is placed across the signal leads from the load cell to reduce high-frequency noise from the nearby stepper motors. The load cell is connected to an in-line amplifier⁹ that is read out by a 12-bit analog-to-digital converter (ADC). The result is a tension reading that permits the detection of changes at a level of 5 mN or less. The commissioning of a source string path involves repeatedly raising and lowering a string through that path to determine the standard profile of the string tension as a function of the string position and direction. During all subsequent deployments, we ensure that the data from each load cell are consistent with its load cell profile; if there are any sustained deviations, the string is stopped, withdrawn, and redeployed. In testing with intentionally misaligned guide tubes, we have found that this method is able to reliably detect when a string has failed to enter a guide tube during deployment.

Each string enters and leaves its spool horizontally and passes over a PTFE string guide, which is mounted on the lever actuator of a micro switch. This guides the string into the opening of the guide tube below during the string deployment and, when extracting the string, triggers the motion system to stop if there is significant tension on the string. In normal operation, this switch should never be triggered. After the string passes over the PTFE guide, it passes through an aperture in the lever actuator of a second micro switch that we use as a home position indicator. As the string is withdrawn from the cryostat, when the first larger capsule at the bottom of the string hits this lever actuator, it triggers the micro switch.

Near the bottom of each motion box, an inductive proximity sensor¹⁰ detects and

8. Strain Measurement Devices. Miniature Platform Load Cell S251 (SMD2387-05-020). <http://www.smdsensors.com>

9. Honeywell. Bridge Based Sensor In-Line Amplifier (060-6827-04). <https://measurementsensors.honeywell.com>

10. Proxitron. Inductive Ring Sensor (IKVS-010.23-G-S4). <http://proxitron.de>

counts the copper source capsules entering and leaving the cryostat. We use this signal to ensure that all of the capsules have unspooled successfully and have entered the cryostat, and after calibration has been completed, we use it as an additional verification that all of the source capsules are fully withdrawn from the cryostat.

Our string position uncertainty is dominated by the effects of spooling and unspooling the source strings, and the source capsules in particular. To remove this uncertainty, we redefine the zero position of each string as its final source capsule enters the proximity sensor during the string deployment. This occurs before the first source capsules cross the 4-K plate of the cryostat and results in a final position uncertainty of approximately 1 mm in the lower regions of the cryostat.

All of the motors, motion box sensors, and temperature and pressure monitors are read and controlled by a dedicated server, which enables us to perform the calibration sequence and monitor its status remotely and automatically (see [Section 4.2](#)).

4.1.4 Source string guide tubes

The source strings pass through several different guide tubes en route to the detector region at 10 mK. Each source string has its own set of guide tubes; the strings are thus isolated from each other at all points to avoid any possibility of the strings becoming entangled or stuck. A schematic of the layout of the DCS guide tubes in the CUORE cryostat is shown in [Figure 4.9](#).

The DCS is fully integrated with a complex cryostat, and as such, the system's heat load must not exceed the available cooling power at any thermal stage. The guide tubes in the warmer regions are made of stainless steel, and the source strings are constructed from Kevlar; the low thermal conductivity of these materials limits the heat transfer between stages. During the deployment of a string into the cryostat, the source capsules are cooled before entering the most sensitive parts of the cryostat to avoid placing a large heat load on the colder cryostat stages. During string extraction,

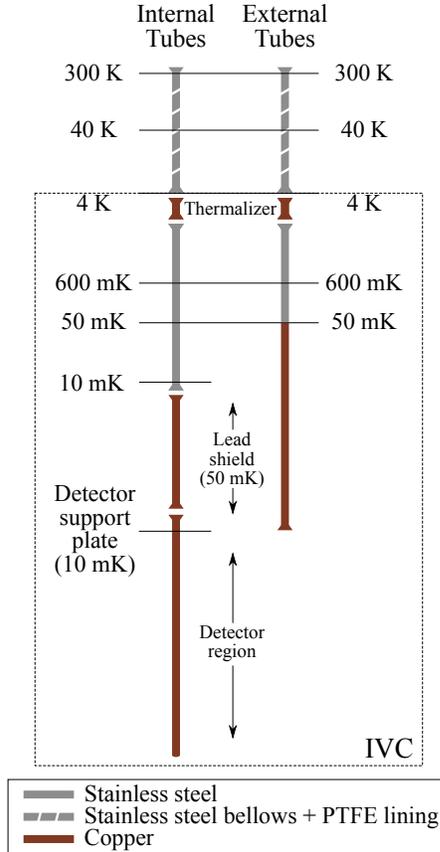


Figure 4.9: Schematic of the DCS guide tubes in the CUORE cryostat. All tubes are fully thermalized to the cryostat plates that they cross. The region surrounded by the dotted line represents the inner vacuum chamber. The thermalizer at 4 K is discussed in [Section 4.1.5](#).

the friction between the strings and the guide tubes accounts for the majority of the heat load; in addition to using PTFE-coated capsules and source strings and highly polished guide tubes, we minimize the frictional heat load by moving the strings very slowly (~ 10 mm/minute) in the coldest regions of the cryostat.

The CUORE cryostat has two vacuum chambers: the inner vacuum chamber (IVC) and the outer vacuum chamber (OVC). The walls of the IVC are the 4-K cryostat plate and vessel, whereas the room-temperature (300 K) plate and vessel form the boundaries of the OVC. The motion boxes and all string paths open directly into the IVC. In practice, this means that the motion boxes and guide tubes above



Figure 4.10: An S-tube assembly being installed in the cryostat. Each such assembly contains three string paths and connects the 300-K steel plate (top) to the 4-K copper plate (bottom). The tubes pass through a hole in the 40-K plate, where they are thermalized to the plate with a clamp (not shown).

the 4-K cryostat plate must be vacuum-tight, but those below the 4-K plate can be open to the IVC around them.

As the source strings are lowered below the gate valves in the motion boxes, the first guide tubes they encounter are S-shaped tubes (“S tubes”) that bring the strings from room temperature to the 4-K cryostat plate (see [Figure 4.10](#)). These S tubes have a coaxial design, with stainless steel vacuum-tight thin-wall formed bellows surrounding fluorinated ethylene propylene (FEP) tubing with an inner diameter of 4 mm. This design accommodates the relative motion of the 300-K, 40-K, and 4-K plates that arises from the cryostat suspension system and from the thermal contraction of the cryostat as it cools. The stainless steel bellows provide structural integrity with low thermal conductivity, and the FEP tubing helps to minimize the friction that the strings experience as they move through the S tubes. The bellows maintain a temperature gradient between 300 K and 4 K that helps to cool down the

string as it is lowered. The FEP tubing is also flared at the ends using a thermal forming process; these flares act as funnels, helping the strings to enter the S tubes without becoming caught on any edges. The S tubes pass through slots in the 40-K cryostat plate, where they are thermalized to the plate with copper braids squeezed around the bellows with an indium shim. The assembly is vacuum sealed to the 300-K cryostat plate with a fluoroelastomer O-ring and to the 4-K plate with indium.

At the bottom of the bellows, the strings enter a thermalizer that mechanically squeezes the source capsules on the string to cool them down to 4 K. We discuss this thermalization mechanism further in [Section 4.1.5](#).

Below the 4-K thermalizer, the strings enter stainless steel guide tubes. These tubes have an inner diameter of 4 mm and a wall thickness of 0.38 mm. They are mechanically and thermally anchored to the 600-mK cryostat plate with copper clamps and separated from the thermalizers by a gap of approximately 1 cm. This gap limits thermal conduction between the 4-K plate, to which the thermalizer is anchored, and the 600-mK plate; however, the tubes do maintain a small temperature gradient above 600 mK because of radiation from the nearby thermalizer. On both sides of the gap, there are funnels to guide each string into the next tube. Above the 600-mK plate, the six tubes that are on inner guide tube paths each have additional bends in the shape of a chicane to improve the contact between the source capsules and the walls of the tube and thus to improve the thermalization of the capsules (see [Figure 4.11](#)).

The guide tubes that contain inner source strings are made of stainless steel until they reach the 10-mK plate. They are split at the 600-mK plate for installation into the cryostat, but the upper and lower portions are clamped together with no gap between them. In addition to being thermalized to the 600-mK plate by copper clamps, they are also similarly thermalized to the 50-mK and 10-mK plates and thus maintain a temperature gradient from 600 mK to 10 mK. These guide tubes are also

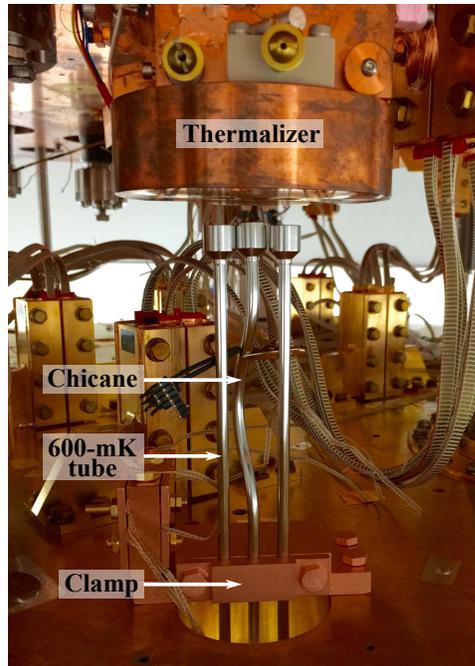


Figure 4.11: Photograph of the 600-mK guide tubes installed below a thermalizer. There is a visible chicane in the center tube. The tubes are thermalized to the 600-mK plate via the copper clamp at the bottom of the image. The copper clamp on the left thermalizes the thermometer wires attached to the tubes.

sloped from the 50-mK plate to the 10-mK plate ($45\text{--}59^\circ$ off vertical), encouraging source capsule thermalization with the walls of the tube. Below the 10-mK plate, there is a gap, and the inner source strings then continue into copper guide tubes inside the cryostat's lead shielding. Below the lead shield, there is another gap, and the strings then enter the detector-region guide tubes. The gaps above and below the lead shield allow the detector-region guide tubes to be vibrationally isolated from the lead shielding and from the other guide tubes. The inner source strings are fully contained in copper tubes at all times when they are in the detector region. These tubes capture thermal radiation from the source strings, prevent the contamination of the detector towers by the source strings, and ensure that the sources remain in the correct position for calibration.

All of the tubes below the 10-mK plate are vertical to minimize the friction be-

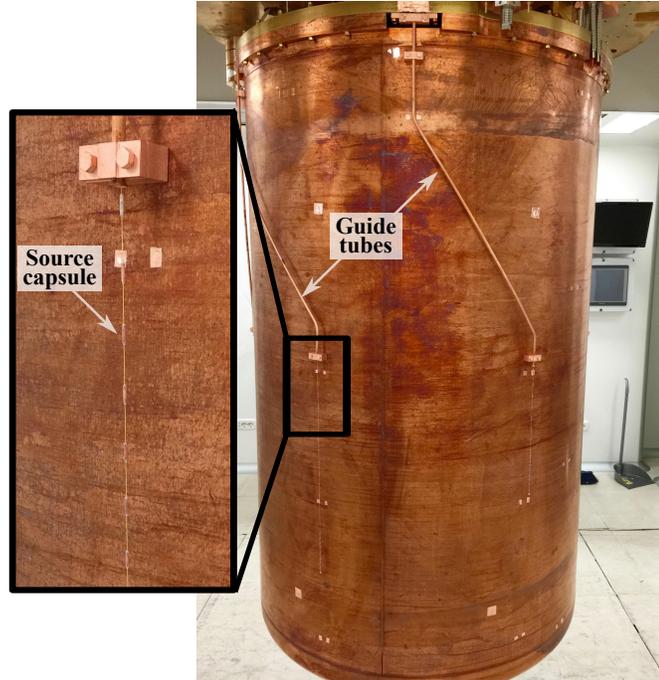


Figure 4.12: Helical outer guide tubes at 50-mK installed in the CUORE cryostat. A source string is shown in a partially deployed position as it is lowered through a guide tube.

tween the source strings and the tubes. The tubes in the detector region (6 mm in inner diameter) are also larger in diameter than those above to further reduce contact with the walls of the tubes and the resulting friction. The tubes that pass through the lead shielding and those in the detector region are composed entirely of electrolytic tough pitch copper subjected to tumbling, electropolishing, chemical etching and plasma etching [117] to respect the strict background radioactivity requirements of the cryostat near the detectors.

The outer source strings continue below the 4-K thermalizer in stainless steel tubes until they reach the 50-mK cryostat plate. These tubes are sloped between the 600-mK and 50-mK plates (34° off vertical) and are thermally coupled to both; thus, they maintain a temperature gradient that cools the strings as they are lowered. Below the 50-mK plate, helical oxygen-free high-thermal-conductivity (OFHC) copper guide tubes bring the strings down to the detector region (see [Figure 4.12](#)). These

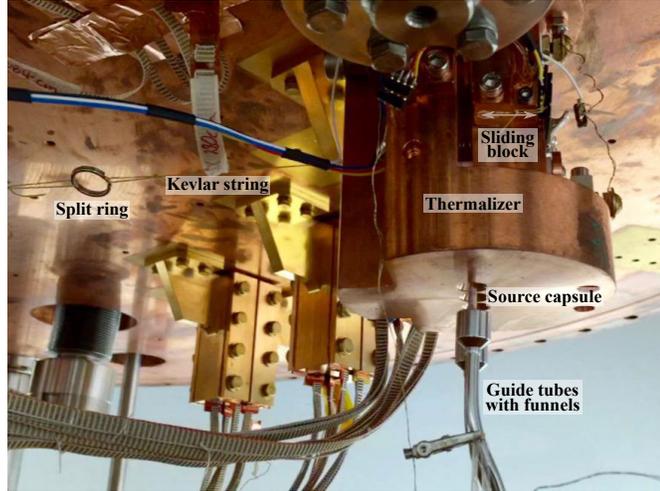


Figure 4.13: A thermalizer mounted in the CUORE cryostat. A capsule is visible entering the stainless steel guide tubes below the thermalizer.

tubes are thermally anchored to the outside of the 50-mK cryostat vessel and to the 50-mK plate, and they connect seamlessly to the stainless steel guide tubes above the plate. The slope of the helical tubes ($15\text{--}46^\circ$ off vertical) allows the strings to thermalize to 50 mK through contact with the tubes as they are lowered. Below the helical tubes, the outer strings are allowed to hang freely because they are spatially separated from the detectors by the 10-mK and 50-mK cryostat vessels.

4.1.5 4-K thermalization mechanism

One of the critical requirements of the DCS is to ensure that lowering the source strings into the cold cryostat does not change the operating temperature of the bolometers; that is, the temperature stabilization system must be able to compensate for the head load introduced by the source strings as they are lowered. It is therefore desirable to have an additional thermalization mechanism in place, as simple contact between the source strings and the guide tubes is not sufficient to cool down the source capsules efficiently. We have created such a mechanism, the 4-K thermalizer, which makes mechanical contact with each source capsule as it is lowered into the

cryostat. There are four such thermalizers in the cryostat; the three strings deployed from a single motion box pass through the same thermalizer. A photograph of a thermalizer is shown in [Figure 4.13](#).

Each thermalizer is well anchored, physically and thermally, to the bottom of the 4-K plate of the cryostat. This location was chosen because it is the coldest cryostat plate at which there is a large amount of cooling power available to the DCS (see [Table 4.1](#)). The thermalizer is primarily composed of a fixed copper body and a sliding copper block on PTFE guides; the sliding block is held apart from the body by a spring but is thermally coupled to it with a copper braid. The sliding block is attached to a Kevlar string, and when tension is applied to this string, the block is pulled against the spring and moves closer to the fixed body of the thermalizer, squeezing any capsules that are in between. This Kevlar string is connected to another Kevlar string that runs to the shaft of a room-temperature rotary vacuum feedthrough¹¹ at the top of the cryostat via a well-polished split ring, for ease of installation. A reproducible, fixed amount of tension is applied from outside the cryostat by means of a hanging mass. To open the thermalizer, a linear actuator lifts the hanging mass to release the tension on the Kevlar string, and the springs cause the sliding block to move away from the body of the thermalizer. Each thermalizer can squeeze two normal source capsules on a single string or four of the more closely spaced heavier source capsules at the bottom of the string. A three-dimensional rendering and a schematic of the thermalization system are shown in [Figure 4.14](#).

All four thermalizers are controlled by individual hanging masses and linear actuators outside the vacuum and at room temperature above the cryostat. The force from each hanging mass is transmitted to the string inside the cryostat that closes the thermalizer via a rotary feedthrough, in which a hermetic ferrofluid seal around

11. Ferrotec. Ferrofluidic Seal Thread Mounted Feedthrough (SS-250-SLAB). <https://www.ferrotec.com>

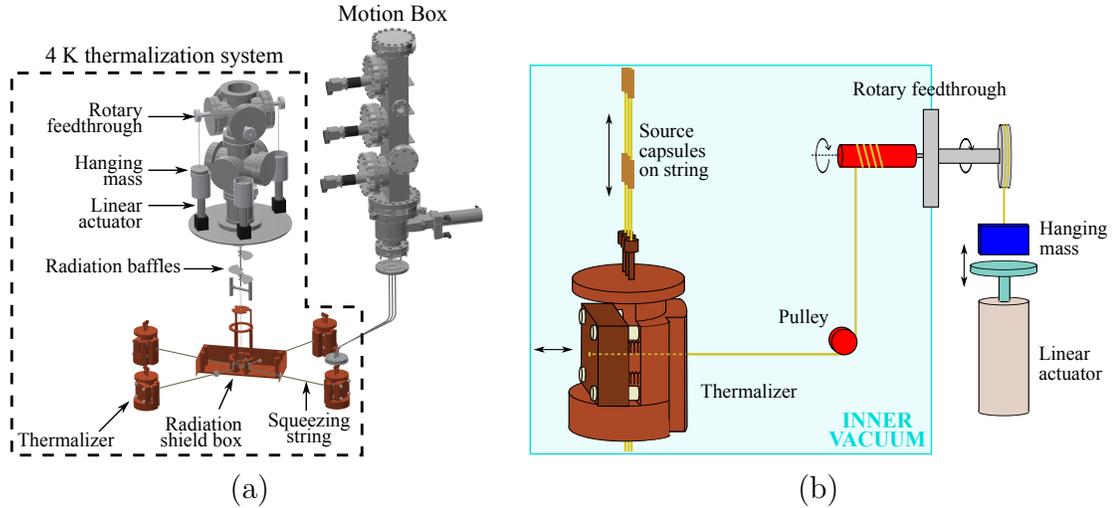


Figure 4.14: (a) Rendering of the 4-K thermalization hardware (boxed), with a motion box and an S-tube assembly containing source strings above one of the four thermalizers. (b) A schematic of one thermalizer and its accompanying hardware.

the rotating shaft maintains the vacuum inside the cryostat.

The four linear actuators are located together above the center of the cryostat, and the four Kevlar strings that cause the thermalizers to close pass through a single port on top of the cryostat (see Figure 4.14). This port opens directly into the IVC via stainless steel tubing and bellows, and it has line-of-sight access through the 4-K plate. The strings pass through this tubing and break out toward the four thermalizers below the 4-K plate of the cryostat, turning 90° from vertical to horizontal on small pulleys. Baffles in the tubing and a small copper box around the pulleys, internally coated with polyimide film, minimize the radiation that reaches below the 4-K stage. The strings emerge from small cutouts in the sides of the box via PTFE string guides.

The 4-K thermalization system is instrumented in several ways to ensure that an operating thermalizer has closed onto a capsule inside the cryostat. First, outside the cryostat, a potentiometer indicates the position of the linear actuator, with which we can verify that the hanging mass is being raised or lowered. Second, gold contact pads on either side of the sliding block of the thermalizer are grounded when the thermalizer

is either fully open or fully closed, whereas they are ungrounded (floating) when the thermalizer is between these extremes. Thus, we can verify that the sliding block has moved, and we can also distinguish between squeezing on a source capsule and squeezing on nothing or only the Kevlar string; if there is a source capsule in the thermalizer, it will not close fully, and if there is no capsule in the thermalizer, it will. Finally, the sliding block of the thermalizer exhibits a measurable rise in temperature when the cold thermalizer squeezes onto a warmer source capsule, which we measure using dedicated thermometers.

In tests at 4 K, we squeezed a PTFE-encapsulated silicon diode thermometer in the cryostat with a variety of different forces to test the cooling power of the thermalizer. We observed that the cooling time decreased as the force on the capsule increased, up to approximately 32 N. Larger forces did not result in significantly improved cooling times and are more likely to deform the capsules. Thus, a force of 32 N was selected for the final installation.

4.2 Electronic control and monitoring system

The DCS is a fully automated system, controlled and monitored by a dedicated server located next to the CUORE cryostat. A schematic of the DCS control system is shown in [Figure 4.15](#).

4.2.1 Hardware

The server communicates directly with four pieces of hardware: two temperature monitors, a vacuum gauge controller, and a PXI chassis¹². The PXI chassis contains

12. National Instruments. 5-Slot PXI Chassis with Integrated MXI-Express Controller (PXI-1033). <http://www.ni.com>

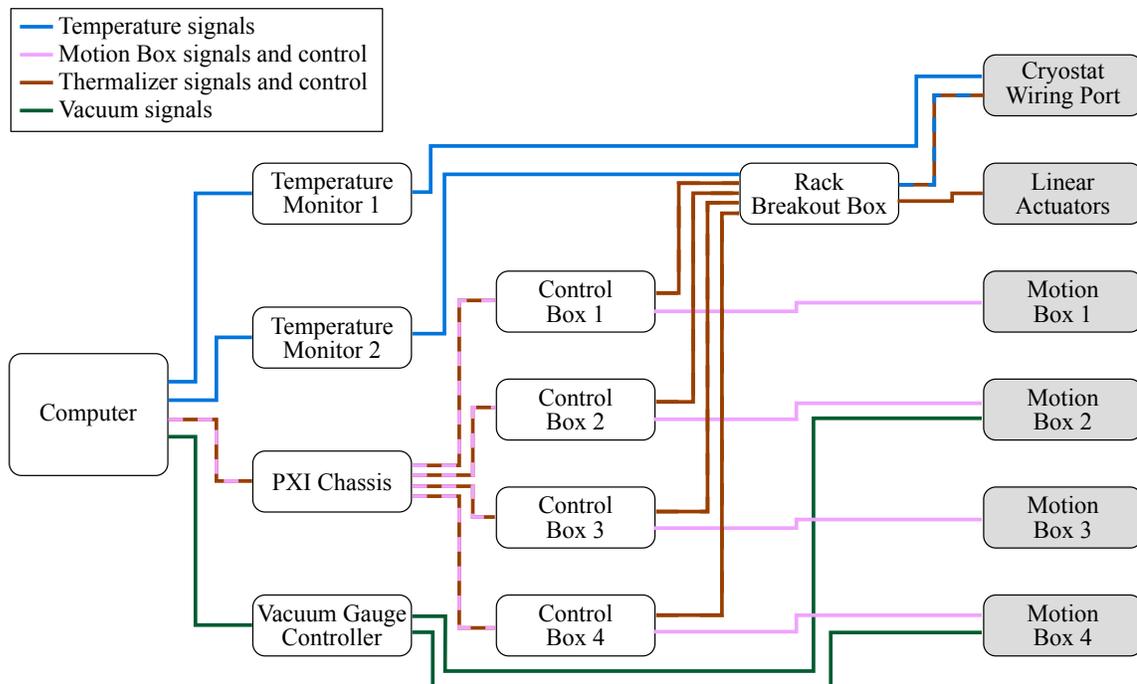


Figure 4.15: Overview of the DCS control system. White boxes represent components hosted in the control rack; shaded boxes represent components mounted on the 300-K plate of the cryostat. Because all of the motion boxes are part of the same vacuum system, there is no need for an individual gauge on each.

four stepper motor controller PXI cards¹³, each of which controls four motors and contains an integrated four-channel 12-bit ADC and 32 digital input/output lines. Each PXI card is in turn connected to a custom-designed control box, which interfaces with the DCS hardware and with the computer through these motor controller cards.

Each motion box and its accompanying linear actuator are powered by and communicate with the computer through a single control box. One control box contains four stepper motor drives¹⁴ and four 300-W power supplies, one for each of the three motors on the corresponding motion box and one for the linear actuator that controls the motion of the accompanying thermalizer. It also contains a National Instruments

13. National Instruments. 4-Axis Mid-Range Stepper/Servo Motion Controller (PXI-7340). <http://www.ni.com>

14. Kollmorgen. CTDC Stepper Drive (P70530). <http://www.kollmorgen.com>

Universal Motion Interface (UMI)¹⁵, which serves as a bridge between the motor controller in the PXI chassis and the hardware in the control box and in the cryostat. The UMI connects to the stepper motor drives, the three motor encoders on each motion box, and the potentiometer that provides position feedback on the linear actuator. It also passes to the PXI card the two micro switch signals from each of the motors in the motion box, the amplified load cell signals that reflect the string tensions, and the state of the thermalizer contact pads (grounded or ungrounded).

All electronic signals for the DCS from inside the cryostat pass through a single wiring port on the cryostat. There are signals from four Cernox thermometers¹⁶ that measure the temperatures of the sliding blocks of the thermalizers, signals from 12 Cernox thermometers that measure the temperatures of the 12 guide tubes mounted to the top of the 600-mK cryostat plate, and signals indicating the four thermalizer contact pad states.

The control boxes each have several output cables that run directly to their respective motion boxes on the cryostat. The linear actuator power and feedback potentiometer signals from all four control boxes, however, are consolidated into only two cables because the linear actuators are all located together near the center of the cryostat. This consolidation occurs inside a small rack breakout box. Also inside the rack breakout box, the thermalizer thermometer signals are directed to a temperature monitor, and the four thermalizer contact pad states are broken out to the four control boxes.

15. National Instruments. 4-Axis Universal Motion Interfaces for Industrial Applications (UMI-7774). <http://www.ni.com>

16. Lake Shore Cryogenics. Cernox thin-film resistance cryogenic temperature sensors (CX-1010-SD). <http://www.lakeshore.com>. Cernox is a trademark of Lake Shore Cryotronics, Inc.

4.2.2 Software

Custom software written with LabVIEW¹⁷ controls the entire DCS. With this software, we can operate the system remotely and perform fully automated deployments of the calibration sources into the cryostat.

On the front end, the software displays the complete system status in one of two modes. In the “visual overview” mode, the location and direction of motion of the source capsules in the cryostat are shown on a schematic diagram, along with indicators that show when each proximity sensor is recording a capsule and when each thermalizer is in a closed position. In the “details” mode, all parameters of the system that are measured and recorded by the software are presented.

The software receives input from the user as a series of text-based commands (“steps”). These steps are collected into procedures, through which the software sequentially progresses when it is in operation. The software is capable of executing multiple procedures in parallel and can synchronize procedures at various points to coordinate the deployments of multiple strings. Thus, we can perform fully automated deployments of the 12 source strings into the cryostat with predefined procedures running in parallel.

On the back end, the software contains a series of interlocks to ensure that the system is operating safely and to prevent the user from accidentally causing any harm to the hardware. Before a string begins to move, the software verifies that the gate valve between the motion box and the cryostat is open and that the thermalizer is also open. If the direction of the requested movement is upward, it checks that the two micro switches in the motion box (i.e. the one that indicates that the string is already in its home position and the one that indicates that the string tension is abnormally high) are not triggered. Before closing a thermalizer, the software also verifies that all

17. LabVIEW is a trademark of National Instruments.

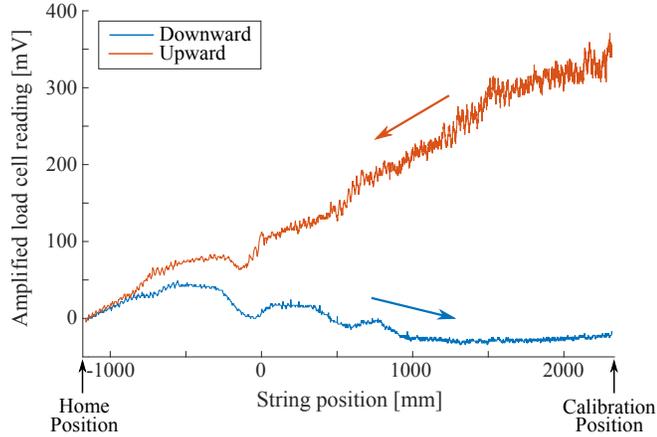


Figure 4.16: Load cell readings for a single string during deployment (downward motion) and extraction (upward motion). Downward motion, indicated with the blue arrow, is in the positive direction (left-to-right on this plot). Upward motion, indicated with the red arrow, is in the negative direction (right-to-left on this plot). The tension is much greater in the upward direction because of friction.

three strings that pass through that thermalizer are not moving, and before closing a gate valve, it verifies that all three strings of the corresponding motion box are in their home positions.

The software continuously reads the digital and analog signals from the DCS hardware to ensure that the system is operating as expected. The primary indication that a source string is being lowered or raised through the cryostat correctly is the string tension measured by the load cell. The load cell readings are continuously compared to the load cell profile, which represents the expected load cell reading as a function of the position of the string in the cryostat during a normal deployment. Because each string path through the cryostat is different, there is a unique load cell profile for each string. If the load cell value deviates from a predefined range around this profile for more than 10 seconds, then the string is stopped. The system operator can restart the string motion after assessing the situation, either continuing with or aborting the deployment.

A representative load cell plot recorded during the deployment of a single string

is shown in [Figure 4.16](#). As the string is first lowered, the tension is primarily determined by the amount of string that is hanging freely and thus increases approximately linearly. Near the position of -500 mm, the bottom of the string encounters the sloped portion of the S tube, which supports some of the string weight through friction. Once the string reaches higher position values (lower in the cryostat), all capsules have passed completely through the sloping guide tubes and the weight-supporting friction provided by the Kevlar string sliding on the guide tubes is approximately constant. When the string is being withdrawn and is moving in the reverse direction, the tension begins at a much higher value because of the change in the direction of the friction. As the string spools up, progressively less string is deployed in the cryostat, decreasing the total string mass and total friction of the string in the guide tubes, and the string tension steadily decreases. The tensions for both the upward and downward directions are similar near the home position because once the string is withdrawn from the S tubes, it is hanging freely and there is little friction. The fluctuations in string tension as the string moves in both directions are repeatable and result from the complex path that each string takes through the cryostat. In particular, the small periodic structures observed during upward motion result from the source capsules moving across bends and other structures as they are withdrawn.

The load cell readings are significantly steadier when the strings are being lowered compared with when they are being raised (see [Figure 4.16](#)). A very tight band of acceptable load cell values is necessary during string lowering to ensure that the string is moving correctly; the string is not actuated from both ends, so the load cell is our primary indicator of successful motion inside the cryostat. We use a band of ± 10 mV around the expected position from the load cell profile, whereas the profile itself varies over ~ 80 mV during motion in the downward direction. In the upward direction, the band can be much wider, as the only possible exceptional occurrence during a string withdrawal is that the string could become caught or tangled, in which

case the tension would significantly increase very rapidly.

Another indicator that the source capsules are successfully entering the guide tubes leading into the cryostat is the proximity sensor at the bottom of each motion box. The software counts the source capsules as they pass through the proximity sensors and into the cryostat, and it resets the corresponding string position value to 0 as the last capsule enters its proximity sensor. It also counts the capsules as they are withdrawn from the cryostat to ensure that all capsules are safely removed from the cryostat after the calibration has concluded.

Finally, the software records the temperatures of the sliding blocks of the thermalizers and the temperatures of the guide tubes anchored to the 600-mK cryostat plate. The thermometer on the sliding block of a thermalizer shows a characteristic spike when it squeezes on a warm capsule, which demonstrates that the thermalizer is working correctly. The thermometers on the guide tubes reveal how much heat was not removed by the thermalizer through their temperature rises, and these signals allow us to verify that the sources have actually entered the tubes.

4.3 Performance in commissioning runs

We have fully deployed all 12 source strings into the CUORE cryostat in cryostat commissioning runs. The cryostat base temperature during the source deployment tests was approximately 7 mK. Our testing was designed to find the maximum possible deployment and extraction speeds without causing the base temperature to rise above ~ 10 mK, which, in these tests, corresponds to a heat load on the mixing chamber of approximately 2–3 μW [132]. This provides the temperature stabilization system with sufficient leeway to maintain a constant base temperature above 10 mK during string deployment, leaving the detectors undisturbed.

4.3.1 Deployment

Cooling the sources down from room temperature to 4 K is a slow process. To this end, we left the source capsules in the S tubes for a full day before lowering them further into the cryostat, a process we refer to as “precooling.” The S tubes maintain a thermal gradient from 300 K to 4 K, and precooling of the strings is therefore achieved through contact with the sloped walls of the tubes. Simulations indicate that the background contribution in the $0\nu\beta\beta$ region of interest from having the sources in the cryostat but fully above the thermalizers is under 0.2% of the CUORE background budget. Thus, normal low-background data taking can continue during precooling.

The deployment procedures for the inner and outer source strings begin in a similar manner. Following precooling, the 4-K thermalizer squeezes pairs of adjacent capsules on the strings, each for a period of time ranging from 10 to 20 minutes. The string is lowered by 58 mm (twice the pitch of the capsules on the string) between each squeeze. Longer squeezes are required for the higher capsules, which are precooled to higher temperatures because of their relative positions along the thermal gradient in the S tubes. In total, the thermalizer squeezes last for approximately 4 hours per string.

For the inner strings, following the final thermalizer squeeze, the bottom of the source string is approximately at the level of the 10-mK cryostat plate. Before the strings reach the coldest stage of the cryostat (the mixing chamber level), they have very little effect on the cryostat base temperature. Indeed, a single inner string causes the cryostat to warm up only from 6.7 mK to 7.2 mK during this phase of deployment (see [Figure 4.17](#)). We are able to perform this thermalization sequence on four strings simultaneously while keeping the cryostat temperature below 9 mK. The outer strings have little to no effect on the cryostat base temperature during this phase of deployment.

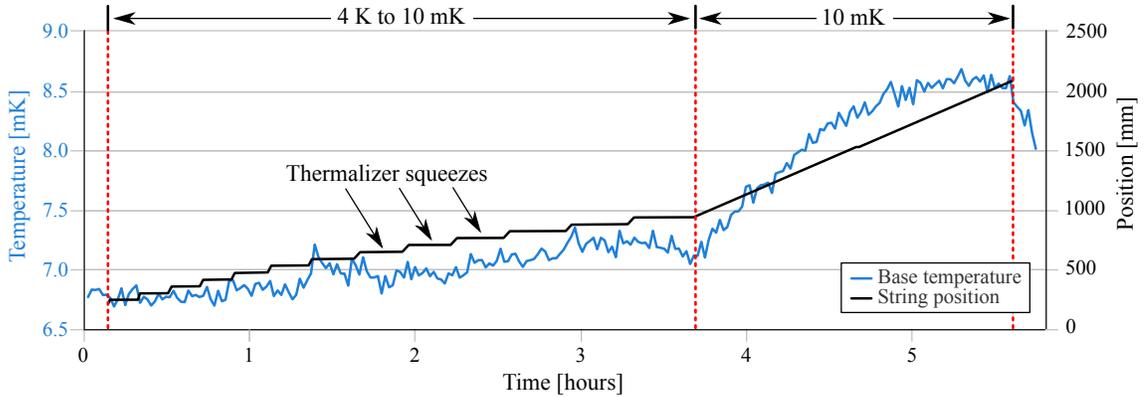


Figure 4.17: Cryostat base temperature and string position during a full inner-string deployment. Downward motion is in the positive direction. Two regions are identified; the first consists of progressively longer thermalizer squeezes while the bottom of the string moves from the 4-K cryostat plate to the 10-mK plate, and the second is the movement of the string through the 10-mK plate to its final deployment position.

The second deployment phase for each string involves lowering it from its final squeeze position to its full deployment position. It is in this phase that the inner strings are at or below the level of the mixing chamber on the 10-mK plate. We must bring the strings into this region very slowly to avoid exceeding the cooling power of the dilution refrigerator and causing the base temperature to spike. In testing, we determined that a single inner string moving at 10 mm/minute into this region raised the base temperature to 8.6 mK (see [Figure 4.17](#)).

The outer strings are deployed outside of the 50-mK vessel and thus do not impact the base temperature directly. The important parameter for the outer strings is the 50-mK vessel temperature (and, by proxy, the temperature of the heat exchanger of the dilution refrigerator), which begins to affect the base temperature if it rises and remains above $\sim 80\text{--}100$ mK. We determined that up to four outer strings can be simultaneously deployed at speeds of 15 mm/minute with only minor effects on the cryostat base temperature.

The most sensitive and time-consuming part of the deployment is the second deployment phase of the inner strings, which must be done serially. Following a par-

allelized four-hour first deployment phase, this second phase requires approximately two hours for each inner string, corresponding to a total of approximately 12 hours for all six inner strings. During the entire process, whenever a motion box is not moving an inner string, we can deploy one of its outer strings. The motion boxes are capable of moving all 12 strings simultaneously, but the 4-K thermalizers squeeze on all three string paths below a motion box simultaneously, thus making it impossible to move one string while squeezing on another from the same motion box. With a properly planned strategy, this is not a limiting factor in our deployment time.

4.3.2 Extraction

Following the calibration period, we extract the source strings from the cryostat. Because the sources have had time to equilibrate with the cryostat and fully cool, the heat load on the cryostat from the string extraction originates almost entirely from friction. Although there is no need to squeeze on the source capsules with the 4-K thermalizers during string extraction, the slow speeds required to avoid excessive frictional heating limit the total extraction time to only slightly less than the deployment time.

As during string deployment, the most time-consuming part of the string extraction is the movement of the inner strings in the region below the 10-mK plate. A single inner string extracted at 10 mm/minute, in parallel with an outer string extracted at 15 mm/minute, raised the base temperature to approximately 9 mK during the time when it was partially below the mixing chamber plate. Thus, we extract the strings in pairs, each consisting of one inner and one outer string, and can begin to extract the next pair of strings when the previous pair is above the mixing chamber. In the warmer parts of the cryostat, extraction can continue at significantly greater speeds without impacting the base temperature. A 12-string, 16-hour extraction is represented in [Figure 4.18](#).

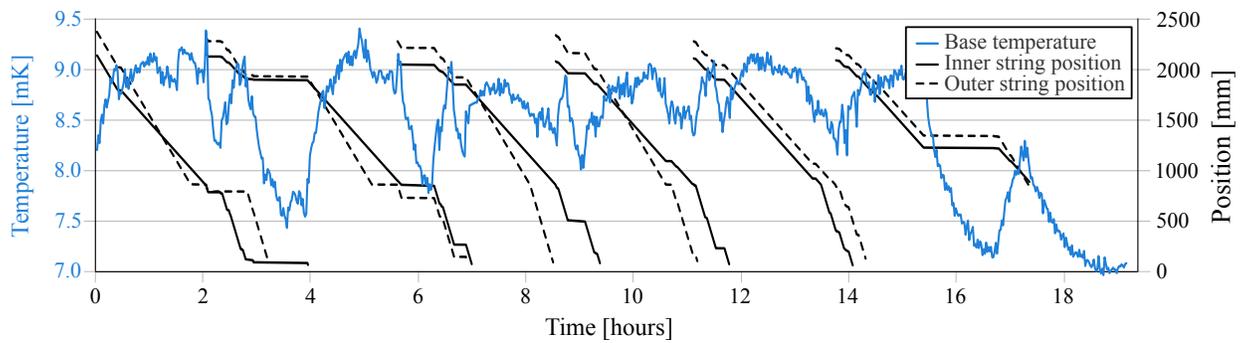


Figure 4.18: Cryostat base temperature and string positions during a 12-string extraction. The 12 downward-going black lines represent the positions of the 12 strings (upward motion is in the negative direction). In this procedure, the strings were extracted from their deployment positions to their precooling positions near the 4-K plate. Horizontal segments in the lines representing the string positions correspond to periods of no motion to observe the temperature response and recovery time of the cryostat.

Chapter 5

Data acquisition and processing

“As soon as we venture on the paths of the physicist, we learn to weigh and measure, to deal with time and space and mass and their related concepts, and to find more and more our knowledge expressed and our needs satisfied through the concept of number.”

— D’Arcy Wentworth Thompson

Compared with the challenges presented by many particle physics experiments, data acquisition with CUORE is, superficially, simpler. The CUORE detectors are slow, after all, with signal pulses lasting several seconds. And despite the fact that we have 988 individual detectors, each detector has only one signal channel: its temperature. The real challenge, however, is in the details. To achieve energy resolutions at the 0.2% level and energy thresholds orders of magnitude below our main signal region, we need extremely clean, clear signals. And with a heat capacity that depends on the cube of the detector temperature and a thermistor with a resistance that varies with temperature over orders of magnitude, we need a stable operating environment and a robust method of compensating for thermal gain variations.

All told, the techniques and electronics for acquiring and processing our data have been developed over two decades of experiments with macroscopic cryogenic bolometers [82, 110–113, 115]. Building on this experience, we have been able to

develop and implement the systems necessary to achieve our physics objectives with a significantly larger experiment than in the past.

5.1 Acquiring data

CUORE acquires data in runs that last up to 24 hours. Each run is one of three basic types: test runs, calibration runs, and physics runs. Test runs are used to measure or test various parts of the electronics that are described in the sections below. Calibration runs are when we acquire data with the calibration sources in place near the cryostat. Physics runs are when we are operating in a mode appropriate for observing $0\nu\beta\beta$ decay.

Calibration runs and physics runs together make up a dataset. Each dataset is generally composed of an initial calibration of a few days, followed by approximately one month of physics data, and finally a closing calibration. Calibration runs can serve as part of both as the closing calibration of one dataset and the initial calibration of a subsequent dataset.

5.1.1 Electronics

The CUORE electronics provide a low-noise signal readout for the 988 detector channels, and they control the thermistor bias and heater currents. They are designed for room-temperature operation, with the front-end boards located directly above the cryostat to minimize the distance that the signal travels before processing.

The thermistor signals run from the top of the detector towers to the top of the cryostat in twisted-pair constantan wires. Outside the cryostat, low-noise linear power supplies [133] bias the thermistors to convert their temperature-dependent resistances into a measurable voltages. This voltage is passed through a differential voltage-sensitive preamplifier [134] to a programmable gain amplifier. Next, the signal passes

through an antialiasing filter to prepare it for digitization; we use six-pole Bessel–Thomson low-pass filters with a roll-off of 120 dB/decade [135]. Following this, the signals are digitized by 18-bit differential ADCs produced by National Instruments¹. We do not use any hardware triggers, and instead acquire samples continuously from all channels at a rate of 1000 samples/second.

When operating CUORE, we need to optimize the “working point” of each detector; that is, we set each thermistor’s bias voltage, gain, and offset. We first set the bias voltage in order to maximize the detector sensitivity. At low bias voltages, the thermistor signal height is approximately proportional to the bias. At higher voltages, however, the thermistor begins to heat up and its temperature increases, thus decreasing its resistance and sensitivity. We set the bias voltage at the optimal point between these two extremes. We can then set the gain of the front-end electronics such that the signals use the full dynamic range of the ADCs, and finally the offset, such that the signal is inside the allowed ADC voltage range (−10 V to 10 V). We set the offset such that the baseline voltage is negative and set the gain such that the amplitude of the 2615 keV peak is approximately 1–2 V, which allows us to record events with energies of 10–20 MeV.

The heater pulses are also sent through constantan wires to the detector towers, where they cause silicon wafers to deposit fixed amounts of heat periodically into the crystals. The precise energy of these pulses is different for each channel and is not used for energy calibration; the pulses are used only for gain stabilization [136]. They are generated by extremely temperature-stable pulser boards [137], one per column of 13 crystals, and they are square waves with a very precise width, with a spread of ~ 1 ppm [138]. We use heater pulse widths of 1 ms, much less than the rise time of the detector signal pulses (100 ms), to imitate a standard energy deposition.

1. High-Accuracy M Series Multifunction DAQ (PXI-6284). <http://www.ni.com>

5.1.2 Signal pulses and triggering

CUORE data is acquired and saved continuously with APOLLO [139], software written in C++ by the CUORE Collaboration and based on the ROOT framework [140]. We use APOLLO for a wide variety of tasks related to data-taking, including electronics setup, data monitoring and acquisition, and event triggering.

We use two trigger algorithms for CUORE. The first, the Derivative Trigger, is used for the standard analysis. It is based on a simple principle: we trigger every time the signal amplitude derivative rises above a certain level for a certain amount of time. The threshold for the trigger is set on each channel to be as low as possible while preventing an excessive number of triggers on noise spikes. The second, the Optimum Trigger, is optimized for low-energy pulses, which are used for dark matter and other rare events searches [141]. The Optimum Trigger uses the noise power spectrum and average pulse shape, acquired with the Derivative Trigger or from previous data, to filter the incoming data, improving the triggering performance for small pulses.

A full CUORE waveform consists of the 3 seconds before a trigger and the 7 seconds following a trigger, for a total of 10 000 samples. We acquire the full pulse shape and the baseline waveform before the pulse for future analysis and event selection.

In addition to signal events, we also have heater events. There is a heater pulse on each channel every 380 s. These pulses are flagged and triggered by APOLLO automatically, with input from the pulser boards, but are also flagged with the automatic trigger, which allows us to test the efficiency of our triggering.

Finally, we have noise events. To evaluate the noise behavior of the detectors, we trigger on every channel in a tower simultaneously every 80–100 s. These waveforms allow us to build up a noise power spectrum for each detector and allow us to study correlated noise between detectors. We use these noise events to determine our baseline energy resolution at 0 keV as well.

5.2 Processing data

Now that we have triggered waveforms, we can begin to process the data. For this, we use DIANA, another software package written in C++ by the CUORE Collaboration and based on ROOT. DIANA is highly modular, and each part of the analysis is performed with a specific DIANA module. The modular structure allows each part of the analysis to be developed, tested, and run independently. The modules can add information to each event (e.g. the amplitude of the pulse or its energy) and can also calculate global quantities that are stored in an analysis database.

5.2.1 Preprocessing

We refer to the first step of the data processing as “preprocessing.” We begin by analyzing the the detector baseline. In particular, we look at the first 75% (2.25 seconds) of the pre-trigger time in each waveform; that is, we look from 3.00 s to 0.75 s before the trigger. We measure the baseline voltage as the average voltage during this period, which is a proxy for the starting temperature of the bolometer. We measure the baseline slope by fitting a line to these 2.25 seconds of data, which we we use in the pulse shape analysis. And we measure the baseline noise as the RMS of the voltage around this best-fit line, which is a measurement of the detector performance. An example of this is shown in [Figure 5.1](#).

The other function of the preprocessing is to count the number of pulses in each event. We do this by taking the smoothed derivative of the pulse waveform and looking for positive peaks in the derivative plot that far exceed the RMS of the derivative. Any peak whose derivative is five times the derivative RMS is counted as a separate pulse, and the number of pulses is stored with the event.

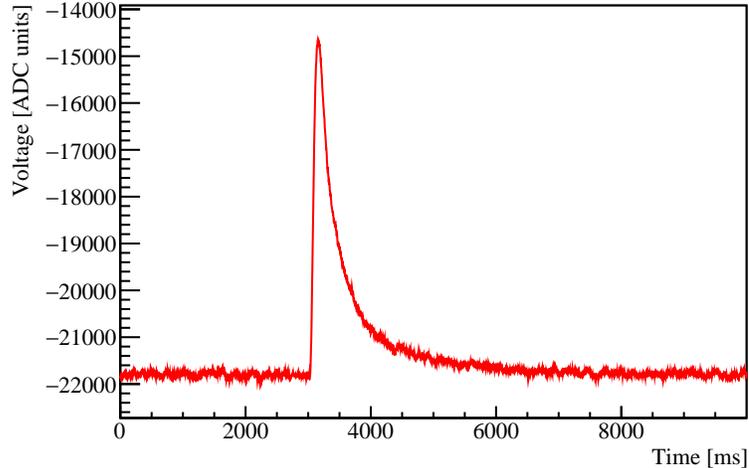


Figure 5.1: A representative raw waveform from a single detector, starting 3 seconds before a pulse trigger and ending 7 seconds after the trigger. This waveform has a baseline of $-21\,800$ ADC units, a baseline RMS of 60 ADC units, and a baseline slope of approximately 0. The window contains 1 pulse. These parameters are determined in the preprocessing using the first 2250 ms of the waveform.

5.2.2 Pulse amplitude evaluation

The raw digitized waveforms are noisy, and it is desirable to reduce the noise before evaluating the pulse heights. Because we find that evaluating pulse heights yields a better energy resolution than evaluating pulse areas, we do not evaluate the pulse areas at all, and the uncertainty in our determination of the pulse height directly affects our energy resolution.

The pulse $p(t)$ can be described, to a good approximation, as

$$p(t) = a \cdot s(t - t_0) + n(t) + b, \quad (5.1)$$

where a is the signal amplitude, $s(t)$ is the signal response function, $n(t)$ is the noise, and b is the detector baseline. The signal response function $s(t)$ is, for the most part, independent of energy, so the signal amplitude a is the only energy-dependent quantity. To evaluate a , we need to know the signal response function and filter out the noise as much as possible. Generally speaking, to achieve a resolution of 5 keV

at 2615 keV, we need to know the signal amplitude to an accuracy of 0.2%.

The first requirement is to determine $s(t)$ for each channel in each dataset, which we call the “average pulse” of that channel. We build the average pulse from the events with an amplitude greater than 40 times the noise RMS. Assuming the noise $n(t)$ is random, averaging together many pulses removes the majority of the noise and allows us to compute a reasonable $s(t)$.

We now want to filter out the noise from individual pulses as much as possible to decrease the uncertainty on the pulse amplitude. To do this, we use one of two filters: a standard optimum filter or a decorrelating optimum filter. The optimum filter is based on the matched filter technique and is the linear filter that maximizes the signal-to-noise ratio given a stochastic background. The filtering is done in frequency space, where frequencies tied to noise are reduced in weight and those tied to signal are amplified. Specifically, we implement a digital filter with transfer function

$$H(\omega) = \frac{S^*(\omega)}{N(\omega)} e^{-i\omega t_M}, \quad (5.2)$$

where $S(\omega)$ is the Fourier transform of the signal response function $s(t)$ from the average pulse, $N(\omega)$ is the average noise power spectrum, and t_M is the time of the pulse maximum. The average noise power spectrum $N(\omega)$ is determined from the random triggers on each channel every 80–100 s.

A generalization of the optimum filter, which we call the decorrelating optimum filter, is designed to remove correlated noise between detectors [79]. To do this, we build a covariance matrix,

$$C_{ij}(\omega_p, \omega_q) = \langle N_i(\omega_p) N_j^\dagger(\omega_q) \rangle, \quad (5.3)$$

which relates the noise in channel i at frequency ω_p to the noise in channel j at frequency ω_q . We then subtract the correlated frequencies from a particular channel

before applying the optimum filter. Although the performance decorrelating optimum filter should match or exceed that of the standard optimum filter in theory, it was only superior for some channels in practice. For this first analysis, we use only the standard optimum filter as we work to optimize the decorrelating optimum filter for future datasets.

Once we have the filtered pulse, it is straightforward to extract the pulse amplitude. Starting from the trigger time, we look for the first local maximum in the pulse. We take the maximum point and one point on either side of the maximum and fit a parabola to these three points. We define the pulse amplitude as the maximum value of this parabola.

5.2.3 Thermal gain stabilization

Because the thermistor resistance and crystal heat capacity are strongly temperature-dependent (see [Section 3.1](#)), the signal amplitude at a given energy is affected by the temperature-dependent gain. To correct for this, we need to create a mapping from detector temperature (i.e., the baseline voltage) to gain.

The most straightforward way to create this mapping is by using the heater pulses. These pulses, which occur on each channel every 380 seconds, remain constant in energy despite the changing baseline. They are also automatically flagged as heater pulses by APOLLO, making it simple to pick them out of the data. We plot the pulse amplitude vs baseline voltage for all of the heater events in the dataset in a particular channel and perform a linear fit to the scatter plot (see [Figure 5.2](#)). This linear fit is then a mapping from pulse baseline to mean heater pulse amplitude. We then calculate a unitless stabilized amplitude for each signal event, with

$$\frac{\text{Pulse stabilized amplitude}}{5000} = \frac{\text{Pulse raw amplitude}}{\text{Mean heater amplitude at pulse baseline}}. \quad (5.4)$$

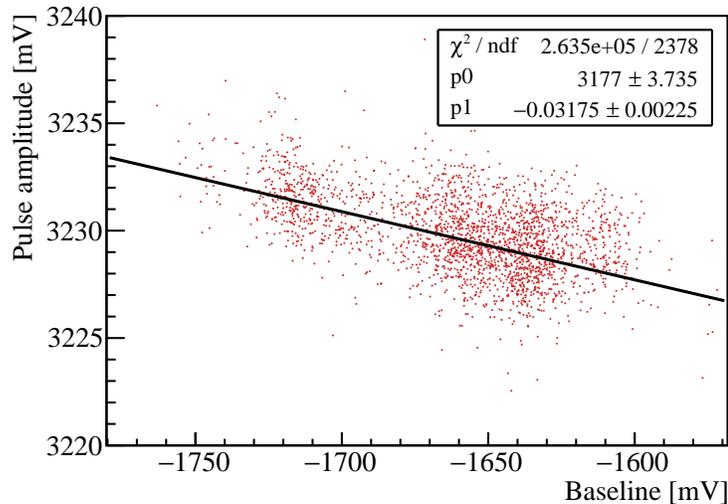


Figure 5.2: Scatter plot of heater pulse amplitude vs. pulse baseline for channel 108 in dataset 3018. The linear fit (in black) to the scatter plot approximates the mean heater pulse amplitude at a given detector baseline.

In essence, since we do not know the true heater energy, we arbitrarily fix the heater pulses to be centered on 5000 stabilized amplitude units and scale the other pulses accordingly.

An alternate method of gain stabilization is with the 2615 keV events from the calibration data instead of with the heater pulses. This is necessary for channels in which the heater has failed and optional elsewhere. The concept is similar to the heater stabilization, except we must now search for the 2615 keV events in the calibration data, as they are not flagged automatically. We perform this stabilization procedure for every channel and return to the question of choosing a stabilization procedure in [Section 5.3.4](#).

5.2.4 Calibration

We now have a gain-stabilized amplitude for each event, but no indication of the event's true energy. For this, we need to use the known energy of lines in the calibration spectrum to generate a map from stabilized amplitude to true energy for each

Energy [keV]	Source	Energy 2 [keV]	Source 2	Amplitude ratio
238.632 ± 0.002	^{212}Pb	—	—	—
338.320 ± 0.003	^{228}Ac	328.000 ± 0.006	^{228}Ac	0.262
583.187 ± 0.002	^{208}Tl	—	—	—
911.204 ± 0.004	^{228}Ac	904.20 ± 0.04	^{228}Ac	0.030
968.971 ± 0.017	^{228}Ac	964.766 ± 0.010	^{228}Ac	0.316
2614.511 ± 0.010	^{208}Tl	—	—	—

Table 5.1: The peaks used for calibrating the detectors. “Energy 2” and “Source 2” refer to the energy and source of weaker, nearby contaminant peaks, and “Amplitude ratio” is the ratio of the intensity of the contaminant peak to that of the primary peak. Energies and branching ratios are from Nuclear Data Sheets for ^{208}Tl [142], ^{212}Pb [143], and ^{228}Ac [144], all of which are daughters of ^{232}Th .

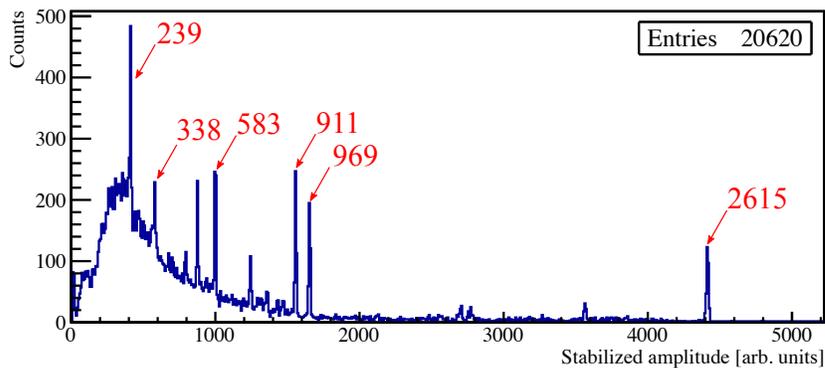


Figure 5.3: Uncalibrated stabilized amplitude spectrum for one channel in one dataset (channel 108 in dataset 3018). The six γ lines used for calibration are marked.

channel in each dataset. We call this map a “calibration function.” This is necessary even for the channels stabilized with the 2615 keV line, for which the precise energy is known, because the calibration function is slightly non-linear.

Starting from the stabilized amplitude spectrum, DIANA automatically locates the 6 lines used for calibration (see Table 5.1 and Figure 5.3). It first looks for the highest-energy peak, which is from the 2615 keV line of ^{208}Tl . With a guess for the location of the 2615 keV peak, DIANA then guesses where the strongest peak in the spectrum, at 239 keV from ^{212}Pb , would be located if the calibration function were linear from the origin to the 2615 keV peak. It refines this guess to be the stabilized

amplitude of the tallest bin in an area around this original guess. Now with three points (the origin, the stabilized amplitude and energy of the 2615 keV peak, and the stabilized amplitude and energy of the 239 keV peak), DIANA begins to build up a temporary calibration function. For each subsequent peak, the software makes a first guess as to where the peak should be located based on the calibration function, and then refines the guess to be the stabilized amplitude of the tallest bin near its original guess. It then updates its temporary calibration function and looks for the next peak. When the module has a refined guess for the location of each peak in the spectrum, the temporary calibration function is discarded.

It is worth noting that the automatic peak-finding in the spectrum works quite well, but occasionally DIANA fails to locate the correct peaks in the spectrum, particularly on poorly-performing channels. For these cases, we have a dedicated graphical user interface that allows the user to click on the locations of the relevant γ lines in the uncalibrated spectra. And following the first calibration, as long as the detector conditions have not changed significantly, we can use calibration functions from previous datasets to point DIANA to the approximate location of the γ lines in subsequent datasets. The combination of the automated peak finding, especially when aided by previous calibration functions, and the graphical user interface for manually selecting the peaks allows us to calibrate all of the channels of CUORE reliably and efficiently.

The next step is to fit the spectrum around the guessed location of each peak to get an accurate value for the mean stabilized amplitude of the peak. We do an unbinned extended maximum likelihood fit to fit each of the peaks to the following function:

$$f(x) = n_p p(x) + n_b b(x) + n_c c(x) \tag{5.5}$$

where $p(x)$ is the normalized peak shape, $b(x)$ is the normalized background continuum, and $c(x)$ is the normalized multiscatter Compton background. For those peaks with a nearby “contaminating” peak, we use $p(x) = r p_1(x) + (1 - r) p_2(x)$, where r is

fixed based on the branching ratio of the two lines, and the parameters of $p_2(x)$ are fixed to those of $p_1(x)$, with the exception of the stabilized amplitude, which is fixed to the ratio of the energy of the peaks. This ratio ignores the slight non-linearity of the calibration function — the ratio of the stabilized amplitude of the two peaks is not equal to the ratio of the energies of the two peaks — although this is a negligibly small effect for nearby peaks and we can safely ignore it.

For the peak shape, we use the Crystal Ball function

$$p(x) = N \cdot \begin{cases} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), & \text{for } \frac{x-\mu}{\sigma} > -\alpha \\ A \cdot \left(B - \frac{x-\mu}{\sigma}\right)^{-n}, & \text{for } \frac{x-\mu}{\sigma} \leq -\alpha \end{cases} \quad (5.6)$$

with

$$A = \left(\frac{n}{|\alpha|}\right)^n \cdot \exp\left(-\frac{|\alpha|^2}{2}\right), \quad (5.7)$$

$$B = \frac{n}{|\alpha|} - |\alpha|, \quad (5.8)$$

and N a normalization constant. The Crystal Ball function is a function with a Gaussian core and a power-law low-end tail (see [Figure 5.4](#)). It has four parameters: μ , which we use as the fitted location of the peak; σ , the Gaussian width of the peak; α , a measure of how far into the tail of the Gaussian we transition to the power-law tail; and n , a measure of the steepness of the power-law tail. Although the Crystal Ball function does not describe our peak shape perfectly, it tends to minimize the effect of our low-energy tail on the fitted mean of the center of the peak compared to using a simple Gaussian function. For the 2615 keV peak, we treat μ , σ , α , and n as free parameters; for the other peaks in the spectrum, we use the value of α and n from the 2615 keV peak and have only μ and σ floating.

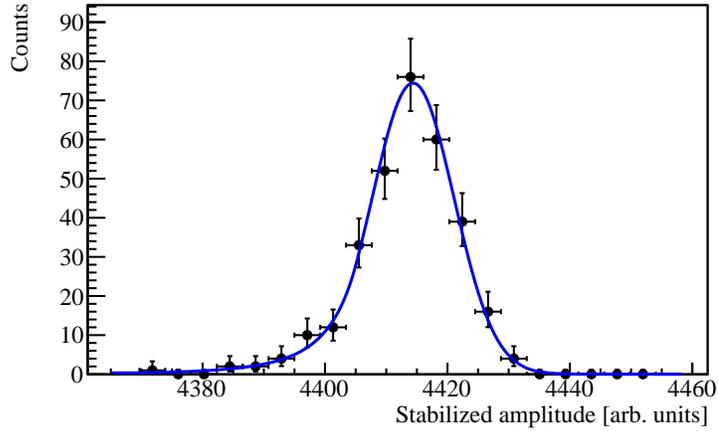


Figure 5.4: Crystal Ball fit to the 2615 keV line in the uncalibrated stabilized spectrum of one channel in one dataset (channel 108 in dataset 3018), showing the Gaussian core with a power-law low-energy tail.

We use a linear background continuum

$$b(x) = a_0 + a_1x \quad (5.9)$$

to represent a generic background of γ rays and α particles. In addition, we use a multiscatter Compton background

$$c(x) = \frac{1}{2} \operatorname{Erfc} \left(\frac{x - \mu}{\sqrt{2}\sigma} \right). \quad (5.10)$$

This is essentially a step function smeared by the detector resolution, which represents a background of γ rays that scatter multiple times in the detectors before exiting and γ rays with slightly degraded energy due to scattering in the detector shielding before hitting the detectors. Note that the μ and σ in $c(x)$ are the same as the μ and σ in $p(x)$, above.

At this point the module has a value for the stabilized amplitude, with uncertainty, of as many of the 6 calibration lines as it was able to fit. With this, the module makes a scatter plot of the energy vs. stabilized amplitude of each peak. By fitting a second

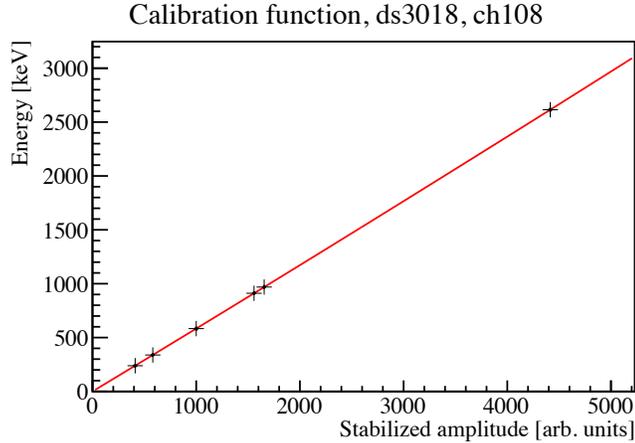


Figure 5.5: A representative CUORE calibration function, mapping stabilized amplitude to true energy for channel 108 in dataset 3018. The second-degree polynomial fit is in red, and the 6 points used to determine the calibration function are marked in black.

degree polynomial with zero intercept ($y = ax^2 + bx$) to this plot, we obtain the final calibration function for the particular channel and dataset we are working with (see Figure 5.5). And with calibration functions determined for each channel, we can assign a true energy to each event in the calibration and physics data.

5.2.5 Data blinding

At this point, we have an energy associated with each event, but little else. We will need to do a significant amount of additional data processing in order to produce a good collection of events for our analyses. But before we move forward with the analysis sequence, we must blind the data. We cannot be completely unbiased while analyzing data, and we do not want to preferentially select and process data to conform with our expectations for the result of the experiment. In this case, blinding means processing and analyzing the data without the knowledge of how it’s affecting the data at the Q -value for $0\nu\beta\beta$ decay, $Q_{\beta\beta}$, in the so-called “region of interest” (ROI).

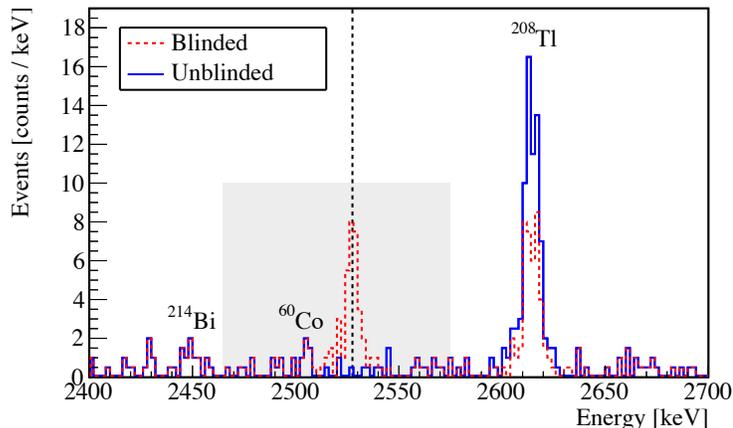


Figure 5.6: Example of the CUORE blinding technique on dataset 3018. A random fraction of events in the ^{208}Tl line at 2615 keV are exchanged with those near the Q -value at 2528 keV, thus creating a fake peak in the region of interest (shaded in gray).

Some experiments choose to ignore the ROI altogether for analysis, but for CUORE we follow a slightly different approach. We instead exchange a small, but unknown, fraction of events between the ROI and the region around the 2615 keV background line. This has the effect of producing a false peak in the ROI, which we can analyze as we are observing $0\nu\beta\beta$ decay. With this method, we can refine all of our analysis tools on a continuous energy spectrum, including how we set a limit on or calculate the value of the $0\nu\beta\beta$ decay half-life in the event that we do make a discovery.

The blinding process is illustrated in [Figure 5.6](#). The procedure is reversible, so we can unblind the data before setting our limit without reprocessing the data from scratch.

5.3 Selecting data

Every triggered event is not equally likely to be a signal event. Because of this, we use a variety of data selection criteria to improve our signal-to-noise ratio while

removing as small a percentage of good signal events as possible. The exact selection criteria will vary depending on the desired analysis, but are generally quite similar. We discuss here the selection criteria used for the $0\nu\beta\beta$ decay analysis.

5.3.1 Removing problematic data

As with any complex detector, many things can result in us collecting poor-quality data. It is crucial to detect when this is occurring so that we can remove this data and not spoil our results. In addition, a search for $0\nu\beta\beta$ decay requires a knowledge of the live time of the detector, and we do not want to include periods of detector malfunction in our evaluation of the live time. Some of the reasons for data removal are:

- Noise. Occasionally, certain channels become very noisy. This can be caused by something as simple as someone opening the Faraday cage door or walking near the cryostat, but can also occur without explanation. Periods of noise can spoil the detector resolution if they are not removed from the dataset.
- Instability. Rarely, there are jumps in the detector baseline that are not correlated with temperature. These can be caused by electronics problems or can occur without any obvious cause. These jumps can limit our ability to perform thermal gain stabilization if they are not removed.
- Saturation. This occurs when the detector baseline reaches the upper or lower limit of the ADC. Often this occurs when a cryogenics problem or an external factor, like an earthquake, causes a significant and sudden change in temperature. We cannot collect any useful data when the detector baseline has saturated the ADC.
- Calibration problems. When a channel cannot be accurately calibrated, we cannot use it for physics analysis. This can occur if the calibration runs at the

beginning and end of a dataset are incompatible, which is often due to failed thermal gain stabilization. While this occurred occasionally in CUORE-0, it has not yet been observed in CUORE.

Overall, data cut for these reasons results in a negligible loss of exposure of approximately 1%.

5.3.2 Pulse shape analysis

Even at times when the detector is otherwise behaving, every triggered pulse is not necessarily good for analysis. Noise spikes or energy deposition directly into the thermistors can cause errant triggers. A small period of detector noise can distort the shape of a pulse. And pile-up, where a new pulse begins before the previous pulse has ended, can result in significantly mistaken energy measurements. To combat these problems, we compute a set of shape parameters for each detector pulse. To evaluate each parameter, we begin with the filtered pulse, as described in section [Section 5.2.2](#). We then set limits on these pulse shape parameters and remove pulses that do not meet the criteria.

The first two parameters of the pulse are the rise time and decay time. The rise time is defined as the time taken by a pulse to rise from 10% to 90% of its maximum amplitude, and for good signal pulses is approximately 100 ms, depending on the particular bolometer. The decay time is the the time taken by a pulse to fall from 90% to 30% of its maximum amplitude, and is approximately 400 ms. These time parameters are designed to eliminate pulses that do rise or decay as one would expect from an energy deposition in the bolometer.

The next parameter is the baseline slope. As was done in the preprocessing, we fit a line to the first 2.25 seconds of each waveform, from 3.00 s to 0.75 s before the trigger. Events with abnormal baseline slope generally occur when the detector is still recovering its baseline following a previous pulse, thus resulting in a mistaken

energy reading.

Next we look at the peak delay, or the time between the start of the window, 3 second before the trigger, and the maximum of the pulse. This is fairly consistent for standard energy depositions at approximately 3.1 s.

Finally, we have two parameters that quantify the shape of the rising edge of the pulse and the shape of the decaying edge. They are χ^2 tests that evaluate how similar the filtered pulse is to the template (average) pulse on the left and right side of the peak, respectively.

The six pulse shape parameters allow us to remove many bad events, but all vary slightly between different channels and datasets, and generally vary significantly with energy. In order to avoid setting channel-, dataset-, and energy-dependent parameter limits, we instead create normalized versions of the parameters that are independent of channel, dataset, and energy. To do this, we take a pulse shape parameter from one channel in one dataset, divide it into energy bins, and then evaluate the median and median absolute deviation (MAD) for that parameter in each energy bin (see [Figure 5.7](#)). We then use empirically-determined functions (see [Table 5.2](#)) to create a mapping between the parameter median vs. energy and the parameter MAD vs. energy. Finally, we normalize the parameter for an event with energy E as

$$\text{normalized value} = \frac{\text{raw value} - \text{median at energy } E}{\text{MAD at energy } E}. \quad (5.11)$$

In order to reject outlier pulses, we need to establish the normal variation of normalized pulse shape parameters in good signal events. For this, we look at events in the strongest lines in the background spectrum, those from ^{40}K (1461 keV) and ^{60}Co (1173 and 1332 keV). This undoubtedly includes some background events, but these regions are dominated by good signal events. From these events, we obtain a mean μ_i for each normalized pulse shape parameter and a covariance matrix S

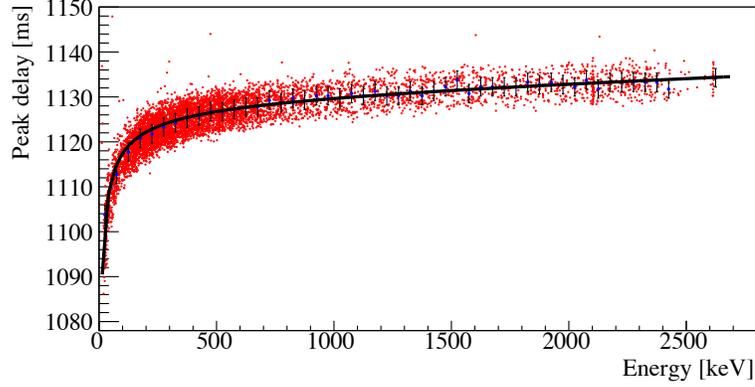


Figure 5.7: Example of a pulse shape parameter (in this case, the peak delay) dependence on energy for a single channel in CUORE-0. The median of the parameter is evaluated in each energy bin, and a fit to these points determines the mapping between the median and the event energy. Figure from S. Pozzi.

Parameter	Median fit	MAD fit
Baseline slope	a_0	a_0
Rise time	$a_0 + a_1x + a_2x^{-1/2}$	$a_0 + a_1x^{-a_2} + a_3x$
Decay time	$a_0 + a_1x + a_2x^{-a_3}$	$a_0 + a_1x^{-a_2} + a_3x$
Peak delay	$a_0 + a_1x^{-1/2} + a_2x$	$a_0 + a_1 \exp(a_2(a_3 + x))$
TVL	$a_0 + a_1x^{-a_2} + a_3x$	$a_0 + a_1x^{a_2} + a_3x^2 $
TVR	$a_0 + a_1x^{a_2} + a_3(x - a_4)^2$	$a_0 + a_1x^{a_2} + a_3x^2 $

Table 5.2: The empirically-determined functions used to evaluate the dependence of the median and median absolute deviation (MAD) of the pulse shape parameters on energy. Note that the baseline slope, which is evaluated from the detector baseline prior to the pulse, does not have any energy dependence.

relating all of the parameters.

With the means and covariance matrix, we can determine how much of an outlier any given event is from the standard distribution of good signal events. For this, we use the Mahalanobis distance, which is essentially determining how much of an outlier the parameter set is by computing the number of standard deviations it is from the mean, in multiple dimensions [145]. For an observation \vec{x} , the Mahalanobis distance is defined as

$$D_M(\vec{x}) = \sqrt{(\vec{x} - \vec{\mu})^T S^{-1} (\vec{x} - \vec{\mu})}, \quad (5.12)$$

where $\vec{\mu}$ is the vector of parameter means and S is the covariance matrix. By setting an upper limit on this distance, we can effectively remove outlier events from our spectrum. We return to a discussion of choosing this upper limit in [Section 6.1](#).

We have not found a way to distinguish the pulses that result from α , β , and γ interactions in the detector. As such, the pulse shape analysis is a useful tool for rejecting spurious pulses, but cannot reject true backgrounds in the $0\nu\beta\beta$ decay ROI.

5.3.3 Geometric coincidence analysis

The CUORE detector is comprised of 988 individual detectors. While this is a downside in terms of detector complexity, it is a significant boon to our analysis. Most importantly, simulations tells us that $0\nu\beta\beta$ decay events would be completely contained in a single crystal 88% of the time, while external muons almost always deposit energy in multiple crystals. Other types of events, particularly γ interactions, sometimes show up in one crystal and sometimes in multiple; properly simulating this helps us to ensure that we understand our background spectrum well.

To perform a coincidence analysis, we must understand the time scales of the CUORE detector. Compared to the time it takes for a particle to travel from one crystal to another (nanoseconds), the response time of the bolometers and our electronics is very slow (milliseconds). Thus, any particles that pass through multiple crystals should appear in our detectors simultaneously, and likewise for any physical event that causes particles to deposit energy in more than one crystal.

In practice, even though events should appear to be simultaneous in multiple channels, their measured times (determined by the time of the peak of the waveform) can differ quite significantly (on the order of 10–50 ms). This is due to differences in the response functions and average pulse shapes of different bolometers. We call this difference in response time between different channels “jitter,” and it is remarkably stable across different energies and over time.

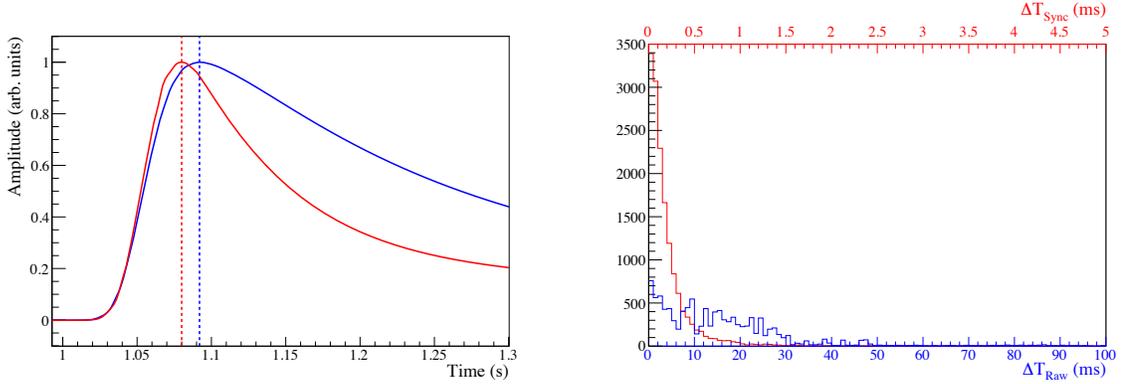


Figure 5.8: Left: The average pulses from two different channels in a CUORE-0 data, demonstrating how different response functions can lead to differences in measured peak timing. The jitter is evaluated as the measured time difference between the peaks of the waveforms for truly simultaneous energy depositions in two different bolometers. Right: The distribution of the time delay between pairs of coincident events before (blue) and after (red) correcting for jitter. Figure from J. Ouellet.

To measure the jitter, we look at events that we believe to be truly coincident; in particular, we look in the calibration data for events occurring in two crystals in a short time window with a summed energy of approximately 2615 keV. These events are highly likely to be caused by a 2615 keV γ ray from ^{208}Tl decay Compton scattering in multiple crystals, and should therefore be perfectly simultaneous up to the timing resolution of the detector. By looking at the difference in measured time between these simultaneous events, we can deduce the jitter of each channel and correct for it when determining the timing of each event. Using this jitter-adjusted time, we can define coincident events as those happening within a 10 ms window, instead of those within a 50 ms window (see Figure 5.8). This helps to greatly reduce accidental coincidences, which hurt our detector live time when we select for only non-coincident events.

Once the jitter has been corrected, we assign each event a “multiplicity.” A multiplicity-1 event happens alone in the detector, a multiplicity-2 event happens at the same time as one other event, and so on. Because each detector has a different

energy threshold, in order to standardize our evaluation of the multiplicity, we only consider events with energy greater than 150 keV for this calculation.

5.3.4 Combining energy estimators

Throughout the data processing, there are some decisions on methodology that need to be made that affect our ultimate energy spectrum, but that we choose to postpone until the processing is complete. We do this in order to evaluate how these decisions affect our detector performance after all processing is complete. In addition, we may want or need to make different decisions for different channels. To accomplish this, we branch our analysis and end up with multiple different evaluations of each pulse’s energy, which we refer to as “energy estimators.” For this analysis, there is one such decision: whether to perform thermal gain stabilization using the heater pulses or using the 2615 keV line from calibration. In CUORE-0, we also followed this procedure for the decision on whether to perform noise decorrelation, which we anticipate using on the full CUORE detector array in future analyses. This is also the standard procedure by which we evaluate the effect of new analysis techniques on our data.

The decision on how to perform thermal gain stabilization is necessarily made on a channel-by-channel basis, since there are several channels for which heater stabilization is impossible due to dead or malfunctioning heaters. For those channels that have both options available, we calculate a ratio R that incorporates the ratio of the ultimate sensitivities (see [Equation 2.29](#)), namely

$$R^{1/4} = \frac{\sqrt{t_c/\sigma_c}}{\sqrt{t_h/\sigma_h}},$$

where σ_c and σ_h are the resolutions at the 2615 keV line using the calibration and heater stabilization, respectively, and t_c and t_h are the corresponding live times using

the same. We define this as $R^{1/4}$ so that we can work with the easily computed variance σ^2 instead of $\sqrt{\sigma}$.

With some exceptions, stabilization with the heater and with the 2615 keV calibration line function similarly, though the heater stabilization performs slightly better on average. We therefore choose the heater stabilization as our baseline approach. After calculating R and δR (the uncertainty on R) for each channel in each dataset, we use the stabilization with the 2615 keV calibration line for those channels with $R > 1$ and $(R - 1)/\delta R > 1$. That is, we require an improvement in the expected sensitivity and require that the improvement is statistically significant at the $\sim 90\%$ confidence level. In practice, we find that the vast majority of channels are clustered at or below $R = 1$. For the first dataset, we had some difficulties with the constant-energy pulsers, and 40 channels were switched to using the calibration stabilization with this method. For the second dataset, only 2 channels were switched.

After the final selection of the energy estimators, every event in the data has an energy and multiplicity, and we are ready to begin the physics analysis.

Chapter 6

Neutrinoless double-beta decay analysis

“Science is competitive, aggressive, demanding. It is also imaginative, inspiring, uplifting. You can do it, too.”

— Vera Rubin

The CUORE detector towers were installed in the cryostat in the summer of 2016, and detector operations began in January 2017. Following a period of detector commissioning, calibration, and optimization, we began taking physics data in May 2017. CUORE is scheduled to run for 6 years, with 5 years of live time.

The analysis presented below is based on the first neutrinoless double-beta decay search data from CUORE, acquired in two month-long datasets between May 2017 and September 2017, which we will call Dataset 1 and Dataset 2. Between the two datasets, an optimization campaign was performed to reduce detector noise, so some of the parameters in the analysis vary between these two datasets. In this period, we acquired 86.3 kg·yr of TeO₂ exposure, corresponding to a total ¹³⁰Te exposure of 24.0 kg·yr.

The production of the official data, as described in [Chapter 5](#), is a collaboration

effort, with analysis software written and executed by numerous members of the CUORE collaboration, including me. The analysis below, however, was performed by me, building off of the official CUORE data, except where specifically noted. It differs in some respects from the official CUORE $0\nu\beta\beta$ decay analysis, which is the merged product of several independent analyses performed by many collaborators. In general, all of the methods used produce consistent results. I will point out and discuss the places in which my methods differ from the official CUORE analysis in footnotes throughout this chapter.

6.1 Detection efficiency

The first step in our search for $0\nu\beta\beta$ decay is selecting the candidate events from our energy spectrum. In our quest to remove backgrounds from our energy region of interest and produce the cleanest energy spectrum, however, we will inevitably discard real $0\nu\beta\beta$ decay events — if they exist — with some probability. Our goal is to quantify that probability so that we can properly calculate a true event rate from our observed number of events.

This quantification involves two steps: data selection optimization and a calculation of the data selection efficiency. To avoid biasing our evaluation of the selection efficiency, we optimize our data selection criteria (“cuts”) on different data than we use to evaluate the selection efficiency. The optimization and efficiency calculations are described below.

6.1.1 Optimization

Our main tool for rejecting noise and other spurious pulses is our pulse shape analysis. The six pulse shape parameters that we use are the baseline slope, rise time, decay time, delay, rise shape, and decay shape. These parameters vary with pulse energy,

so when we are selecting data, we use normalized values; these normalized values are a measure of the deviation from the mean value of that parameter at the evaluated energy of the pulse. We then evaluate a Mahalanobis distance for each parameter set to quantify how much of an outlier each pulse is (see [Section 5.3.2](#)).

To choose the optimal value for the upper limit on the Mahalanobis distance, we examine the 1461 keV line from ^{40}K , since it is the strongest line in the physics spectrum. To date, we have no ability to distinguish between different particles interacting in our detector, so although this is a line from a γ ray, we use it to model the $0\nu\beta\beta$ decay pulse, which would be from electrons. We first calculate the number of events in the peak, after background subtraction, without applying any cut on the Mahalanobis distance, and then calculate this number after applying the cut. The number of remaining events divided by the number of initial events is the signal selection efficiency¹. We then calculate the number of events before and after applying the cut in a region of the spectrum with no known spectral lines; this is our background selection efficiency. For this, we use the energy region 2700–3100 keV, which is free from any γ lines and has approximately the same background rate as the $0\nu\beta\beta$ decay region of interest.

In the ideal situation, we would want to choose a cut that maximizes our signal selection efficiency ϵ_s while minimizing our background selection efficiency ϵ_b . Since these cannot be accomplished simultaneously, we must find the best tradeoff between the two. We thus calculate the signal and background selection efficiencies with various upper limits on the Mahalanobis distance and search for the configuration that maximizes $\epsilon_s/\sqrt{\epsilon_b}$; i.e., we maximize the number of signal events divided by the background fluctuation, assuming a Gaussian-distributed background, as this is proportional to the significance of our result. The results of this search are shown in

1. Note that the signal selection efficiency here is that used to optimize the pulse-shape cut; the efficiency that goes in to the final analysis is calculated on different data (the 2615 keV line) after the cut is finalized, to avoid biasing the result.

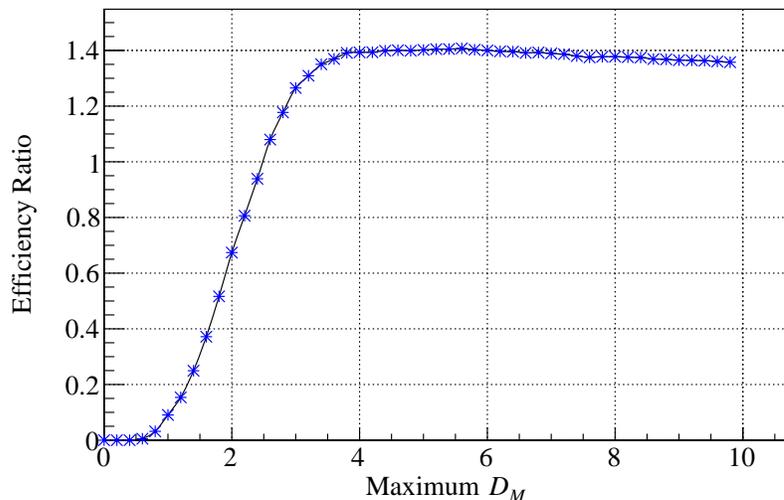


Figure 6.1: Signal-to-background figure of merit ($\epsilon_s/\sqrt{\epsilon_b}$) evaluated on the 1461 keV line from ^{40}K with various limits on the Mahalanobis distance of the pulse shape parameters. The choice $D_M < 5$ may be refined as more data is acquired in subsequent datasets, but in the end, the sensitivity is not highly dependent on this value for all maximum values of D_M greater than 4.

Figure 6.1. In the end, we place a cut on the Mahalanobis distance of $D_M < 5$.

6.1.2 Efficiency evaluation

Now that our pulse shape cuts are finalized, we can estimate our selection efficiency. That is, we estimate the probability of us recording a $0\nu\beta\beta$ decay event given that one has occurred. This depends on a variety of factors, including whether the electrons from the decay are fully contained in one bolometer, whether the resulting pulse is triggered and accurately reconstructed, and whether the pulse survives our analysis cuts. A summary of the three analysis cuts that we use for the $0\nu\beta\beta$ decay analysis (pile-up, pulse shape, and coincidence) is in [Table 6.1](#).

To evaluate the selection efficiency for the pile-up and pulse shape cuts, we examine the 2615 keV line from ^{208}Tl in the physics spectrum. We use this line because it is close in energy to the $0\nu\beta\beta$ decay Q -value at 2528 keV and fairly prominent in

Name	Description
<u>Pile-up:</u>	
Single pulse	There is only one derivative sign change and one trigger in the pulse window.
<u>Pulse shape:</u>	
Baseline slope	The baseline before the pulse is sufficiently flat.
Rise time	The rise time of the pulse is as expected for its energy.
Decay time	The decay time of the pulse is as expected for its energy.
Delay	The delay between the trigger and pulse peak is as expected for its energy.
Rise shape	The shape of the pulse during its rise is as expected for its energy.
Decay shape	The shape of the pulse during its decay is as expected for its energy.
<u>Coincidence:</u>	
Anti-coincidence	There are no other simultaneous energy depositions in nearby crystals.

Table 6.1: The cuts used in the final $0\nu\beta\beta$ decay analysis. The expected values of the pulse shape parameters are determined by examining the 1461-keV γ line from ^{40}K .

the spectrum². As we do when optimizing the cuts, we first calculate the number of events in the peak, after background subtraction, without applying any cuts, and then calculate this number after applying the pile-up and then pulse shape cuts. The number of remaining events divided by the number of initial events is the selection efficiency, and we calculate a Clopper–Pearson interval [146] at 68% coverage as the uncertainty on the efficiency. In this analysis, we evaluate the selection efficiencies independently for the two datasets, since the detector conditions changed slightly between the datasets. We find the pile-up selection efficiency to be $(97.6 \pm 1.1)\%$ in Dataset 1 and $(96.7 \pm 1.0)\%$ in Dataset 2, and our pulse shape selection efficiency to be $(93.9 \pm 1.6)\%$ in Dataset 1 and $(96.8 \pm 1.0)\%$ in Dataset 2.

As we discussed in Section 5.3.3, $0\nu\beta\beta$ decay events are largely contained in a single crystal; that is, the entire energy of the electrons emitted in the decay is absorbed by the crystal inside which the decay occurred. In contrast, external γ rays often Compton scatter in multiple crystals in a time scale much faster than our detector timing, and external muons almost always deposit energy in multiple crystals

2. In the published CUORE analysis, the pile-up efficiency was evaluated on the pulser line. At the cost of lower statistics and thus a greater uncertainty, I’ve evaluated this efficiency on the 2615 keV line in order to evaluate it using events in the physics spectrum, as the posited $0\nu\beta\beta$ decay peak would be. These two methods agree within their uncertainty.

Source	Selection efficiency (%)	
	Dataset 1	Dataset 2
Pile-up cut	97.6 ± 1.1	96.7 ± 1.0
Pulse shape cut	93.9 ± 1.6	96.8 ± 1.0
Anti-coincidence	99.8 ± 0.1	$100. \pm 0.1$
Trigger and reconstruction	99.0 ± 0.1	99.0 ± 0.1
Total excluding containment	90.5 ± 1.9	92.7 ± 1.4
Containment	88.3 ± 0.1	88.3 ± 0.1

Table 6.2: The selection efficiencies for the cuts used in the $0\nu\beta\beta$ decay analysis.

simultaneously as well. For these reasons, we impose a multiplicity cut and require that each event is coincident with no other energy deposits in nearby crystals. To evaluate the efficiency of this anti-coincidence cut with respect to accidental coincidences between detectors, we examine the 1461 keV line from ^{40}K decay instead of the 2615 keV line. Unlike the 2615 keV γ ray, which is always emitted along with one or more other γ rays as the ^{208}Pb nucleus deexcites, the 1461 keV γ ray is always emitted alone; that is, ^{40}K electron capture always populates the same excited state of the ^{40}Ar nucleus, which always decays directly to the ground state. This makes the 1461 keV line ideal for evaluating whether a single event, like $0\nu\beta\beta$ decay, would be in random, accidental coincidence with a background event in another crystal. We evaluate the anti-coincidence selection efficiency to be $(99.8 \pm 0.1)\%$ in Dataset 1 and $(100. \pm 0.1)\%$ in Dataset 2. This very high efficiency is in line with our expectations, due to the extremely low event rate in the detectors.

All of the selection efficiencies discussed above are based on triggered events, but we need to consider two other effects as well. The first is the fraction of real events that actually generate a trigger. We call this our trigger efficiency, and we evaluate it with our periodic heater pulses.³ Specifically, we look at all of the heater pulses

³. The trigger efficiency is energy-dependent, as it relies on picking out pulses, whose heights depend on energy, from background fluctuations. However, this is a significant effect only near the energy threshold, which is orders of magnitude below the $Q_{\beta\beta}$ and the energy of the heater pulses.

flagged as such by the data acquisition system, and we count how many also passed our signal trigger. The second is the fraction of events that are reconstructed with the “correct” energy. This reconstruction efficiency is evaluated by fitting the peak in the energy spectrum of heater pulses in each channel with a Gaussian distribution and counting the fraction of events reconstructed within 3σ of the mean⁴. Together, the trigger and reconstruction efficiency is independently evaluated to be $(99.0 \pm 0.1)\%$ in both datasets⁵.

There is one final component of the selection efficiency that cannot be calculated from our data. Because we look only at the energy recorded in single crystals, any $0\nu\beta\beta$ decay events in which the two electrons were not fully contained in the crystals would not appear in the energy spectrum at $Q_{\beta\beta}$. We evaluate this probability with Monte Carlo simulations and obtain a $0\nu\beta\beta$ decay full containment probability of $(88.3 \pm 0.1)\%$ ⁶.

Our overall data selection efficiency is the product of each of these individual efficiencies; we evaluate this to be $(79.9 \pm 1.9)\%$ in Dataset 1 and $(81.9 \pm 1.4)\%$ in Dataset 2. In other words, we expect that $\sim 80\%$ of real $0\nu\beta\beta$ decay events would appear in our final spectrum in a distribution centered on $Q_{\beta\beta}$.

6.2 Detector response and resolution

Ultimately, we are seeking a small peak in the energy spectrum at $Q_{\beta\beta}$. To look for this peak, we must know three things: the expected shape of the peak, the expected width of the peak, and the expected location of the peak. For these, we turn to the γ lines in our physics spectrum.

4. The few channels without a functioning heater are excluded from the trigger and reconstruction efficiency evaluation, and are thus assumed to have average trigger and reconstruction efficiency.

5. The trigger and reconstruction efficiency was computed by A. Caminata and I. Nutini.

6. Monte Carlo simulations were performed by S. Capelli.

6.2.1 2615 keV calibration line

As we have seen, the 2615 keV line from ^{208}Tl decay is very useful for the $0\nu\beta\beta$ decay analysis since it lies so close in energy to $Q_{\beta\beta} = 2528$ keV. Here, we use it to establish our energy spectrum line shape so we know what to expect from a $0\nu\beta\beta$ decay signal. We use the 2615 keV line in the calibration data, since there are significantly more events in the peak there than in the physics data.

We perform a simultaneous unbinned extended maximum likelihood (UEML) fit to the 2615 keV calibration line over all channels in all datasets. We approximate the line shape in each channel with three Gaussian functions: one central peak and two small peaks (one on each side), which have the effect of widening the tails of the central peak. Each channel–dataset pair has its own peak width (i.e., its own resolution) and its own peak mean (to account for any residual calibration biases), but all channel–dataset pairs are constrained to have the same intensity ratio between the three Gaussians and the same distance between the center peak and the two side peaks⁷. The other components of this fit are:

- A flat background, resulting from uncaught pile-up and a variety of background particles, including degraded α particles resulting from ^{232}Th , ^{238}U , and ^{210}Po decay. Each channel has its own floating background rate.
- A complementary error function,

$$\frac{1}{\sqrt{\pi}} \int_{\frac{x-\mu}{\sqrt{2}\sigma}}^{\infty} e^{-t^2} dt, \quad (6.1)$$

where μ and σ are those from the central Gaussian function, representing the

7. This is a significant differences between the analysis presented here and the official CUORE analysis. The official analysis allows each channel to have its own ratios between the intensities and energies of the three peaks, but here they are shared by all channels. The official analysis thus fits this calibration line better, but at the expense of introducing over 7,000 additional parameters in the fit. The fits to the lines in the physics spectrum are functionally equivalent without these additional parameters.

multiscatter Compton background from 2615 keV γ rays. The relative intensity of this complementary error function to the intensity of the main peak is common to all channels.

- A Gaussian function located 28 keV below the main peak at 2615 keV, representing energy depositions in which an X-ray from Te escaped⁸. The relative intensity of this line to the intensity of the main peak is common to all channels.

Overall, the fit function for this peak is

$$f(E) = a \left[\mathcal{N}(\mu, \sigma; E) + \kappa_L \mathcal{N}(\epsilon_L \mu, \sigma; E) + \kappa_R \mathcal{N}(\epsilon_R \mu, \sigma; E) + \frac{\kappa_C}{\sqrt{2\pi} \sigma} \int_E^\infty e^{-(t-\mu)^2/(2\sigma^2)} dt + \kappa_X \mathcal{N}(\epsilon_X \mu, \sigma; E) \right] + b, \quad (6.2)$$

where a and b (scaling parameters) and σ and μ (line shape parameters) are floating and split by channel and dataset; κ_L , κ_R , κ_C , and κ_X (scaling parameters), and ϵ_L and ϵ_R (energy scaling factors for the lower and upper Gaussian functions) are floating and shared by all channels in all datasets⁹; ϵ_X is an energy scaling factor for the Te X-rays; and \mathcal{N} is a Gaussian function. The results of this fit are shown in [Figure 6.2](#). In the end, for the shared parameters we obtain the best-fit values $\kappa_L = 0.040$, $\kappa_R = 0.001$, $\kappa_C = 0.003$, $\kappa_X = 0.003$, $\epsilon_L = 0.996$, and $\epsilon_R = 1.005$.

The fit to the 2615 keV line gives us a resolution σ_i and mean μ_i for each channel–dataset pair i , as well as the shared parameters described above. We save these parameters and use them to fit other lines in the physics data, including the potential $0\nu\beta\beta$ decay peak. But first, we can look at the distribution of σ_i to get an idea of what our detector performance is. As is customary, we actually look at $2\sqrt{2\ln 2} \sigma_i \approx 2.355\sigma_i$, which is the FWHM of a Gaussian distribution with standard deviation σ_i

8. This is really 6 individual lines 27–30 keV below the main peak, but because they are close in energy relative to the detector resolution and altogether very weak, we model them as a single line.

9. For computational reasons, these parameters are computed on a tower-by-tower basis and then averaged together, with the average being assigned to each channel in a second iteration of the fit.

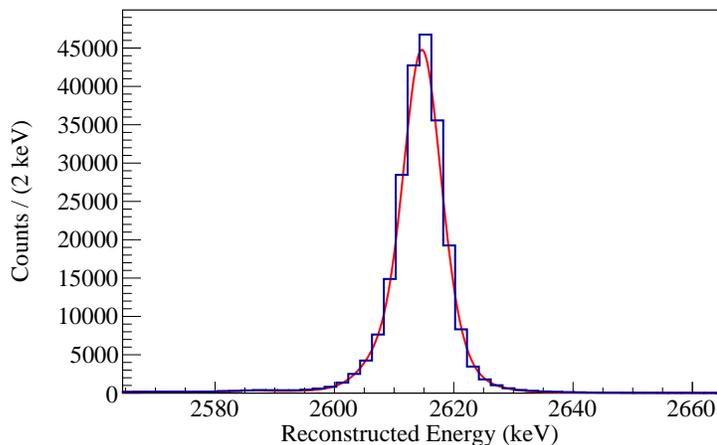


Figure 6.2: Overall fit to the 2615 keV line, summed over all channels and datasets.

(see Figure 6.3). Because our sensitivity is approximately proportional to $\sqrt{1/\sigma_i}$, we compute an effective FWHM for our detector, at the 2615 keV line in calibration, as the physics-exposure-weighted harmonic mean of the individual channel–dataset resolutions; this gives us $\Delta E_{\text{eff}}^{\text{cal}} = 8.3$ keV (FWHM)¹⁰, or 9.6 keV in Dataset 1 and 7.5 keV in Dataset 2.

6.2.2 Full physics spectrum analysis

We use the 2615 keV line from calibration to give us the line shape in each channel and the relative resolution of the channels, but it cannot tell us the resolutions in our physics runs. For this, we perform a UEML fit to the 2615 keV line in the physics data. Each resolution σ_i is taken from the fit to the 2615 keV calibration line and is scaled by a factor η , shared by all channel–dataset pairs, which we call the resolution scaling parameter. The overall mean μ of the fit is floating and shared by all channels, though it is expressed as an energy scaling factor multiplied by μ_i from the 2615 keV

¹⁰. The published CUORE analysis found an effective resolution of 8.0 keV, as opposed to 8.3 keV. With floating channel-dependent parameters, the additional Gaussian peaks on the left and right sides of the main peak had a high intensity in some channels (up to $\sim 25\%$ of that of the main peak), so the reported mean of the central Gaussian was lower. In the end, this does not affect the results, just the quoted resolution, since the full line shape is used in the analysis regardless.

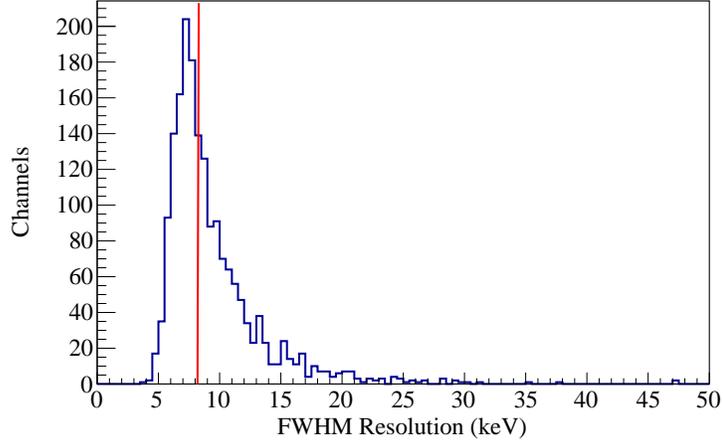


Figure 6.3: Distribution of resolutions at the 2615 keV peak in calibration data for all channel–dataset pairs. The physics-exposure-weighted harmonic mean resolution of 8.3 keV is marked in red.

calibration line fit, so that any residual miscalibration is accounted for.

Because we have many fewer events in the 2615 keV peak in physics data (generally ≤ 1 event per channel–dataset pair), and almost no background events, we can not have the signal and background rates floating for each channel–dataset pair. At the same time, we expect channels closer to the large copper shielding to have more background events than channels in the center of the detector array. We therefore split the channels into two layers, each of which has a floating signal rate (in counts/(kg·yr)) and background rate (in counts/(keV·kg·yr)). The outer layer consists of the outer ring of 12 towers as well as the top and bottom 2 floors from the inner 7 towers, for a total of 736 channels; the inner layer consists of the other 262 channels.

In all, the fit function for the 2615 keV peak in the physics spectrum is

$$\begin{aligned}
 f(E) = \xi \epsilon s \left[\mathcal{N}(\chi\mu, \eta\sigma; E) + \kappa_L \mathcal{N}(\epsilon_L \chi\mu, \eta\sigma; E) + \kappa_R \mathcal{N}(\epsilon_R \chi\mu, \eta\sigma; E) \right. \\
 \left. + \frac{\kappa_C}{\sqrt{2\pi} \eta\sigma} \int_E^\infty e^{-(t-\chi\mu)^2/(2\eta^2\sigma^2)} dt \right] + \epsilon b \Delta E,
 \end{aligned} \tag{6.3}$$

where σ and μ (line shape parameters) are split by channel and dataset and fixed

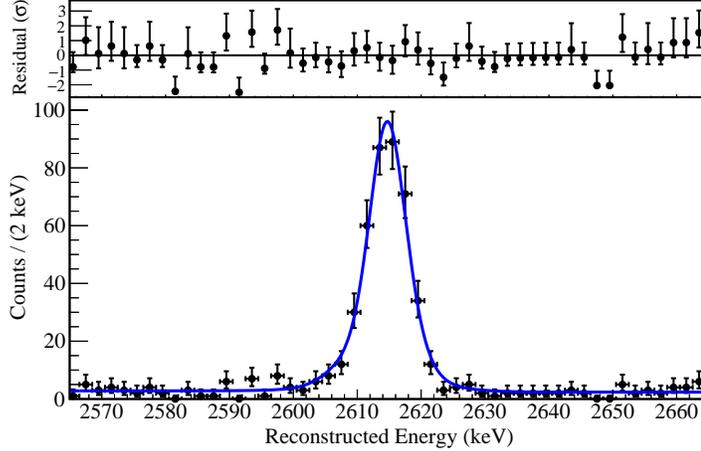


Figure 6.4: Overall fit to the 2615 keV background line in all channels and datasets.

to the calibration line fit results; ε (exposure) is split by channel and dataset and fixed by the data; χ (the peak mean scaling factor) and κ_C (Compton scatter scaling factor) are floating and shared by all channels in all datasets; η (resolution scale factor) is floating and split by dataset; ξ (efficiency) is split by dataset and fixed from the efficiency calculated above (see Table 6.2); s (signal rate) and b (background rate) are split by layer and floating; κ_L and κ_R are fixed from the calibration line fit results; and ΔE is the width of the fit window (in order to properly scale the background rate b). No X-ray escape peak is included, as was done for the calibration line fit, because it is too weak to appear in the spectrum.

The result of this fit to the 2615 keV line in the physics spectrum is shown in Figure 6.4. We find that the resolution of the 2615 keV background line is $(87.5 \pm 6.7)\%$ of the 2615 keV calibration line resolution in Dataset 1 and $(95.6 \pm 5.7)\%$ in Dataset 2. We also find that the fit mean is (2614.53 ± 0.18) keV, which is consistent with the nominal value of (2614.511 ± 0.010) keV [142].

We then repeat this fit on the other prominent γ background lines in the physics spectrum (see Figure 6.5). In particular, we look at the 1461 keV line from ^{40}K , the 1173 keV and 1332 keV lines from ^{60}Co , the 911 keV line from ^{228}Ac , and the 835 keV

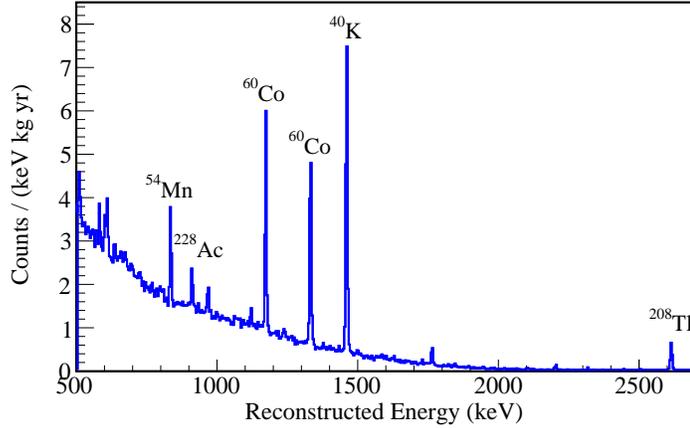


Figure 6.5: Physics spectrum, including all cuts used in the $0\nu\beta\beta$ decay analysis. Visible lines in the spectrum used to determine the resolution and energy bias are labelled.

line from ^{54}Mn . The results from these fits are shown in [Table 6.3](#). To estimate the energy resolution and energy reconstruction bias at $Q_{\beta\beta}$, we extrapolate from this data with a second-order polynomial (see [Figure 6.6](#)). We estimate the resolution at $Q_{\beta\beta}$ to be (8.2 ± 0.4) keV in Dataset 1 and (7.1 ± 0.7) keV in Dataset 2. We see no evidence of any statistically significant bias, with an extrapolation showing a bias of (0.01 ± 0.17) keV at $Q_{\beta\beta}$.

6.3 Neutrinoless double-beta decay fit

We now have all the ingredients necessary to look for any evidence of a signal at $Q_{\beta\beta}$. We choose the region from 2465 keV to 2575 keV as the region of interest (ROI), starting above the 2448 keV line from ^{214}Bi and below any influence of the tail of the 2615 keV line from ^{208}Tl . Inside this ROI, we have only a few components: a flat background, a peak at approximately 2506 keV representing the coincident absorption of two ^{60}Co γ rays, and a posited peak centered at $Q_{\beta\beta}$. Both peaks use the full line-shape developed in the fit to the 2615 keV calibration line.

Energy [keV]	Signal rate [counts/(kg·yr)]	Background rate [counts/(keV·kg·yr)]	FWHM resolution [keV]	Bias [keV]
2614.511	2.04 ± 0.33 (inner)	0.011 ± 0.002 (inner)	8.40 ± 0.64 (DS 1)	0.02 ± 0.18
	6.07 ± 0.34 (outer)	0.015 ± 0.002 (outer)	7.18 ± 0.43 (DS 2)	
1460.822	44.0 ± 1.5 (inner)	0.418 ± 0.073 (inner)	5.54 ± 0.13 (DS 1)	−0.02 ± 0.07
	64.1 ± 1.2 (outer)	0.453 ± 0.099 (outer)	5.58 ± 0.11 (DS 2)	
1332.492	14.0 ± 1.0 (inner)	0.561 ± 0.018 (inner)	5.23 ± 0.15 (DS 1)	−0.12 ± 0.05
	47.2 ± 1.0 (outer)	0.581 ± 0.022 (outer)	5.68 ± 0.16 (DS 2)	
1173.228	13.6 ± 1.0 (inner)	0.790 ± 0.022 (inner)	4.90 ± 0.18 (DS 1)	−0.08 ± 0.05
	44.4 ± 1.0 (outer)	0.981 ± 0.037 (outer)	4.88 ± 0.14 (DS 2)	
911.204	3.64 ± 0.78 (inner)	1.18 ± 0.03 (inner)	3.74 ± 0.50 (DS 1)	0.06 ± 0.16
	7.88 ± 0.67 (outer)	1.56 ± 0.03 (outer)	4.45 ± 0.51 (DS 2)	
834.848	4.87 ± 0.87 (inner)	1.29 ± 0.03 (inner)	4.12 ± 0.27 (DS 1)	0.12 ± 0.11
	18.3 ± 0.8 (outer)	1.77 ± 0.05 (outer)	4.68 ± 0.23 (DS 2)	

Table 6.3: Fit results for the prominent γ background lines. The effective resolution is the exposure-weighted harmonic mean FWHM of the peak in all of the channels. The bias is the fit result minus the nominal energy of the peak.

The function used to fit the ROI is thus given by

$$\begin{aligned}
f(E) = & \xi_C \xi_\gamma \varepsilon R_{\beta\beta} [\mathcal{N}(\chi_{\beta\beta}\mu, \eta\sigma; E) + \kappa_L \mathcal{N}(\epsilon_L \chi_{\beta\beta}\mu, \eta\sigma; E) + \kappa_R \mathcal{N}(\epsilon_R \chi_{\beta\beta}\mu, \eta\sigma; E)] \\
& + \xi_\gamma \varepsilon R_{\text{Co}} e^{-t/\tau_{\text{Co}}} [\mathcal{N}(\chi_{\text{Co}}\mu, \eta\sigma; E) + \kappa_L \mathcal{N}(\epsilon_L \chi_{\text{Co}}\mu, \eta\sigma; E) + \kappa_R \mathcal{N}(\epsilon_R \chi_{\text{Co}}\mu, \eta\sigma; E)] \quad (6.4) \\
& + \varepsilon b \Delta E.
\end{aligned}$$

There are only five free parameters in this fit: $R_{\beta\beta}$, the $0\nu\beta\beta$ decay signal rate (in counts/(kg·yr)); R_{Co} , the ^{60}Co decay signal rate in the first dataset (also in counts/(kg·yr)); b , the dataset-dependent background rate (in counts/(keV·kg·yr)), counting as two free parameters; and χ_{Co} , the energy ratio between the ^{60}Co sum line (nominally at 2505.7 keV) and the 2615 keV line. The energy of the ^{60}Co sum line is allowed to float due to a non-unity quenching factor for coincident γ rays on the same crystal¹¹. Conversely, the energy ratio $\chi_{\beta\beta}$ is fixed to (2527.515 keV)/(2614.511 keV), because it is not believed to have a different quenching factor from the single- γ lines used to set the energy scale. The other parameters of the fit are all fixed. These are ξ_C , the $0\nu\beta\beta$ decay full containment probability; ξ_γ , the dataset-dependent selection

11. This effect was seen in Cuoricino and in CUORE-0, and was verified with a dedicated ^{60}Co calibration with CUORE-0 [126].

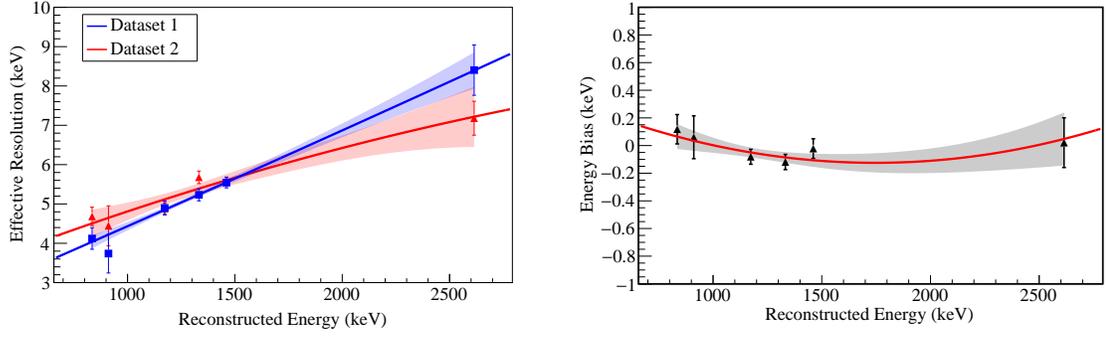


Figure 6.6: Energy resolution and energy bias of the 6 prominent lines in the physics spectrum, along with a quadratic fit and uncertainty, used to evaluate the resolution at bias at $Q_{\beta\beta}$.

efficiency evaluated on γ lines; ε , the channel- and dataset-dependent exposure; μ and σ , the channel- and dataset-dependent Gaussian line shape parameters from the 2615 keV calibration line fit; η , the dataset-dependent resolution scaling factor determined with Equation 6.3; t , the start time of the dataset relative to the start time of the first dataset; τ_{Co} , the known half-life of ^{60}Co ; κ_L , κ_R , ϵ_L , and ϵ_R , the secondary and tertiary Gaussian parameters determined from the 2615 keV calibration line, and ΔE , the width of the the ROI, used to normalize the background rate b .

We perform an unbinned extended maximum likelihood fit to the data to find the best-fit values for the five floating parameters. The results of this fit are shown in Figure 6.7. We observe a downward fluctuation in the data around $Q_{\beta\beta}$ and obtain the best-fit values

$$\hat{R}_{\beta\beta} = (-0.13 \pm 0.04) \text{ counts}/(\text{kg}\cdot\text{yr}) \quad (6.5)$$

$$\hat{R}_{\text{Co}} = (0.23 \pm 0.08) \text{ counts}/(\text{kg}\cdot\text{yr}) \quad (6.6)$$

$$\hat{b}_1 = (0.016 \pm 0.002) \text{ counts}/(\text{keV}\cdot\text{kg}\cdot\text{yr}) \quad (6.7)$$

$$\hat{b}_2 = (0.015 \pm 0.002) \text{ counts}/(\text{keV}\cdot\text{kg}\cdot\text{yr}). \quad (6.8)$$

The best-fit energy for the ^{60}Co sum peak is (2506.3 ± 1.1) keV, compared to the

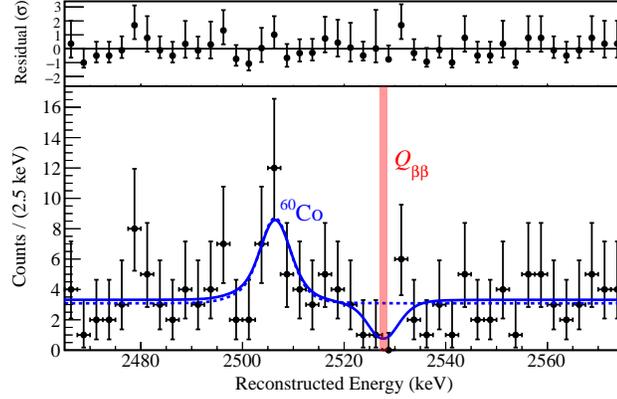


Figure 6.7: The $0\nu\beta\beta$ decay region of interest, along with the best-fit line under the signal hypothesis (solid line) and no-signal hypothesis (dashed line).

nominal value of 2505.7 keV. In the null hypothesis fit, with the signal rate fixed to zero, we obtain a background rate of (0.014 ± 0.002) counts/(keV·kg·yr).

We then convert the signal rate $R_{\beta\beta}$ into a decay rate $\Gamma^{0\nu}$ with

$$\Gamma^{0\nu} = R_{\beta\beta} \times \frac{\mathcal{M}_{\text{TeO}_2}}{aN_A} = R_{\beta\beta} \times \frac{159.6 \text{ g mol}^{-1}}{(0.3417)(6.022 \times 10^{23} \text{ mol}^{-1})}, \quad (6.9)$$

where $\mathcal{M}_{\text{TeO}_2}$ is the molar mass of TeO_2 , a is the ^{130}Te isotopic abundance, and N_A is Avogadro's number. In other words, we compute that there are 1.3×10^{24} atoms of ^{130}Te in 1 kg of TeO_2 , and therefore express the best-fit decay rate as

$$\hat{\Gamma}^{0\nu} = (-0.99_{-0.27}^{+0.37}) \times 10^{-25} \text{ yr}^{-1}, \quad (6.10)$$

including only statistical uncertainty for now.

6.3.1 Statistical limits

With a negative fluctuation at $Q_{\beta\beta}$, it is clear that we have seen no evidence of $0\nu\beta\beta$ decay in these first two datasets. Therefore, the next step is to set a lower limit on the decay rate based on this data.

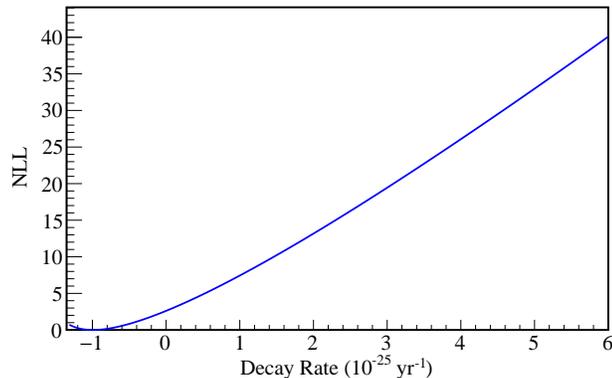


Figure 6.8: Negative log-likelihood (NLL) as a function of signal rate for the best-fit model in the ROI. The constant offset is irrelevant for setting a limit on the half-life, so it is subtracted out.

The first step in setting this limit (in a Bayesian formulation) is to calculate the likelihood of the data given the best-fit model at different signal rates. In reality, we compute the negative log-likelihood (NLL), which gives much more manageable quantities and has some nice properties, which we return to below. The NLL as a function of signal rate is shown in [Figure 6.8](#). We set a flat prior on the signal rate, $\Gamma^{0\nu} \geq 0$, restricting ourselves to the physical region where the decay rate is nonnegative. We then integrate the likelihood function ($e^{-\text{NLL}}$) in the region $\Gamma^{0\nu} \geq 0$ and normalize it so that the integral in this region is 1. By integrating again until we reach 0.9, we can set a 90%-C.L. upper limit on the signal rate. This gives us a 90%-C.L. limit of

$$\Gamma^{0\nu} < 0.48 \times 10^{-25} \text{ yr}^{-1} \text{ (stat. only)}, \quad (6.11)$$

or an equivalent half-life limit of

$$T_{1/2}^{0\nu} > 1.43 \times 10^{25} \text{ yr (stat. only)}. \quad (6.12)$$

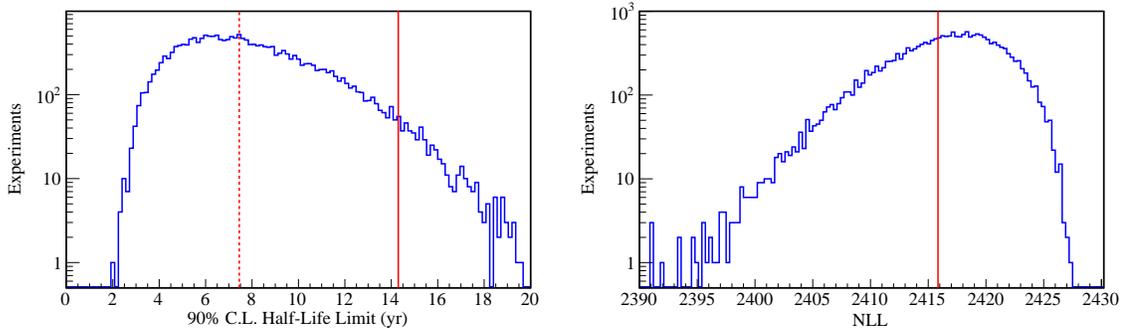


Figure 6.9: Distribution of 90%-C.L. half-life limits (left) and best-fit negative log-likelihoods (NLLs, right) from a collection of 20,000 pseudo-experiments. The median half-life limit, or the sensitivity of this search, is marked with a dotted line, and the half-life limit obtained from the true data is shown with a solid line.

6.3.2 Sensitivity

The next thing to evaluate is the sensitivity of our search, or the median half-life limit we expect from our data. This is important as it represents the overall performance of CUORE better than the actual half-life limit, which is subject to strong fluctuations, and it can also tell us how likely we are to obtain a limit at least as strong as the one we obtained.

To evaluate the sensitivity of our search, we perform a large number of pseudo-experiments, also known as toy Monte Carlo experiments. We repeatedly generate fake datasets using the best-fit model under the no-signal hypothesis and perform a fit, including a floating signal rate, to each. From each dataset, we obtain a 90%-C.L. limit on the decay rate, following the same procedure described above for the true data, and the minimum value of the NLL.

The results of these pseudo-experiments is shown in [Figure 6.9](#). The median sensitivity of this search is 7.45×10^{24} yr, and there is a 2.6% probability of obtaining a more stringent limit than the one obtained with our data. While our limit is unlikely, the fit itself is reasonable; our fit is in the 41st percentile of the distribution of NLL vales obtained in pseudo-experiments.

Although we observe a fairly strong downward fluctuation near $Q_{\beta\beta}$, we have no reason to believe that what we are observing is anything other than a normal statistical fluctuation. There is no evidence of any structure in the ROI, such as other γ lines, and from examining a wide region of the spectrum beyond that included in the ROI, it is clear that a flat background approximation is appropriate.

6.3.3 Evaluation of systematic errors

Until now, we have considered only statistical uncertainties in the evaluation of our half-life limit. However, there are several sources of potential systematic errors that we have not yet examined. For each, we evaluate an absolute systematic uncertainty σ_{abs} and relative (rate-dependent) systematic uncertainty σ_{rel} with a large number of pseudo-experiments generated with different signal rates. These uncertainties are:

- Resolution. The resolution of the $0\nu\beta\beta$ decay peak lends a relative uncertainty to the final fit result, which we evaluate by generating pseudo-experiments after varying the resolution by $\pm 1\sigma$.
- Energy reconstruction. If there is a bias in the energy reconstruction, it can move the $0\nu\beta\beta$ decay peak closer or further away from the ^{60}Co line, affecting the final results. We conservatively generate pseudo-experiments with $Q_{\beta\beta}$ shifted by ± 0.5 keV to account for any such uncertainty, although we see no evidence of any significant energy bias.
- Line shape. We fit the 2615 keV calibration line with a triple-Gaussian line shape, but this is an empirical (not physically motivated) fit. To account for the possibility that the $0\nu\beta\beta$ decay peak would appear without this structure, we generate pseudo-experiments with a Gaussian signal peak.
- Background shape. It is possible that the background is slightly sloped, although we see no evidence for this. Regardless, we generate pseudo-experiments

	Absolute uncertainty [yr ⁻²⁵]	Relative uncertainty [%]
Resolution	—	1.5
Energy reconstruction	—	0.2
Line shape	0.02	2.4
Background shape	0.05	0.8
Efficiency	—	1.8

Table 6.4: The systematic uncertainties used in the $0\nu\beta\beta$ decay analysis.

with the best-fit linear background from the data to see how our fitted signal rate is affected.

Finally, there is a several-percent uncertainty in our efficiency due to the low statistics in the physics spectrum lines that we use for the efficiency calculations. For this, we do not generate pseudo-experiments, and we simply assign the fractional uncertainty in the efficiency as a relative systematic uncertainty. The systematic errors used in this analysis are shown in [Table 6.4](#)¹².

Our general approach for incorporating systematic uncertainties into the limit-setting procedure is to modify the NLL curve to account for these uncertainties before integrating. We do this by first considering the NLL to be a χ^2 statistic, a consequence of Wilks' Theorem [147]. That is, we define

$$\chi_{\text{stat}}^2 = -2(\text{NLL} - \text{NLL}_0), \quad (6.13)$$

where the NLL of the best fit, NLL_0 , is subtracted away. In reality, this test statistic is not a perfect χ^2 distribution with one degree of freedom, because there are too few events in the ROI for Gaussian approximations to apply, but it is a reasonable approximation of one. Next, we calculate a χ_{syst}^2 value for our systematic uncertainty,

12. Systematic errors were evaluated by C. Davis.

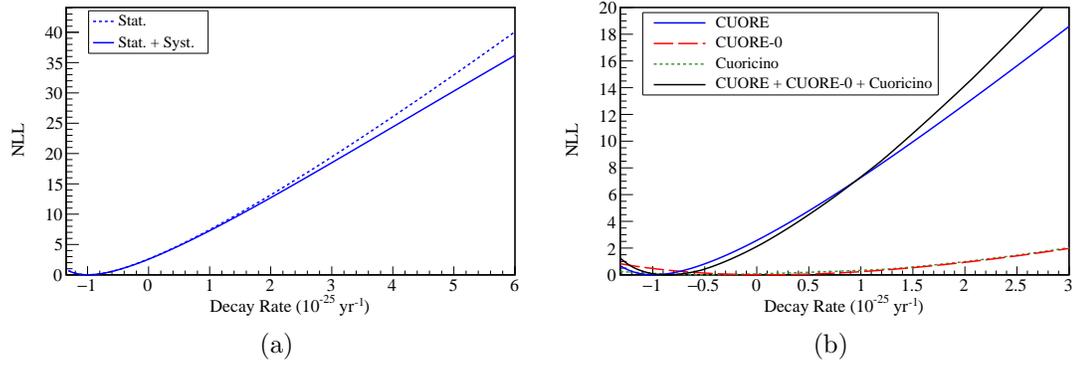


Figure 6.10: (a) NLL as a function of signal rate for the best-fit model in the CUORE ROI, including systematic uncertainties. (b) NLL curves from CUORE, CUORE-0, and Cuoricino, and their sum.

with

$$\chi_{\text{syst}}^2 = \frac{\Gamma^{0\nu} - \hat{\Gamma}^{0\nu}}{\sigma_{\text{syst}}(\Gamma^{0\nu})}, \quad \sigma_{\text{syst}}(\Gamma^{0\nu}) = \sum_i (\sigma_{i,\text{abs}} + \Gamma^{0\nu} \sigma_{i,\text{rel}})^2. \quad (6.14)$$

We combine these two as

$$\frac{1}{\chi_{\text{tot}}^2} = \frac{1}{\chi_{\text{stat}}^2} + \frac{1}{\chi_{\text{syst}}^2}, \quad (6.15)$$

which is effectively combining the statistical and systematic uncertainties in quadrature in the approximation of Gaussian-distributed uncertainties and a small systematic uncertainty [79]. If we then calculate a new NLL curve with

$$\text{NLL} = -\frac{1}{2}\chi_{\text{tot}}^2, \quad (6.16)$$

we see that this NLL curve is almost identical to the statistical one, but is slightly broadened to account for the systematic uncertainty. The original and modified NLL curves are shown in [Figure 6.10a](#).

6.3.4 Final fit results

We are now ready to set a lower limit on the rate of ^{130}Te $0\nu\beta\beta$ decay, incorporating both our statistical and systematic uncertainties. If we integrate the systematics-adjusted NLL curve following the same procedure as in [Section 6.3.1](#), we arrive at a limit of

$$\Gamma^{0\nu} < 0.49 \times 10^{-25} \text{ yr}^{-1} \text{ (90\% C.L.)}, \quad (6.17)$$

or an equivalent half-life limit of

$$T_{1/2}^{0\nu} > 1.40 \times 10^{25} \text{ yr (90\% C.L.)}. \quad (6.18)$$

We can also combine this result with the results from CUORE-0 and Cuoricino to slightly improve our limit. We do so by simply adding the NLL curves from these three experiments (see [Figure 6.10b](#)) and integrating the summed curve. If we do so, we obtain a limit of

$$T_{1/2}^{0\nu} > 1.53 \times 10^{25} \text{ yr (90\% C.L.)}. \quad (6.19)$$

6.3.5 Interpretation

With these first two CUORE datasets, we have set the strongest limit to date on the half-life of ^{130}Te $0\nu\beta\beta$ decay. We have far surpassed the sensitivity and limit set by CUORE-0 with just two months of data. This is in part due to the 19-times larger detector mass, but is also due to greatly improved background levels in the ROI. [Figure 6.11](#) shows clearly the improvement in these backgrounds moving from CUORE-0 to CUORE, including a 4-fold reduction in the ROI. This is almost entirely due to the new cryostat constructed for CUORE, as the CUORE-0 and CUORE towers are functionally identical.

The only physical parameter that CUORE can definitively measure is the half-life

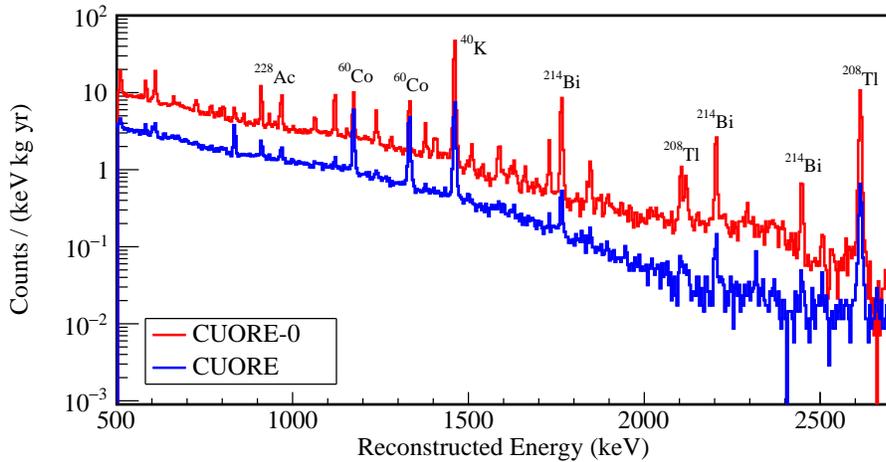


Figure 6.11: Comparison of the CUORE-0 and CUORE energy spectra, showing the large background reductions moving from CUORE-0 to CUORE.

of ^{130}Te $0\nu\beta\beta$ decay. However, as discussed in [Section 2.2](#), we can interpret this half-life as an effective Majorana neutrino mass $m_{\beta\beta}$ in the context of $0\nu\beta\beta$ decay mediated by light Majorana neutrino exchange. With a half-life limit of $T_{1/2}^{0\nu} > 1.53 \times 10^{25}$ yr, we obtain

$$m_{\beta\beta} > 140 - 390 \text{ meV}, \quad (6.20)$$

depending on the nuclear matrix elements used. Note that without any measurements of g_A in $0\nu\beta\beta$ decay, we must choose a value for the sake of comparison with other experiments, so we use the free-space value $g_A = 1.27$. This interpretation is shown in [Figure 6.12](#), along with the results from experiments using other isotopes and the results assuming we reach the ultimate 5-year sensitivity of CUORE, 9.0×10^{25} yr.

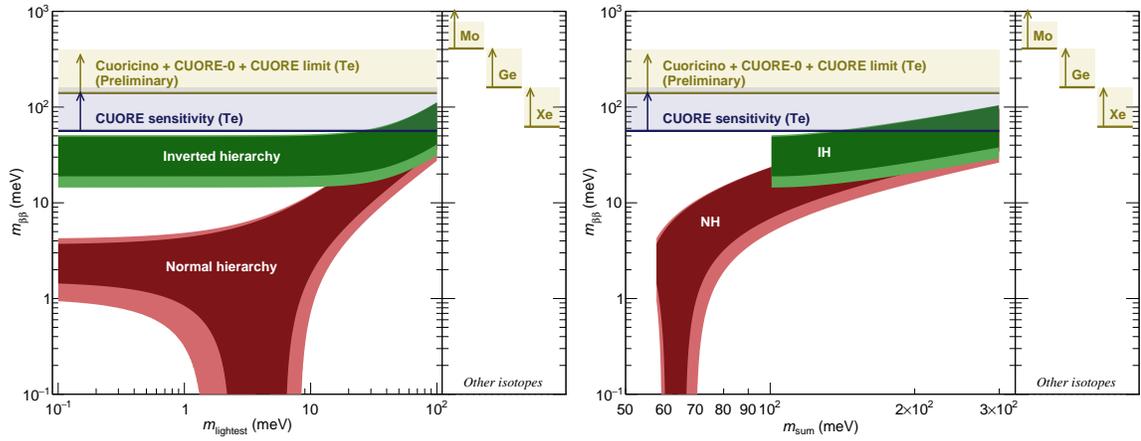


Figure 6.12: Limits on the effective Majorana mass from CUORE, plotted against the lightest neutrino mass (left) and the sum of the neutrino masses (right). Nuclear matrix elements are from [69, 74, 76, 78, 148], and phase space factors are from [68].

Chapter 7

Conclusion

“The most important discoveries will provide answers to questions that we do not yet know how to ask.”

— John Bahcall

CUORE has just begun to take data. The results presented here are based on 2 months of data, just a drop in the bucket compared to the scheduled 5 years of data collection. At the same time, it is never too early to look at the future past CUORE, both for bolometric detectors and for the $0\nu\beta\beta$ decay field as a whole. Below, I will discuss the possibilities for a successor experiment to CUORE, and what I see as the future for the field.

7.1 CUPID: The future for bolometric detectors

A successor experiment to CUORE would need to improve on the sensitivity of CUORE to $0\nu\beta\beta$ decay by at least an order or magnitude or more. Recalling the sensitivity figure of merit for $0\nu\beta\beta$ decay searches (Equation 2.29), there are 5 parameters that directly affect the sensitivity of an experiment like CUORE: the isotopic abundance of the double-beta emitter, the mass of the detector, the exposure time, the background index, and the energy resolution. The most straightforward feature

to improve is the isotopic abundance. In CUORE, we use TeO_2 with a ^{130}Te isotopic abundance of 34%; moving to material enriched to 92% in ^{130}Te , as has already been demonstrated [149], could nearly triple our sensitivity. But the other features are not as ripe for improvement with current technologies. CUORE already has a mass of 742 kg, composed of individual 750 g bolometers; significantly scaling up the detector mass is not feasible. The exposure time of CUORE is 5 years, and an experiment cannot operate for orders of magnitude longer than that. The background index has been reduced with over 20 years of work into making extremely radiopure materials, and there is no path towards a clear improvement in the energy resolution, which is already among the best of any $0\nu\beta\beta$ decay technologies.

Thus, we can point to one feature that is absolutely required on a successor experiment to CUORE: particle identification. The background near the Q -value for $0\nu\beta\beta$ decay is dominated by α particles in CUORE, and if we cannot clean the materials emitting these particles better than we are now, we must be able to detect them as background signals so they can be removed from the analysis. The bolometers as they are instrumented in CUORE have no ability to perform such discrimination, so additional detector technologies are required.

The proposed successor experiment to CUORE is named CUPID (CUORE Upgrade with Particle IDentification) [150, 151]. The technology for CUPID has not been finalized, but several possibilities are the focus of intense R&D programs. These efforts generally break down into two families: those using TeO_2 crystals and those using alternative crystals.

The main reason to switch away from TeO_2 crystals is that TeO_2 does not scintillate. The most straightforward way to differentiate between α particles and β/γ particles is to measure the scintillation light. Scintillation light produced from β/γ particles is significantly stronger than that produced by nuclear recoils from α particles, and the light is detected for a significantly longer amount of time [152]. An

experiment using 10 kg of ZnSe crystals, enriched in the double-beta emitter ^{82}Se , has shown almost total discrimination between α and β/γ particles using standard germanium bolometric light detectors for the scintillation light [153]. Other experiments have demonstrated similar rejection with smaller experiments using CdWO_4 enriched in ^{116}Cd [154] and $\text{ZnMoO}_4/\text{Li}_2\text{MoO}_4$ enriched in ^{100}Mo [155].

Perhaps the most compelling argument for staying with TeO_2 is that we already have significant expertise in producing large quantities of highly radiopure TeO_2 crystals, and we know that we can operate them with a good energy resolution. However, if we do use TeO_2 , we are left with detecting Cherenkov light or surface effects, both of which are significantly more difficult than detecting scintillation light. Specifically, detecting Cherenkov light is difficult because scintillation light is on the order of several keV, while Cherenkov light releases photons of a few eV. At the same time, α particles interacting in CUORE would produce no Cherenkov light at all [156], so discrimination is a simple binary of whether or not Cherenkov light is detected in a given interaction.

The Cherenkov light emitted as a particle passes through the TeO_2 bolometer can be detected in a number of ways. The first detection was with a simple germanium bolometric detector, but this does not have sufficient resolution to discriminate efficiently between α and β/γ particles [157]. Very efficient rejection in TeO_2 crystals using transition edge sensors has been demonstrated [158], as has rejection using electric fields across germanium bolometric detectors [149, 159, 160], taking advantage of the Neganov–Luke effect [161]. R&D is also underway for kinetic inductance detectors [162] and magnetic metallic calorimeters [163].

Finally, there are ways of performing background rejection without detecting either the scintillation or Cherenkov light from the crystals themselves. These rely on the fact that most of the backgrounds are from surface contamination, while true $0\nu\beta\beta$ decay events would be entirely internal to the crystals. In one example, it has

been shown that by surrounding a TeO_2 crystal with a scintillating foil, it is possible to detect the scintillation from surface events due to contamination with α emitters on the crystal itself [164]. It has also been shown that superconducting aluminum films surrounding small TeO_2 crystals can modify the recorded pulse shape when the incident particles are external to the crystals [165]. These films absorb some of the incident particle energy and then release heat into the crystal at a different time scale than the original energy deposition.

Whichever technology is ultimately chosen for CUPID, it will be a significant challenge to scale up the detector technology to the ton scale. Assuming we use take advantage of isotopic enrichment, all new crystals will need to be produced, be it TeO_2 crystals enriched in ^{130}Te or crystals of another type. An efficient, reliable, and reproducible method of producing any kind of light sensors or other technologies will need to be developed and tested. And ultimately, CUPID will either need to be operated in the CUORE cryostat or will need to invest significant resources in producing a new environment for the experiment.

There is one other concern for CUPID that is not a concern for CUORE: muons and muon-induced radioactivity. CUORE operates without a muon tagger or muon veto system, as the muon rate at LNGS is already quite low (~ 1 per square meter per hour [166]), and the vast majority of muons deposit energy in multiple bolometers, so they can be efficiently rejected by the anti-coincidence cuts in the final analysis. Some, however, survive, producing an estimated background of approximately 10^{-4} counts/(keV \cdot kg \cdot yr) in the $0\nu\beta\beta$ decay region of interest in CUORE [167]. This is only $\sim 1\%$ of CUORE's background, but would be likely be large background component of CUPID. Thus, an active muon tagging system around CUPID would be required in order to remove this background. There is also a background due to muon-induced neutrons, which is estimated to be an order of magnitude lower than that due to muons directly. This background would be harder to remove, but we may

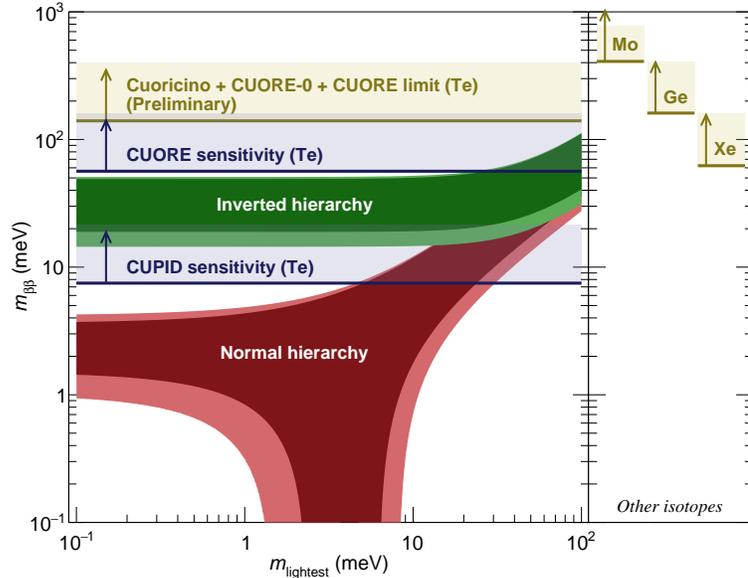


Figure 7.1: Expected effective Majorana mass sensitivity of CUPID.

be able to reduce it by removing data from a certain time period after each muon interaction in the tagger.

7.2 Sensitivity and discovery

CUPID, and other next-generation $0\nu\beta\beta$ decay experiments, are designed to completely probe the so-called inverted hierarchy region (see [Figure 7.1](#)). By the time these experiments have results, however, neutrino oscillation experiments should have a definitive answer as to whether the true hierarchy is normal or inverted. A global analysis of neutrino experiments finds a slight overall preference for the normal hierarchy [168], but it is by no means conclusive. NO ν A, a long-baseline oscillation experiment using accelerator neutrinos, has the best sensitivity to the mass hierarchy of any operating experiment [169], though a definitive word will likely come from future experiments. DUNE, another long-baseline experiment currently under construction, will have 5σ sensitivity to the mass hierarchy no mat-

ter what the CP violation phase is, though it will not be operational for approximately another decade [170]. Other proposed experiments, like PINGU [171], Hyper-Kamiokande [172], and ORCA/KM3NeT [173], will look at atmospheric neutrinos passing through the earth. Still other experiments, like JUNO [174] and the proposed RENO-50 [175], will look at reactor neutrinos. And finally, sensitive cosmological measurements that measure the sum of the neutrino masses could exclude the allowed mass sums in the inverted hierarchy altogether.

Suppose that these experiments show that the true hierarchy is the inverted hierarchy. In this situation, CUPID and other next-generation $0\nu\beta\beta$ decay experiments will make a discovery assuming that the neutrino is a Majorana particle, and our model of the decay mediated by light Majorana neutrino exchange is correct, with $g_A \approx 1.27$, and the theoretically calculated nuclear matrix elements and phase space elements are approximately correct. If the true hierarchy is inverted and no signal is detected, that will be a strong indication that our theoretical models or estimate of g_A are incorrect, or that the neutrino is not, in fact, a Majorana particle.

A more likely scenario, perhaps, is that experiments show that the true hierarchy is the normal hierarchy. In this case, there is no lower limit to the $0\nu\beta\beta$ decay rate. Of course, there is no particular reason to believe that the true values of the Majorana phases would result in a severe suppression of the decay rate, and there is also no particular reason for the lightest neutrino mass to be orders of magnitude less than the masses of the other two neutrino mass states. But the possibility remains that we could find ourselves in this situation in 20 years: no discovery of $0\nu\beta\beta$ decay and strong evidence that the true hierarchy is normal. In this case, what are the options?

A sensitivity to an effective Majorana neutrino mass of 1 meV would mean a $0\nu\beta\beta$ decay half-life of ^{130}Te of $\sim 10^{30}$ yr. With any of today's technology, probing such a half-life is an utter impossibility. A half-life of this order would require over 300 tons of ^{130}Te just to have one decay per year. To obtain 300 tons of ^{130}Te , it would require

triple the annual world production of tellurium, and experimentalists can only dream of operating a 300-ton background-free experiment.

There is much yet to be discovered in the field of neutrino physics. After all, we have done experiments with neutrinos for six decades, and we do not even know their masses! For now, we can hope for an exciting decade of surprises and discoveries, and we experimentalists can see what nature, and theorists, have in store for us.

Bibliography

- [1] C. Jensen. *Controversy and Consensus: Nuclear Beta Decay 1911–1934*. Birkhäuser Basel, Basel (2000). ISBN 978-3-0348-9569-9.
- [2] W. Pauli. Offener Brief an die Gruppe der Radioaktiven bei der Gauvereintagung zu Tübingen (December 4, 1930). Letter.
- [3] E. Fermi. *Versuch einer Theorie der β -Strahlen. I*. *Z. Phys.* **88**, 161–177 (1934).
- [4] W. Heisenberg. *Über den Bau der Atomkerne. I*. *Z. Phys.* **77**, 1–11 (1932).
- [5] I. Joliot-Curie. *Nobel Lecture: Artificial Production of Radioactive Elements* (1935).
- [6] P. A. M. Dirac. *The Quantum Theory of the Electron*. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **117**, 610–624 (1928).
- [7] C. D. Anderson. *The Positive Electron*. *Phys. Rev.* **43**, 491–494 (1933).
- [8] F. Guerra and N. Robotti. *Enrico Fermi’s Discovery of Neutron-Induced Artificial Radioactivity: The Influence of His Theory of Beta Decay*. *Phys. Perspect.* **11**, 379–404 (2009).
- [9] E. Majorana. *Teoria simmetrica dell’elettrone e del positrone*. *Nuovo Cim.* **14**, 171–184 (1937).
- [10] M. Goeppert-Mayer. *Double Beta-Disintegration*. *Phys. Rev.* **48**, 512–516 (1935).
- [11] M. Wang, G. Audi, A. H. Wapstra, F. G. Kondev, M. MacCormick, X. Xu, and B. Pfeiffer. *The AME2012 atomic mass evaluation*. *Chin. Phys. C* **36**, 1603–2014 (2012).
- [12] W. H. Furry. *On Transition Probabilities in Double Beta-Disintegration*. *Phys. Rev.* **56**, 1184–1193 (1939).
- [13] W. Haxton and G. J. Stephenson. *Double Beta Decay*. *Prog. Part. Nucl. Phys.* **12**, 409–479 (1984).

- [14] E. L. Fireman. [A Measurement of the Half-Life of Double Beta-Decay from \${}_{50}\text{Se}^{124}\$](#) . *Phys. Rev.* **75**, 323–324 (1949).
- [15] J. A. McCarthy. [Search for Double Beta Decay in \$\text{Ca}^{48}\$](#) . *Phys. Rev.* **97**, 1234–1236 (1955).
- [16] M. G. Inghram and J. H. Reynolds. [Double Beta-Decay of \$\text{Te}^{130}\$](#) . *Phys. Rev.* **78**, 822–823 (1950).
- [17] C. Alduino *et al.* (CUORE Collaboration). [Measurement of the two-neutrino double-beta decay half-life of \${}^{130}\text{Te}\$ with the CUORE-0 experiment](#). *Eur. Phys. J. C* **77**, 13 (2017).
- [18] S. R. Elliott, A. A. Hahn, and M. K. Moe. [Direct Evidence for Two-Neutrino Double-Beta Decay in \${}^{82}\text{Se}\$](#) . *Phys. Rev. Lett.* **59**, 2020–2023 (1987).
- [19] R. Arnold *et al.* (NEMO-3 Collaboration). [Measurement of the \$2\nu\beta\beta\$ decay half-life and search for the \$0\nu\beta\beta\$ decay of \${}^{116}\text{Cd}\$ with the NEMO-3 detector](#). *Phys. Rev. D* **95**, 170 (2017).
- [20] R. Arnold *et al.* (NEMO-3 Collaboration). [Measurement of the double-beta decay half-life and search for the neutrinoless double-beta decay of \${}^{48}\text{Ca}\$ with the NEMO-3 detector](#). *Phys. Rev. D* **93**, 112008 (2016).
- [21] R. Arnold *et al.* (NEMO-3 Collaboration). [Measurement of the \$2\nu\beta\beta\$ decay half-life of \${}^{150}\text{Nd}\$ and a search for \$0\nu\beta\beta\$ decay processes with the full exposure from the NEMO-3 detector](#). *Phys. Rev. D* **94**, 262 (2016).
- [22] M. Agostini *et al.* (GERDA Collaboration). [Results on \$\beta\beta\$ decay with emission of two neutrinos or Majorons in \${}^{76}\text{Ge}\$ from GERDA Phase I](#). *Eur. Phys. J. C* **75**, 1926 (2015).
- [23] R. Arnold *et al.* (NEMO-3 Collaboration). [Results of the search for neutrinoless double- \$\beta\$ decay in \${}^{100}\text{Mo}\$ with the NEMO-3 experiment](#). *Phys. Rev. D* **92**, 9 (2015).
- [24] J. B. Albert *et al.* (EXO-200 Collaboration). [Improved measurement of the \$2\nu\beta\beta\$ half-life of \${}^{136}\text{Xe}\$ with the EXO-200 detector](#). *Phys. Rev. C* **89**, 4 (2014).
- [25] A. Gando *et al.* (KamLAND-Zen Collaboration). [Measurement of the Double-Beta Decay Half-life of \${}^{136}\text{Xe}\$ in KamLAND-Zen](#). *Phys. Rev. C* **85**, 045504 (2012).
- [26] F. Reines and C. L. Cowan. [Detection of the Free Neutrino](#). *Phys. Rev.* **92**, 830–831 (1953).
- [27] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. [Detection of the Free Neutrino: a Confirmation](#). *Science* **124**, 103–104 (1956).

- [28] G. Feinberg. [Decays of the \$\mu\$ Meson in the Intermediate-Meson Theory](#). *Phys. Rev.* **110**, 1482–1483 (1958).
- [29] G. Danby, J.-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. [Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos](#). *Phys. Rev. Lett.* **9**, 36–44 (1962).
- [30] K. Kodama *et al.* (DONUT Collaboration). [Observation of tau neutrino interactions](#). *Phys. Lett. B* **504**, 218–224 (2001).
- [31] S. Chatrchyan *et al.* (CMS Collaboration). [Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC](#). *Phys. Lett. B* **716**, 30–61 (2012).
- [32] G. o. Aad (ATLAS Collaboration). [Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC](#). *Phys. Lett. B* **716**, 1–29 (2012).
- [33] J. N. Bahcall. [Solar Neutrinos. I. Theoretical](#). *Phys. Rev. Lett.* **12**, 300–302 (1964).
- [34] R. Davis Jr, D. S. Harmer, and K. C. Hoffman. [Search for Neutrinos from the Sun](#). *Phys. Rev. Lett.* **20**, 1205–1209 (1968).
- [35] B. T. Cleveland, T. Daily, R. Davis Jr, J. R. Distel, K. Lande, C. K. Lee, P. S. Wildenhain, and J. Ullman. [Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector](#). *Astrophys. J.* **496**, 505–526 (1998).
- [36] K. S. Hirata *et al.* [Experimental study of the atmospheric neutrino flux](#). *Phys. Lett. B* **205**, 416–420 (1988).
- [37] D. Casper *et al.* [Measurement of atmospheric neutrino composition with the IMB-3 detector](#). *Phys. Rev. Lett.* **66**, 2561–2564 (1991).
- [38] B. Pontecorvo. [Mesonium and Antimesonium](#). *J. Exp. Theor. Phys.* **33**, 549–551 (1957).
- [39] Z. Maki, M. Nakagawa, and S. Sakata. [Remarks on the Unified Model of Elementary Particles](#). *Prog. Theor. Phys.* **28**, 870–880 (1962).
- [40] Y. Fukuda *et al.* (Super-Kamiokande Collaboration). [Evidence for Oscillation of Atmospheric Neutrinos](#). *Phys. Rev. Lett.* **81**, 1562–1567 (1998).
- [41] Q. R. Ahmad *et al.* (SNO Collaboration). [Measurement of the Rate of \$\nu_e + d \rightarrow p + p + e^-\$ Interactions Produced by \$^8\text{B}\$ Solar Neutrinos at the Sudbury Neutrino Observatory](#). *Phys. Rev. Lett.* **87**, 071301 (2001).

- [42] Q. R. Ahmad *et al.* (SNO Collaboration). [Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory](#). *Phys. Rev. Lett.* **89**, 011301 (2002).
- [43] J. N. Bahcall, M. H. Pinsonneault, and S. Basu. [Solar Models: Current Epoch and Time Dependences, Neutrinos, and Helioseismological Properties](#). *Astrophys. J.* **555**, 990–1012 (2001).
- [44] W. Hampel *et al.* (GALLEX Collaboration). [GALLEX solar neutrino observations: results for GALLEX IV](#). *Phys. Lett. B* **447**, 127–133 (1999).
- [45] J. N. Abdurashitov *et al.* (SAGE Collaboration). [Measurement of the solar neutrino capture rate with gallium metal](#). *Phys. Rev. C* **60**, 055801 (1999).
- [46] M. Altmann *et al.* (GNO Collaboration). [Complete results for five years of GNO solar neutrino observations](#). *Phys. Lett. B* **616**, 174–190 (2005).
- [47] S. Adrián-Martínez *et al.* (ANTARES Collaboration). [Measurement of atmospheric neutrino oscillations with the ANTARES neutrino telescope](#). *Phys. Lett. B* **714**, 224–230 (2012).
- [48] M. G. Aartsen *et al.* (IceCube Collaboration). [Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of IceCube DeepCore data](#). *Phys. Rev. D* **91**, 072004 (2015).
- [49] K. Abe *et al.* (Super-Kamiokande Collaboration). [Solar neutrino measurements in Super-Kamiokande-IV](#). *Phys. Rev. D* **94**, 913 (2016).
- [50] S. Abe *et al.* (KamLAND Collaboration). [Precision Measurement of Neutrino Oscillation Parameters with KamLAND](#). *Phys. Rev. Lett.* **100**, 221803 (2008).
- [51] C. Arpesella *et al.* (Borexino Collaboration). [Direct Measurement of the \$^7\text{Be}\$ Solar Neutrino Flux with 192 Days of Borexino Data](#). *Phys. Rev. Lett.* **101**, 091302 (2008).
- [52] F. P. An *et al.* (Daya Bay Collaboration). [Observation of Electron-Antineutrino Disappearance at Daya Bay](#). *Phys. Rev. Lett.* **108**, 171803 (2012).
- [53] Y. Abe *et al.* (Double Chooz Collaboration). [Indication of Reactor \$\bar{\nu}_e\$ Disappearance in the Double Chooz Experiment](#). *Phys. Rev. Lett.* **108**, 131801 (2012).
- [54] J. K. Ahn *et al.* (RENO Collaboration). [Observation of Reactor Electron Antineutrinos Disappearance in the RENO Experiment](#). *Phys. Rev. Lett.* **108**, 191802 (2012).
- [55] P. Adamson *et al.* (MINOS Collaboration). [Combined Analysis of \$\nu_\mu\$ Disappearance and \$\nu_\mu \rightarrow \nu_e\$ Appearance in MINOS Using Accelerator and Atmospheric Neutrinos](#). *Phys. Rev. Lett.* **112**, 984 (2014).

- [56] K. Abe *et al.* (T2K Collaboration). [Precise Measurement of the Neutrino Mixing Parameter \$\theta_{23}\$ from Muon Neutrino Disappearance in an Off-Axis Beam](#). *Phys. Rev. Lett.* **112**, 429 (2014).
- [57] P. Adamson *et al.* (NOvA Collaboration). [First measurement of muon-neutrino disappearance in NOvA](#). *Phys. Rev. D* **93**, 172 (2016).
- [58] P. A. R. Ade *et al.* (Planck Collaboration). [Planck 2015 results - XIII. Cosmological parameters](#). *Astron. Astrophys.* **594**, A13 (2016).
- [59] C. Patrignani *et al.* (Particle Data Group). [Review of Particle Physics](#). *Chin. Phys. C* **40**, 100001 (2016).
- [60] R. N. Mohapatra and G. Senjanović. [Neutrino Mass and Spontaneous Parity Nonconservation](#). *Phys. Rev. Lett.* **44**, 912–915 (1980).
- [61] P. Minkowski. [\$\mu \rightarrow e\gamma\$ at a rate of one out of \$10^9\$ muon decays?](#) *Phys. Lett. B* **67**, 421–428 (1977).
- [62] J. Lesgourgues and S. Pastor. [Neutrino Mass from Cosmology](#). *Adv. High Energy Phys.* **2012**, 1–34 (2012).
- [63] R. Emami *et al.* [Evidence of Neutrino Enhanced Clustering in a Complete Sample of Sloan Survey Clusters, Implying \$\sum m_\nu = 0.11 \pm 0.03\$ eV](#) (2017). arXiv:1711.05210.
- [64] V. N. Aseev *et al.* [Upper limit on the electron antineutrino mass from the Troitsk experiment](#). *Phys. Rev. D* **84**, 112003 (2011).
- [65] C. Kraus *et al.* [Final results from phase II of the Mainz neutrino mass search in tritium \$\beta\$ decay](#). *Eur. Phys. J. C* **40**, 447–468 (2005).
- [66] K. A. Olive *et al.* (Particle Data Group). [Review of Particle Physics](#). *Chin. Phys. C* **38**, 090001 (2014).
- [67] J. Schechter and J. W. F. Valle. [Neutrinoless double- \$\beta\$ decay in \$SU\(2\) \times U\(1\)\$ theories](#). *Phys. Rev. D* **25**, 2951–2954 (1982).
- [68] J. Kotila and F. Iachello. [Phase-space factors for double- \$\beta\$ decay](#). *Phys. Rev. C* **85**, 034316 (2012).
- [69] J. Menéndez, A. Poves, E. Caurier, and F. Nowacki. [Disassembling the nuclear matrix elements of the neutrinoless \$\beta\beta\$ decay](#). *Nucl. Phys. A* **818**, 139–151 (2009).
- [70] A. Arima and F. Iachello. [Collective Nuclear States as Representations of a \$SU\(6\)\$ Group](#). *Phys. Rev. Lett.* **35**, 1069–1072 (1975).
- [71] F. Šimković, A. Faessler, V. Rodin, P. Vogel, and J. Engel. [Anatomy of the \$0\nu\beta\beta\$ nuclear matrix elements](#). *Phys. Rev. C* **77**, 045503 (2008).

- [72] J. Barea, J. Kotila, and F. Iachello. Nuclear matrix elements for double- β decay. *Phys. Rev. C* **87**, 014315 (2013).
- [73] A. Neacsu and M. Horoi. Shell model studies of the ^{130}Te neutrinoless double- β decay. *Phys. Rev. C* **91**, 024309 (2015).
- [74] J. Barea, J. Kotila, and F. Iachello. $0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements in the interacting boson model with isospin restoration. *Phys. Rev. C* **91**, 034304 (2015).
- [75] A. Faessler, M. González, S. Kovalenko, and F. Šimkovic. Arbitrary mass Majorana neutrinos in neutrinoless double beta decay. *Phys. Rev. D* **90**, 096010 (2014).
- [76] J. Hyvärinen and J. Suhonen. Nuclear matrix elements for $0\nu\beta\beta$ decays with light or heavy Majorana-neutrino exchange. *Phys. Rev. C* **91**, 024613 (2015).
- [77] P. K. Rath, R. Chandra, K. Chaturvedi, P. Lohani, P. K. Raina, and J. G. Hirsch. Neutrinoless $\beta\beta$ decay transition matrix elements within mechanisms involving light Majorana neutrinos, classical Majorons, and sterile neutrinos. *Phys. Rev. C* **88**, 064322 (2013).
- [78] T. R. Rodríguez and G. Martínez-Pinedo. Energy Density Functional Study of Nuclear Matrix Elements for Neutrinoless $\beta\beta$ Decay. *Phys. Rev. Lett.* **105**, 252503 (2010).
- [79] J. L. Ouellet. The Search for $0\nu\beta\beta$ Decay in ^{130}Te with CUORE-0. Ph.D. thesis, University of California, Berkeley (2015).
- [80] M. Agostini *et al.* (GERDA Collaboration). Background-free search for neutrinoless double- β decay of ^{76}Ge with GERDA. *Nature* **544**, 47–52 (2017).
- [81] R. Arnold *et al.* (NEMO-3 Collaboration). Search for neutrinoless double-beta decay of ^{100}Mo with the NEMO-3 detector. *Phys. Rev. D* **89**, 111101 (2014).
- [82] K. Alfonso *et al.* (CUORE Collaboration). Search for Neutrinoless Double-Beta Decay of ^{130}Te with CUORE-0. *Phys. Rev. Lett.* **115**, 102502 (2015).
- [83] J. B. Albert *et al.* (EXO-200 Collaboration). Search for Majorana neutrinos with the first two years of EXO-200 data. *Nature* **510**, 229–234 (2014).
- [84] A. Gando *et al.* (KamLAND-Zen Collaboration). Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen. *Phys. Rev. Lett.* **117**, 082503 (2016).
- [85] N. Abgrall *et al.* (MAJORANA Collaboration). The Majorana Demonstrator Neutrinoless Double-Beta Decay Experiment. *Adv. High Energy Phys.* **2014**, 1–18 (2014).

- [86] R. Arnold *et al.* (SuperNEMO Collaboration). [Probing new physics models of neutrinoless double beta decay with SuperNEMO](#). *Eur. Phys. J. C* **70**, 927–943 (2010).
- [87] C. Alduino *et al.* (CUORE Collaboration). [CUORE sensitivity to \$0\nu\beta\beta\$ decay](#). *Eur. Phys. J. C* **77**, 532 (2017).
- [88] A. Maio (for the SNO+ Collaboration). [Search for Majorana neutrinos with the SNO+ detector at SNOLAB](#). *J. Phys. Conf. Ser.* **587**, 012030 (2015).
- [89] M. Auger *et al.* (EXO-200 Collaboration). [The EXO-200 detector, part I: detector design and construction](#). *J. Instrum.* **7**, P05010–P05010 (2012).
- [90] J. Martín-Albo *et al.* (NEXT Collaboration). [Sensitivity of NEXT-100 to neutrinoless double beta decay](#). *J. High Energy Phys.* **2016**, 1–30 (2016).
- [91] X. Chen *et al.* (PandaX-III Collaboration). [PandaX-III: Searching for neutrinoless double beta decay with high pressure \$^{136}\text{Xe}\$ gas time projection chambers](#). *Sci. China Phys. Mech. Astron.* **60**, 061011 (2017).
- [92] K. Asakura *et al.* (KamLAND-Zen Collaboration). [Results from KamLAND-Zen](#). *AIP Conf. Proc.* **1666**, 170003 (2015).
- [93] N. Abgrall *et al.* (LEGEND Collaboration). [The Large Enriched Germanium Experiment for Neutrinoless Double Beta Decay \(LEGEND\)](#). *AIP Conf. Proc.* **1894**, 020027 (2017).
- [94] M. Redshaw, B. Mount, E. Myers, and F. Avignone. [Masses of \$^{130}\text{Te}\$ and \$^{130}\text{Xe}\$ and Double- \$\beta\$ -Decay \$Q\$ Value of \$^{130}\text{Te}\$](#) . *Phys. Rev. Lett.* **102**, 212502 (2009).
- [95] N. D. Scielzo *et al.* [Double- \$\beta\$ -decay \$Q\$ values of \$^{130}\text{Te}\$, \$^{128}\text{Te}\$, and \$^{120}\text{Te}\$](#) . *Phys. Rev. C* **80**, 025501 (2009).
- [96] S. Rahaman, V.-V. Elomaa, T. Eronen, J. Hakala, A. Jokinen, A. Kankainen, J. Rissanen, J. Suhonen, C. Weber, and J. Äystö. [Double-beta decay \$Q\$ values of \$^{116}\text{Cd}\$ and \$^{130}\text{Te}\$](#) . *Phys. Lett. B* **703**, 412–416 (2011).
- [97] E. E. Haller, N. P. Palaio, M. Rodder, W. L. Hansen, and E. Kreysa. [NTD Germanium: A Novel Material for Low Temperature Bolometers](#). In *Neutron Transmutation Doping of Semiconductor Materials*, 21–36. Springer US, Boston, MA (1984). ISBN 978-1-4612-9675-1.
- [98] M. Barucci, C. Brofferio, A. Giuliani, E. Gottardi, I. Peroni, and G. Ventura. [Measurement of Low Temperature Specific Heat of Crystalline \$\text{TeO}_2\$ for the Optimization of Bolometric Detectors](#). *J. Low Temp. Phys.* **123**, 303–314 (2001).
- [99] A. L. Efros and B. I. Shklovskii. [Coulomb gap and low temperature conductivity of disordered systems](#). *J. Phys. C* **8**, L49–L51 (1975).

- [100] M. Vignati. [Model of the Response Function of CUORE Bolometers](#). Ph.D. thesis, Sapienza – Università di Roma, Rome, Italy (2009).
- [101] P. W. Bridgman. [Certain Physical Properties of Single Crystals of Tungsten, Antimony, Bismuth, Tellurium, Cadmium, Zinc, and Tin](#). *Proc. Amer. Acad. Arts Sci.* **60**, 305 (1925).
- [102] Y. Chu, Y. Li, Z. Ge, G. Wu, and H. Wang. [Growth of the high quality and large size paratellurite single crystals](#). *J. Cryst. Growth* **295**, 158–161 (2006).
- [103] C. Arnaboldi *et al.* [Production of high purity TeO₂ single crystals for the study of neutrinoless double beta decay](#). *J. Cryst. Growth* **312**, 2999–3008 (2010).
- [104] F. Alessandria *et al.* (CUORE Collaboration). [CUORE crystal validation runs: Results on radioactive contamination and extrapolation to CUORE background](#). *Astropart. Phys.* **35**, 839–849 (2012).
- [105] B. S. Wang, E. B. Norman, N. D. Scielzo, A. R. Smith, K. J. Thomas, and S. A. Wender. [Cosmogenic-neutron activation of TeO₂ and implications for neutrinoless double- \$\beta\$ decay experiments](#). *Phys. Rev. C* **92**, 024620 (2015).
- [106] A. F. Barghouty *et al.* [Measurements of proton-induced radionuclide production cross sections to evaluate cosmic-ray activation of tellurium](#). *Nucl. Instr. Meth. B* **295**, 16–21 (2013).
- [107] E. Fiorini and T. O. Niinikoski. [Low-temperature calorimetry for rare decays](#). *Nucl. Instr. Meth.* **224**, 83–88 (1984).
- [108] A. Alessandrello *et al.* [Large mass, low temperature, low background detectors](#). *Nucl. Instr. Meth. A* **315**, 263–266 (1992).
- [109] A. Alessandrello *et al.* [A search for neutrinoless double beta decay of ¹³⁰Te with a bolometric detector](#). *Nucl. Phys. B (Proc. Suppl.)* **28**, 229–232 (1992).
- [110] A. Alessandrello *et al.* [A search for neutrinoless double beta decay of ¹³⁰Te with a thermal detector](#). *Phys. Lett. B* **285**, 176–182 (1992).
- [111] A. Alessandrello *et al.* [TeO₂ bolometers to search for Double Beta Decay of ¹³⁰Te: status of art](#). *J. Low Temp. Phys.* **93**, 201–206 (1993).
- [112] A. Alessandrello *et al.* [Improvements in ¹³⁰Te double beta decay search with cryogenic TeO₂ array detectors](#). *Nucl. Phys. B (Proc. Suppl.)* **48**, 238–240 (1996).
- [113] C. Arnaboldi *et al.* [A calorimetric search on double beta decay of ¹³⁰Te](#). *Phys. Lett. B* **557**, 167–175 (2003).
- [114] E. Fiorini. [CUORE: a cryogenic underground observatory for rare events](#). *Phys. Rep.* **307**, 309–317 (1998).

- [115] E. Andreotti *et al.* (Cuoricino Collaboration). [\$^{130}\text{Te}\$ neutrinoless double-beta decay with CUORICINO](#). *Astropart. Phys.* **34**, 822–831 (2011).
- [116] C. Alduino *et al.* (CUORE Collaboration). [CUORE-0 detector: design, construction and operation](#). *J. Instrum.* **11**, P07009–P07009 (2016).
- [117] F. Alessandria *et al.* (CUORE Collaboration). [Validation of techniques to mitigate copper surface contamination in CUORE](#). *Astropart. Phys.* **45**, 13–22 (2013).
- [118] A. Alessandrello *et al.* [Measurements of internal radioactive contamination in samples of Roman lead to be used in experiments on rare events](#). *Nucl. Instr. Meth. B* **142**, 163–172 (1998).
- [119] C. E. Pagliarone *et al.* (CUORE Collaboration). [Status of CUORE Experiment and latest results from CUORE-0](#). *Nuovo Cimento C* **39**, 375 (2016).
- [120] F. Alessandria *et al.* (CUORE Collaboration). [Sensitivity and Discovery Potential of CUORE to Neutrinoless Double-Beta Decay](#) (2011). arXiv:1109.0494.
- [121] J. S. Cushman *et al.* [The detector calibration system for the CUORE cryogenic bolometer array](#). *Nucl. Instr. Meth. A* **844**, 32–44 (2017).
- [122] M. Vignati. [Model of the response function of large mass bolometric detectors](#). *J. Appl. Phys* **108**, 084903 (2010).
- [123] D. R. Artusa *et al.* (CUORE Collaboration). [Initial performance of the CUORE-0 experiment](#). *Eur. Phys. J. C* **74**, 1–7 (2014).
- [124] A. Nucciotti *et al.* [Design of the Cryogen-Free Cryogenic System for the CUORE Experiment](#). *J. Low Temp. Phys.* **151**, 662–668 (2008).
- [125] L. M. Ejzak. [Energy Calibration for a Sensitive Search for Neutrinoless Double-Beta Decay: Using the Cuoricino Experience to Prepare for CUORE](#). Ph.D. thesis, University of Wisconsin – Madison (2013).
- [126] C. Alduino *et al.* (CUORE Collaboration). [Analysis techniques for the evaluation of the neutrinoless double- \$\beta\$ decay lifetime in \$^{130}\text{Te}\$ with the CUORE-0 detector](#). *Phys. Rev. C* **93**, 045503 (2016).
- [127] E. Browne. [Nuclear Data Sheets for \$A = 232\$](#) . *Nucl. Data Sheets* **107**, 2579–2648 (2006).
- [128] K. J. Thomas, A. R. Smith, Y. D. Chan, E. B. Norman, B. S. Wang, and D. L. Hurley. [Low Background Counting at the LBNL Low Background Facility](#). *AIP Conf. Proc.* **1549**, 20–25 (2013).
- [129] G. Ventura, M. Barucci, E. Gottardi, and I. Peroni. [Low temperature thermal conductivity of Kevlar](#). *Cryogenics* **40**, 489–491 (2000).

- [130] W.F. Lake Corporation. *PTFE Coated Aramid Thread (R722-70)*. URL <http://www.wflake.com/thread/Kevlar%20R721-Series.pdf>.
- [131] S. Agostinelli *et al.* *Geant4 — a simulation toolkit*. *Nucl. Instr. Meth. A* **506**, 250–303 (2003).
- [132] N. I. Chott *et al.* *Commissioning of the CUORE Cryostat*. In *Proceedings of “Technology and Instrumentation in Particle Physics”*, PoS(WAM2004)012 (2014).
- [133] C. Arnaboldi, A. Baù, P. Carniti, L. Cassina, A. Giachero, C. Gotti, M. Maino, A. Passerini, and G. Pessina. *Very low noise AC/DC power supply systems for large detector arrays*. *Rev. Sci. Instrum.* **86**, 124703 (2015).
- [134] G. Pessina. *Room temperature differential voltage sensitive preamplifier for large mass bolometric detectors*. *Nucl. Instr. Meth. A* **444**, 132–135 (2000).
- [135] C. Arnaboldi, M. Cariello, S. Di Domizio, A. Giachero, and G. Pessina. *A programmable multichannel antialiasing filter for the CUORE experiment*. *Nucl. Instr. Meth. A* **617**, 327–328 (2010).
- [136] A. Alessandrello *et al.* *Methods for response stabilization in bolometers for rare decays*. *Nucl. Instr. Meth. A* **412**, 454–464 (1998).
- [137] C. Arnaboldi, G. Pessina, and E. Previtali. *A programmable calibrating pulse generator with multioutputs and very high stability*. *IEEE T. Nucl. Sci.* **50**, 979–986 (2002).
- [138] P. Carniti, L. Cassina, A. Giachero, C. Gotti, and G. Pessina. *A High Precision Pulse Generation and Stabilization System for Bolometric Experiments* (2017). arXiv:1710.05565.
- [139] S. Di Domizio. *Search for double beta decay to excited states with CUORICINO and data acquisition system for CUORE*. Ph.D. thesis, Università degli Studi di Genova (2009).
- [140] R. Brun and F. Rademakers. *ROOT — An object oriented data analysis framework*. *Nucl. Instr. Meth. A* **389**, 81–86 (1997).
- [141] S. Di Domizio, F. Orio, and M. Vignati. *Lowering the energy threshold of large-mass bolometric detectors*. *J. Instrum.* **6**, P02007–P02007 (2011).
- [142] M. J. Martin. *Nuclear Data Sheets for A = 208*. *Nucl. Data Sheets* **108**, 1583–1806 (2007).
- [143] E. Browne. *Nuclear Data Sheets for A = 212*. *Nucl. Data Sheets* **104**, 427–496 (2005).
- [144] K. Abusaleem. *Nuclear Data Sheets for A = 228*. *Nucl. Data Sheets* **116**, 163–262 (2014).

- [145] P. C. Mahalanobis. [On the Generalised Distance in Statistics](#). *Proc. Natl. Inst. Sci. India* **49**–55 (1936).
- [146] C. J. Clopper and E. S. Pearson. [The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial](#). *Biometrika* **26**, 404–413 (1934).
- [147] S. S. Wilks. [The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses](#). *Ann. Math. Stat* **9**, 60–62 (1938).
- [148] F. Šimkovic, V. Rodin, A. Faessler, and P. Vogel. [\$0\nu\beta\beta\$ and \$2\nu\beta\beta\$ nuclear matrix elements, quasiparticle random-phase approximation, and isospin symmetry restoration](#). *Phys. Rev. C* **87**, 045501 (2013).
- [149] D. R. Artusa *et al.* [Enriched TeO₂ bolometers with active particle discrimination: Towards the CUPID experiment](#). *Phys. Lett. B* **767**, 321–329 (2017).
- [150] G. Wang *et al.* (CUPID Interest Group). [CUPID: CUORE Upgrade with Particle IDentification](#) (2015). arXiv:1504.03599.
- [151] G. Wang *et al.* (CUPID Interest Group). [R&D towards CUPID \(CUORE Upgrade with Particle IDentification\)](#) (2015). arXiv:1504.03612.
- [152] J. W. Beeman *et al.* [Characterization of bolometric light detectors for rare event searches](#). *J. Instrum.* **8**, P07021–P07021 (2013).
- [153] N. Casali *et al.* [Scintillating bolometric technique for the neutrino-less double beta decay search: The LUCIFER/CUPID-0 experiment](#). *Nucl. Instr. Meth. A* **845**, 342–346 (2017).
- [154] A. S. Barabash *et al.* [First test of an enriched ¹¹⁶CdWO₄ scintillating bolometer for neutrinoless double-beta-decay searches](#). *Eur. Phys. J. C* **76**, 106301 (2016).
- [155] E. Armengaud *et al.* [Development of ¹⁰⁰Mo-containing scintillating bolometers for a high-sensitivity neutrinoless double-beta decay search](#). *Eur. Phys. J. C* **77**, 785 (2017).
- [156] T. T. de Fatis. [Cherenkov emission as a positive tag of double beta decays in bolometric experiments](#). *Eur. Phys. J. C* **65**, 359–361 (2009).
- [157] N. Casali *et al.* [TeO₂ bolometers with Cherenkov signal tagging: towards next-generation neutrinoless double-beta decay experiments](#). *Eur. Phys. J. C* **75**, 12 (2015).
- [158] K. Schäffner *et al.* [Particle discrimination in TeO₂ bolometers using light detectors read out by transition edge sensors](#). *Astropart. Phys.* **69**, 30–36 (2015).
- [159] L. Bergé *et al.* [Complete event-by-event \$\alpha/\gamma\(\beta\)\$ separation in a full-size TeO₂ CUORE bolometer by simultaneous heat and light detection](#) (2017). arXiv:1710.03459.

- [160] L. Pattavina *et al.* [Background Suppression in Massive TeO₂ Bolometers with Neganov–Luke Amplified Light Detectors](#). *J. Low Temp. Phys.* **184**, 286–291 (2015).
- [161] P. N. Luke. [Voltage-assisted calorimetric ionization detector](#). *J. Appl. Phys* **64**, 6858–6860 (1988).
- [162] E. S. Battistelli *et al.* [CALDER: neutrinoless double-beta decay identification in TeO₂ bolometers with kinetic inductance detectors](#). *Eur. Phys. J. C* **75**, 353 (2015).
- [163] A. Giachero. [Assess the neutrino mass with micro and macro calorimeter approach](#). *J. Phys. Conf. Ser.* **841**, 012027 (2017).
- [164] M. Biassoni, C. Brofferio, C. Bucci, L. Canonica, M. L. di Vacri, P. Gorla, M. Pavan, and M. Yeh. [Rejection of Alpha Surface Background in Non-scintillating Bolometric Detectors: The ABSuRD Project](#). *J. Low Temp. Phys.* **184**, 879–884 (2016).
- [165] C. Nones, L. Bergé, L. Dumoulin, S. Marnieros, and E. Olivieri. [Superconducting Aluminum Layers as Pulse Shape Modifiers: An Innovative Solution to Fight Against Surface Background in Neutrinoless Double Beta Decay Experiments](#). *J. Low Temp. Phys.* **167**, 1029–1034 (2012).
- [166] M. Ambrosio *et al.* (MACRO Collaboration). [Vertical muon intensity measured with MACRO at the Gran Sasso laboratory](#). *Phys. Rev. D* **52**, 3793–3802 (1995).
- [167] C. Alduino *et al.* (CUORE Collaboration). [The projected background for the CUORE experiment](#). *Eur. Phys. J. C* **77**, 543 (2017).
- [168] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz. [Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity](#). *J. High Energy Phys.* **2017**, 87 (2017).
- [169] P. Adamson *et al.* (NOvA Collaboration). [First Measurement of Electron Neutrino Appearance in NOvA](#). *Phys. Rev. Lett.* **116**, 151806 (2016).
- [170] R. Acciarri *et al.* (DUNE Collaboration). [Long-Baseline Neutrino Facility \(LBNF\) and Deep Underground Neutrino Experiment \(DUNE\)](#) (2016). arXiv:1601.05471.
- [171] M. G. Aartsen *et al.* (IceCube-Gen2 Collaboration). [PINGU: a vision for neutrino and particle physics at the South Pole](#). *J. Phys. G* **44**, 054006 (2017).
- [172] K. Abe *et al.* (Hyper-Kamiokande Collaboration). [Physics Potentials with the Second Hyper-Kamiokande Detector in Korea](#) (2016). arXiv:1611.06118.

- [173] S. Adrián-Martínez *et al.* (KM3NeT Collaboration). [Letter of intent for KM3NeT 2.0](#). *J. Phys. G* **43**, 084001 (2016).
- [174] F. An *et al.* (JUNO Collaboration). [Neutrino physics with JUNO](#). *J. Phys. G* **43**, 030401 (2016).
- [175] S.-B. Kim. [New results from RENO and prospects with RENO-50](#). *Nucl. Part. Phys. Proc.* **265-266**, 93–98 (2015).

Appendix A

Calibration electronics

Documented here is all of the wiring and electronics for the Detector Calibration System (DCS) electronic controls and sensors. We begin with an overview of the system and the components. Refer to [Figure A.1](#) for the overview diagram of all the components of the DCS electronics system.

A.1 In the rack

The computer that controls the DCS is a Kingstar Supermicro 4U rack-mountable computer (SC743T-665B). The general connection diagram for all DCS components is in [Figure A.1](#), and the rack diagram is in [Figure A.2](#).

The computer connects to the internet with an ethernet cord through a rack-mounted LAN switch. It is connected to the CUORE underground network within the LNGS network. Access to the computer from outside LNGS is only possible with a VPN or through another computer on the LNGS network.

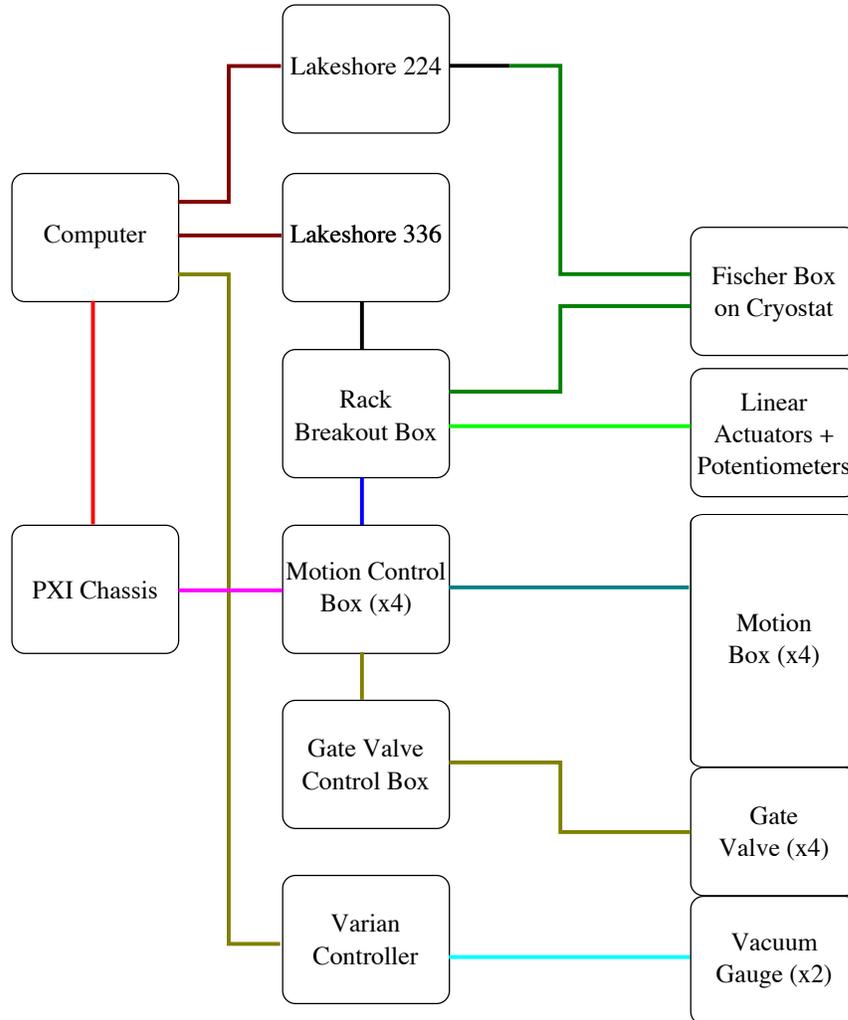
The computer is the main hub for all system controls, data taking, and the user interface. All of the control commands are issued by LabVIEW from the computer. The data is gathered from the various subsystems and put into a PostgreSQL database hosted locally on the computer, and the quantitative data is also placed into JSON

log files to be added to the Slow Monitoring system.

The computer connects to a National Instruments PXI chassis. Through the cards in this PXI chassis, the computer communicates with the stepper motor controllers, which control the motors in the motion boxes and the linear actuators at the center of the cryostat.

The computer reads temperature data from two LakeShore devices via USB cables; a Model 336 measures the temperature of the 4-K thermalizers, and a Model 224 measures the temperature of the thermometers on the 600 mK tubes. It reads pressure data from a Varian XGS-600.

Power to the rack is provided by two outlets on the second floor of the CUORE hut that are connected to a large central UPS. Each outlet powers one rack-mounted power strips. One strip, mounted on the rear of the rack, has only one main power switch, and is for components that will always remain on. The other, mounted on the front of the rack, has a separate power switch for each outlet. All power is at the European standard 230 V, and the outlets accept Schuko plugs only. Each piece of equipment is powered by an IEC/C13-to-Schuko cord.



USB to Firewire
 0931 MXI
 DB-9
 4-pin Molex and DE-9
 Two 68-pin cables to each UMI Box
 27-pin Fischer to DC-25
 18-pin Molex and DA-15
 DD-50, DE-9, and 16-pin Molex
 Ethernet-to-DE-9
 Four DIN cables

Figure A.1: The major components of DCS electronic controls outside the cryostat. The connection from the Fischer box to the cold components is shown in [Figure A.16](#). The wiring inside the motion box is shown in [Figure A.33](#).

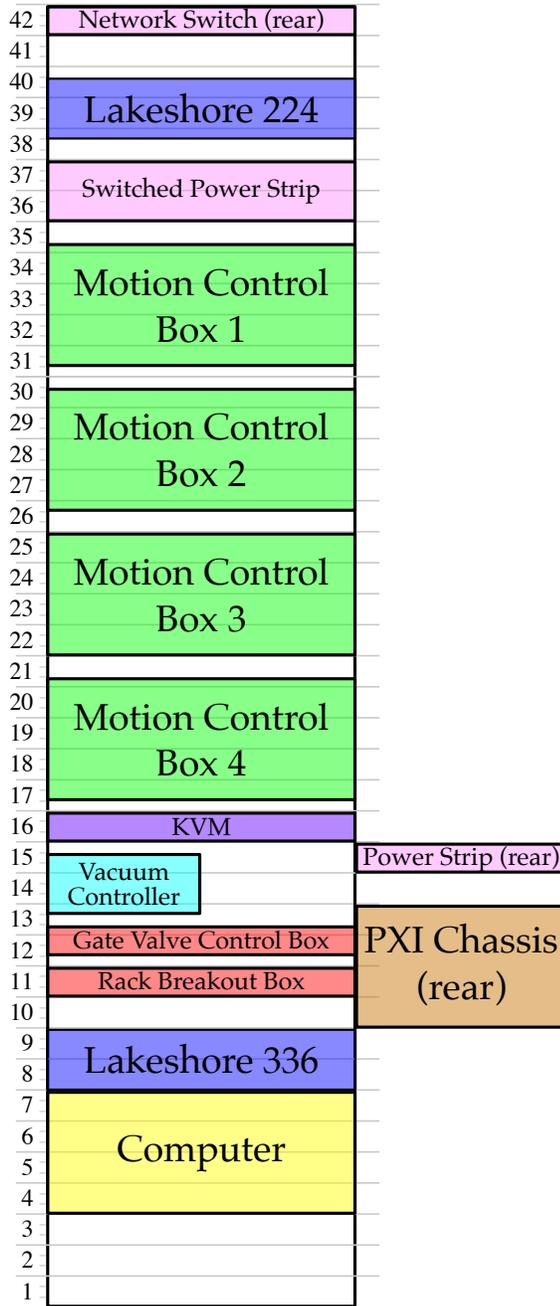


Figure A.2: The rack space required for all of the DCS components. All components are designed to be mounted in a 19" rack.

A.1.1 National Instruments PXI chassis

The National Instruments PXI Chassis (NI PXI-1033) connects to the computer with a MXI-Express/Express Card MXI Cable (NI 779500-01). The PXI chassis has four controller cards (NI PXI-7344), each of which connects to a UMI through 2 digital cables (NI SHC68-68-RDIO). These two cables are for Motion I/O and the custom-modified Digital I/O.

The Digital I/O cables are altered to initialize some subsystems in an active low state, so the Digital I/O and Motion I/O cables are not interchangeable. In order to have the Digital I/O started in active low, the DPull pin (pin 38) must be connected to ground (pin 37). The two wires are soldered together inside one of the connector housings and insulated with electrical tape.

The PXI Chassis AC input is 100–240 V, 50–60 Hz, and 4–2 Amps.

A.1.2 motion control and gate valve control boxes

The 4U aluminum motion control boxes serve as the housing for all of the motor control systems and power supplies. [Figure A.3](#) shows the box diagram layout and [Figure A.4](#) shows the back panel. The electricity is fed in through a standard panel-mount C14 plug ([Figure A.5](#)). AC transmission in the box is through 14 AWG wires, with black for line voltage, white for neutral, and green for ground. DC transmission in the box is through 18 AWG wires, with red for line voltage, brown for neutral, and green for ground. The peak power intake of the box is ~ 1.5 kW; this is ~ 12 A at 120 V or ~ 6.5 A at 230 V.

The UMI 7774 in each box uses a TDK-Lambda LS-100-24 (100 W, 24 V) power supply. This is connected to the V/V_{iso} and C/C_{iso} inputs. On the UMI itself, there are 3 DIP switches for each axis. For each, FAULT should be set to ACTIVE LOW; ENABLE and LIMIT LED should be set to ACTIVE HIGH.

Each axis of the UMI controls a different motor controller. Axes 1, 2, and 3 control

a drive spool motor, while axis 4 controls the corresponding linear actuator. All of the axes are connected to their own motor controller (NI P70530-SDN) through the CONTROL DA-15 on the UMI. This connection uses a specialty cable from National Instruments (NI 198141-2R5). The motor controllers are powered by their own TDK-Lambda SWS-300-48 (300 W, 48 V) power supplies.

Axes 1, 2, and 3 (the three motors on a single motion box) utilize the FEEDBACK DB-25 connector on the UMI. This connection receives information from the motor encoder and home switch. It also provides digital high voltage (+5 V) for the logic in the global stop and home switches. Inside the electronics box, the circuits branch out to the necessary positions via the PCB (Figure A.6). The PCB is connected to the UMI FEEDBACK connector with custom cables (Figure A.7). All motor signals (including the encoders), except for the linear actuator signal, go directly to a motion box.

The GLOBAL STOP DE-9 on the UMI receives information from all of the global stop switches in a motion box simultaneously. If any of the global stop switches are triggered, all motors on the UMI stop and the motors can no longer be moved. To defeat this, the motor shaft can either be turned manually on the motion box or a global stop defeater cable can be inserted in-line with the 50-pin DD-50 cable on the back of the motion control box. This special cable is a normal feedthrough cable except pin 33 is removed, which is the pin that connects to the INHIBIT ALL input on the UMI (see Figure A.6).

The ANALOG INPUT DE-9 on the UMI reads in the data from the 3 load cells on each motion box (via preamps) and the potentiometer from the linear actuator. The potentiometer signal comes via the rack breakout box (see Section A.1.3). The Axis 4 FEEDBACK DB-25 uses a custom logic circuit to read the state of the thermalizer (TM) (see Figure A.8). The TM Contact signal comes via the rack breakout box. When the TM is fully open or fully closed, the TM Contact wire is grounded to the

4-K cryostat plate (through a $75 - \Omega$ ribbon cable), and when it is between open and closed, it is an open switch. [Figure A.9](#) shows the custom cable that connects all of the analog signals and the TM Contact signal.

The DIGITAL I/O DB-25 on the UMI (not to be confused with the 68-pin VHDCI Digital I/O that connects to the PXI cards) is used to read in data from the proximity sensor and gate valves, as well as to control the power to the gate valves via 24 V control, 6 A non-latching relays (Phoenix Contact 2903361). There is a multi-part custom cable that connects to the in-box PCB, the Digital I/O port on the UMI, the Gate Valve feedthrough on the back panel of the box, both relays, and a power supply ([Figure A.10](#)). The power supply is a TDK-Lambda LS 50-12 (50 W, 12 V) supply that triggers the gate valves to open or close.

Each gate valve has two solenoids, one that opens the valve and one that closes it. When a current is applied to the “open” solenoid, it opens the gate valve, and when current is applied to the “close” solenoid, it closes the gate valve; when current is applied to neither, the gate valve remains in its current state. The gate valve control circuit is designed such that current flows through either solenoid only if the computer sends a positive control signal and, simultaneously, the operator holds down a momentary switch on the rack. This is designed to prevent the user accidentally opening or closing the gate valve when it is not safe to do so, and it prevents the computer from shutting the gate valve when the user is not present to give active confirmation. An overall schematic of the gate valve control wiring is shown in [Figure A.11](#). The momentary switches (one per gate valve) are on the front of the gate valve control box, a small plastic box mounted in the rack.

The load cell preamps are powered by a TDK-Lambda LS 25-24 (25 W, 24 V) power supply. The motor power from the controller for axes 1, 2, and 3 is sent to the motion box as part of a 16-pin Molex cable; 4 wires for each of 3 motors, 2 wires for the preamp power, and a ground wire (see [Figure A.12](#)). The power from the

controller for axis 4 goes to the rack breakout box, where it is combined with the linear actuator power signals from other motion control boxes into a 18-pin Molex cable.

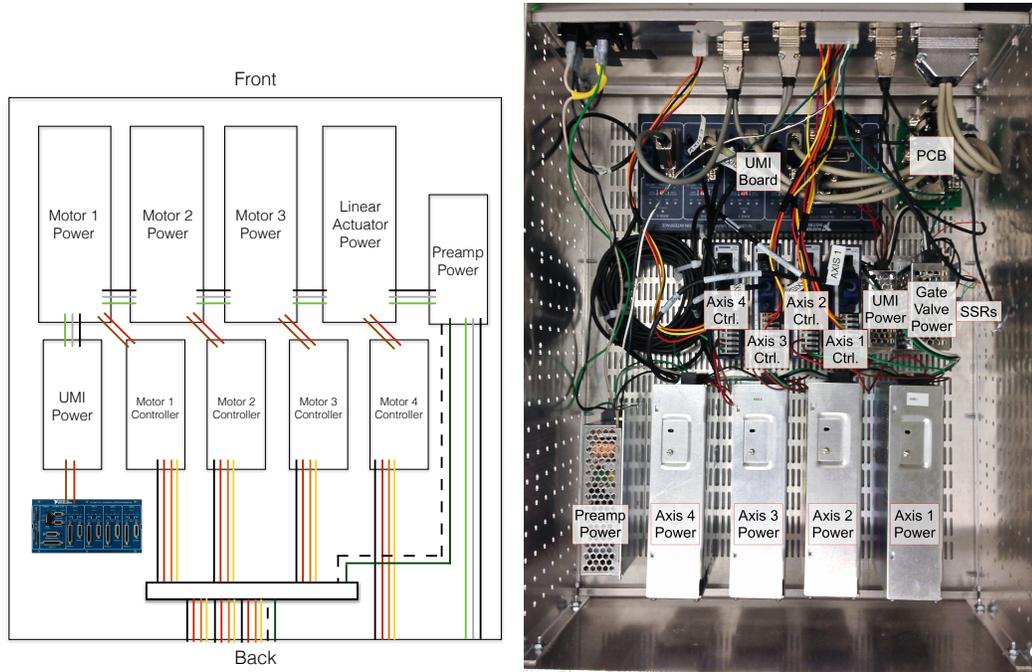


Figure A.3: The layout of the UMI electronics control box, showing all of the 14 AWG power wiring.

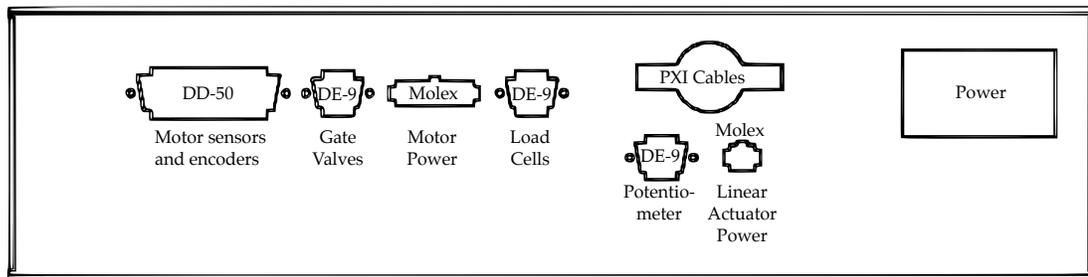


Figure A.4: The back panel of the motion control box.

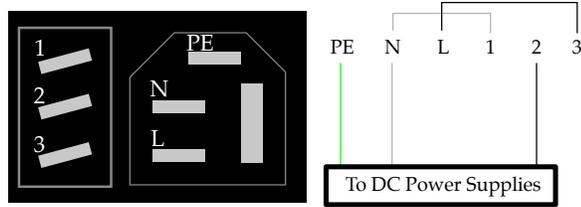


Figure A.5: The C14 inlet diagram (back view). It is a single-fused non-filtered inlet with a SPST lighted switch. The green PE (protective earth) wire connects directly to the DC power supplies. The white N (neutral) wire going to the DC supplies is split to go to connector 1 of the switch as well. The light on the switch is powered by connectors 1 and 2. When the switch is turned on, connectors 2 and 3 are connected and power flows to the DC power supplies.

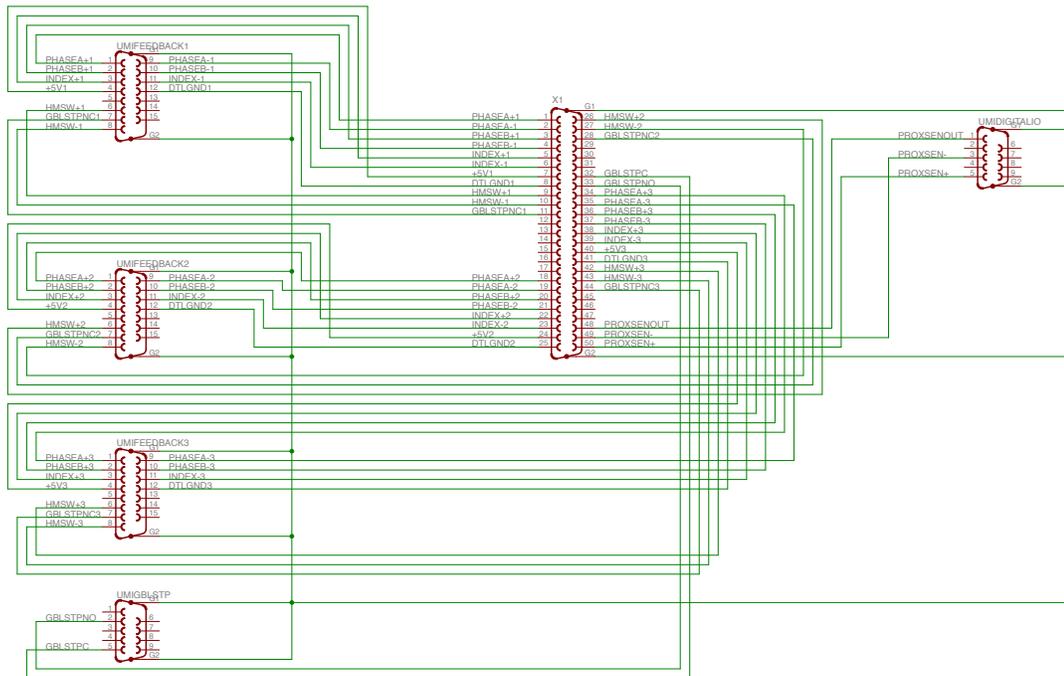


Figure A.6: The schematic of the PCB that connects all the encoders, sensors, and safety controls for the drives spools in a single motion box.

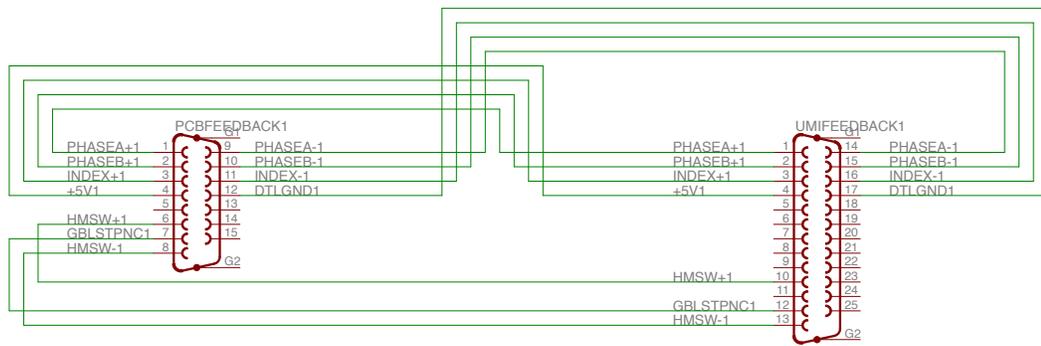


Figure A.7: The custom cable that connects the UMI FEEDBACK DB-25 port to the PCB DA-15 port.

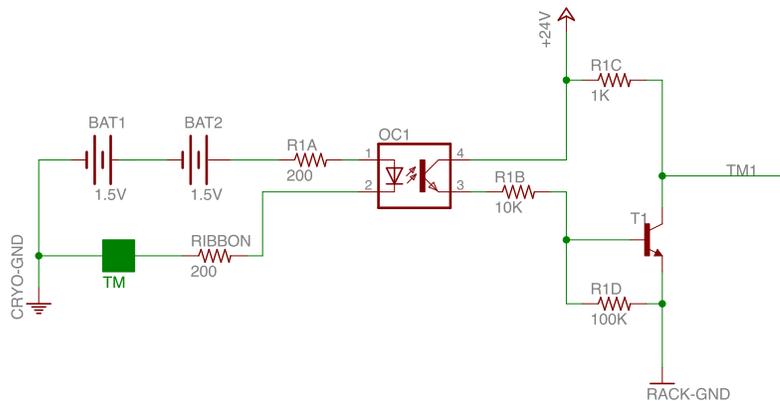


Figure A.8: A schematic of the TM Contact circuit for a single gate valve. The part of the TM Contact circuit that goes through the cryostat is optically isolated from the other electronics in the rack.

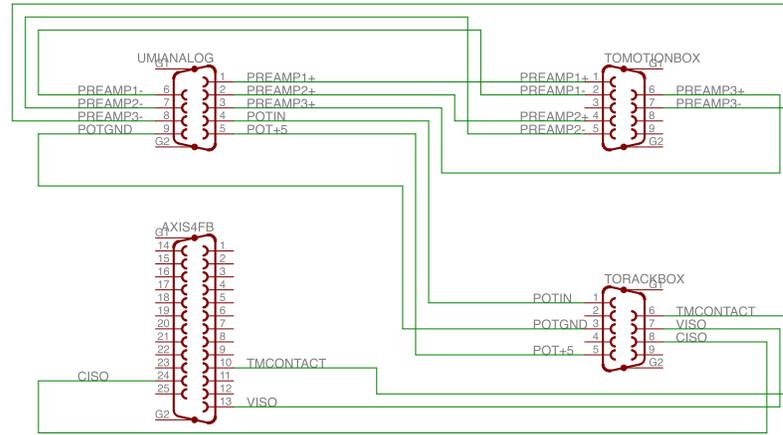


Figure A.9: The custom cable that connects the UMI ANALOG DE-9 port and Axis 4 FEEDBACK DB-25 port to two motion control box feedthroughs.

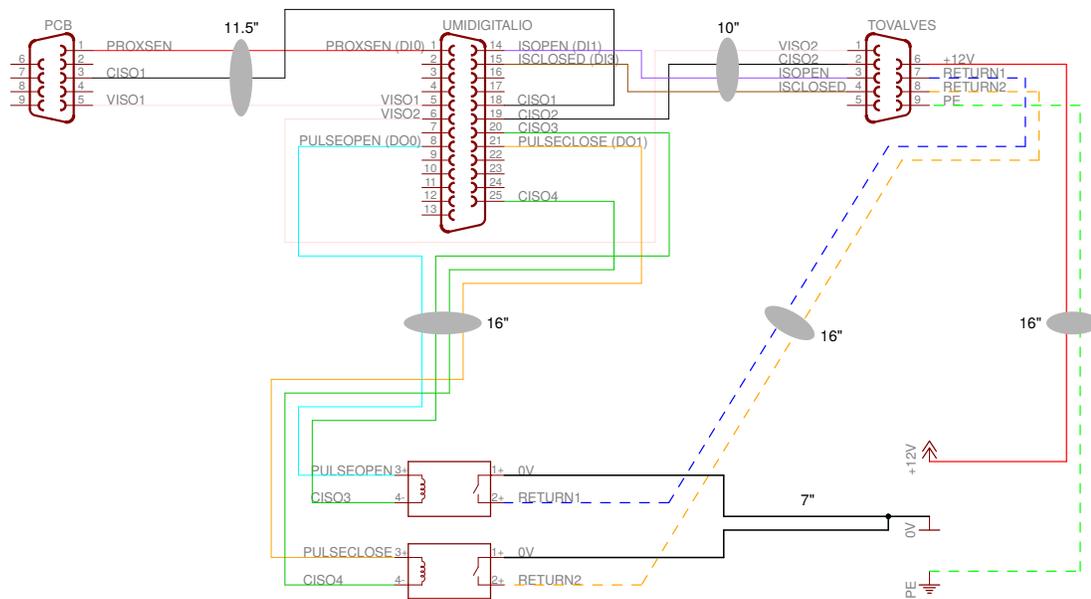


Figure A.10: The custom cable that contains the signals and power for the proximity sensor and gate valve. The gray ovals indicates wires that are bunched together in a cable. The bare wires are 24 AWG. Of the Digital I/O pins that are used, pin 1 is Digital Input 0; pins 5 and 6 are V_{iso} (24 V); pin 8 is Digital Output 0; pin 14 is Digital Input 1; pin 15 is Digital Input 3; pins 18, 19, 20, and 25 are C_{iso} (ground); and pin 21 is Digital Output 1. A positive pulse in PULSEOPEN will open the gate valve and a positive pulse in PULSECLOSE will close it.

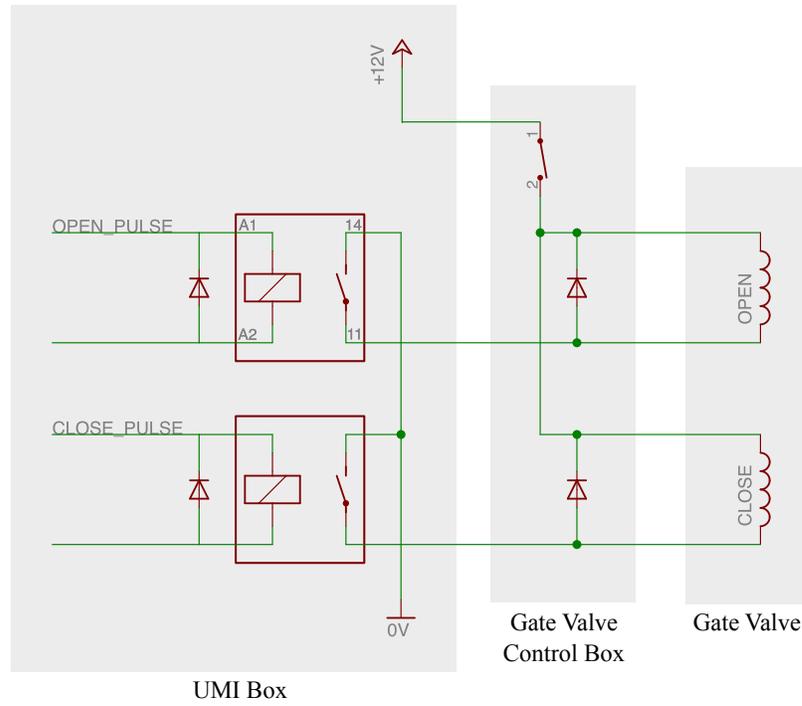


Figure A.11: A schematic of the gate valve control circuit for a single gate valve.

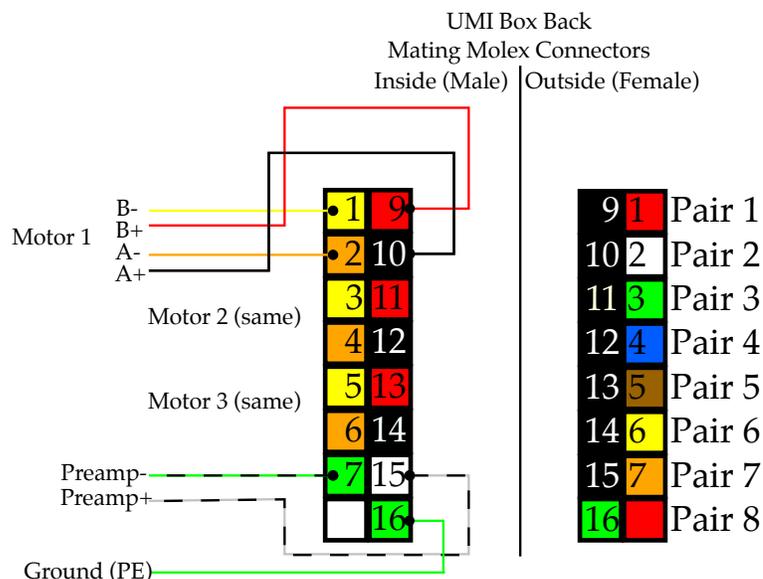


Figure A.12: The Molex connector from the motor controllers to the 16-pin connector, which carries 15 wires to the motion box. The colors on the right are the wire colors in the twisted pair cable going to the motion box; the colors on the left are the colors of the cable going to the motor controllers and preamp power supply. On the motion box, the connections are the same from the twisted pair wire to a female Molex receptacle, and from the male Molex connector to the motors and preamps. All wires are 18 AWG.

A.1.3 Rack breakout box and temperature measurements

All cables that go to the cryostat in general but not to a specific motion box on the cryostat pass through the rack breakout box, with the exception of the 600 mK Cernox thermometers. It handles the thermalizer temperature signals, thermalizer contact signals, linear actuator potentiometer readings, and linear actuator power. The back panel of the breakout box is in [Figure A.13](#), and the non-power wiring inside the rack breakout box is in [Figure A.14](#). The linear actuator power Molex color code inside the breakout box is shown in [Figure A.15](#).

The LakeShore Model 336 Temperature Controller measures the temperature on the four thermalizer sliding blocks. It connects to the computer with a USB cable. The major components of the thermalizer temperature sensor system are in

Figure A.16. The LakeShore 336 connects to 4 temperature sensors using DIN 45322 connections to the back of the unit, labeled Channel A, B, C, and D (Figure A.17). These are connected with straight-through four-wire cables to the rack breakout box, which is in turn connected to the sensors in vacuum with a Fischer27 (S 105 A102-130+) connector.

The LakeShore Model 224 Temperature Monitor connects to the computer similarly to the LakeShore 336. There are 12 channels on this temperature monitor, each of which monitors the temperature of a separate 600 mK stainless steel tube in the cryostat, below the thermalizers. Because the temperature signals will be separated into two feedthroughs when they enter the cryostat, they are sent in two different cables, each of which has a breakout cable that has a DB-25 connector on one end and 6 DIN 45322 connectors on the other (Figure A.18).

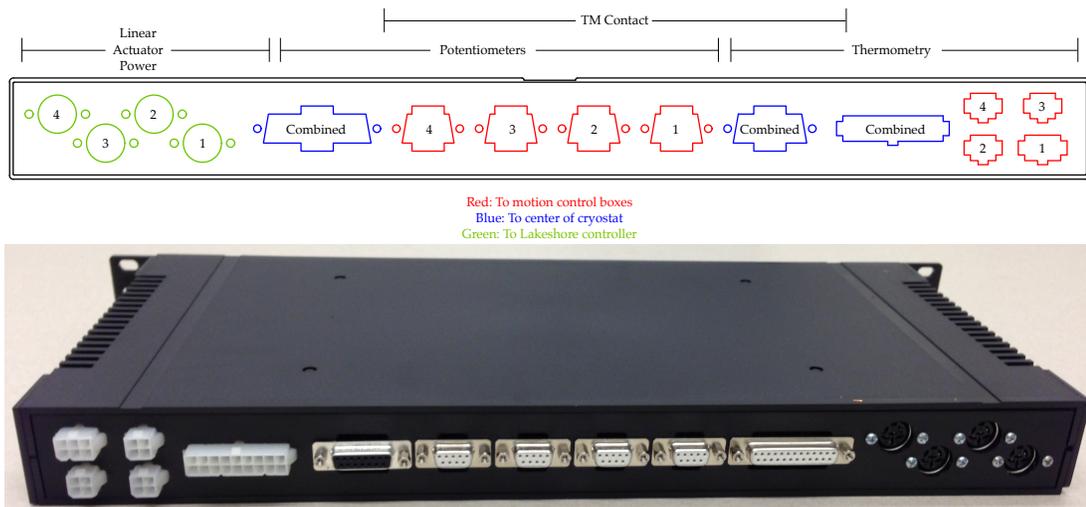


Figure A.13: The back panel of the rack breakout box.

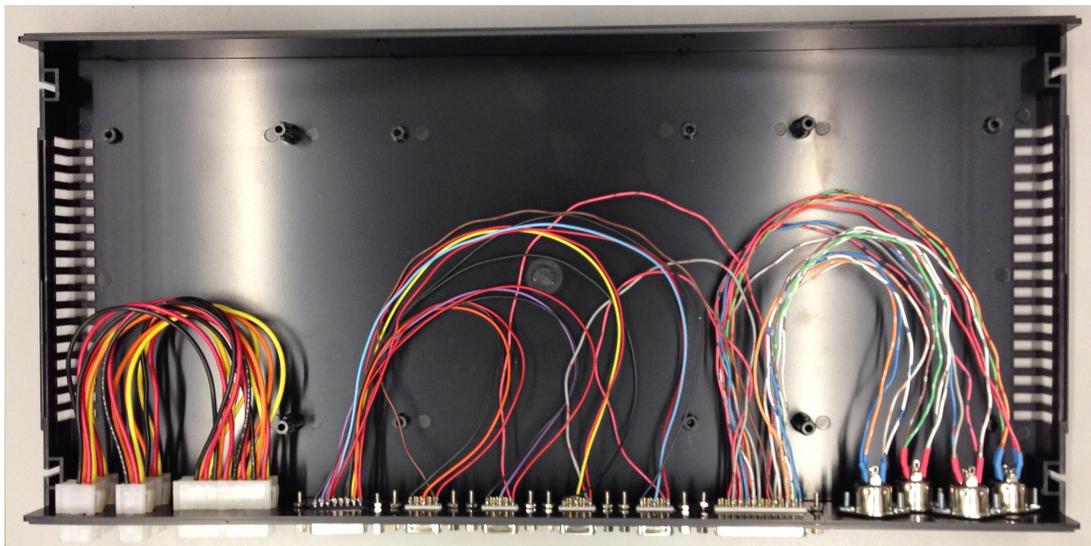
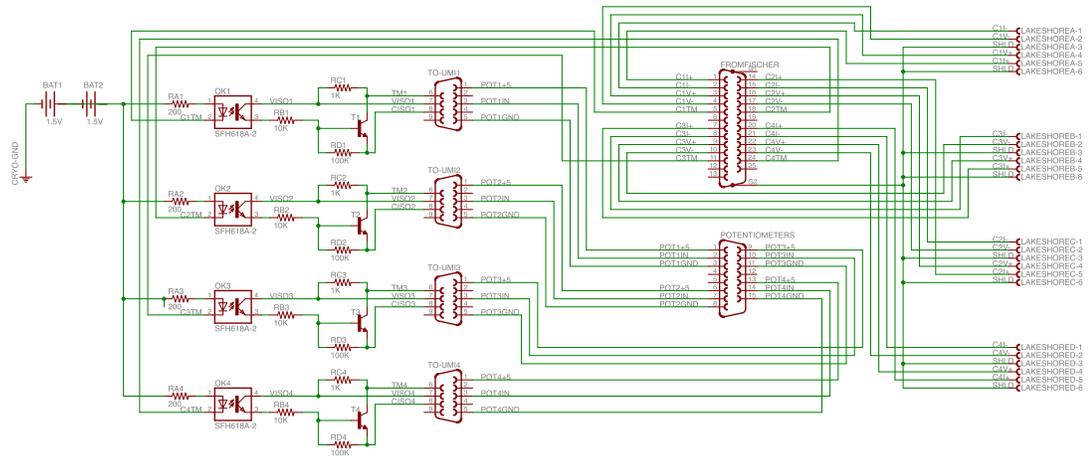


Figure A.14: The wiring inside the rack breakout box. All the feedthroughs are on the back plate of the box, while all the connections are wire soldered to the feedthroughs inside the boxes.

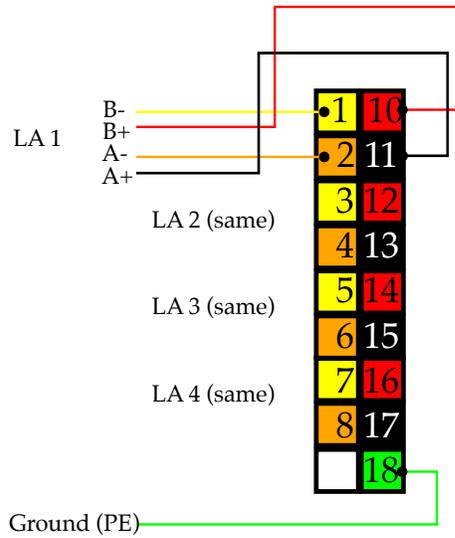


Figure A.15: The Molex connector and color code for all of the linear actuator cables. All wires are 18 AWG.

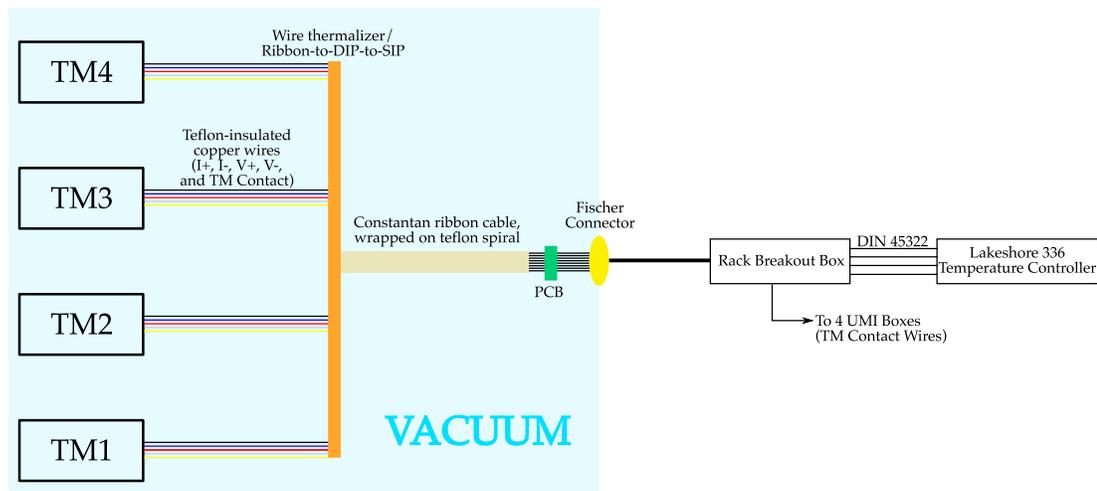


Figure A.16: An overview diagram of the thermalizer temperature sensor connections.

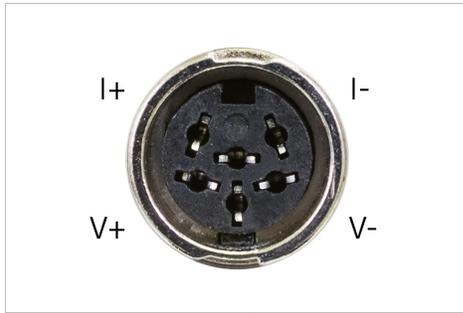


FIGURE 3-3 Sensor input connector

Pin	Symbol	Description
1	I-	-Current
2	V-	-Voltage
3	None	Shield
4	V+	+Voltage
5	I+	+Current
6	None	Shield

Figure A.17: The pinout for the DIN connectors on the back of the LakeShore temperature monitors (from the LakeShore Model 336 manual).

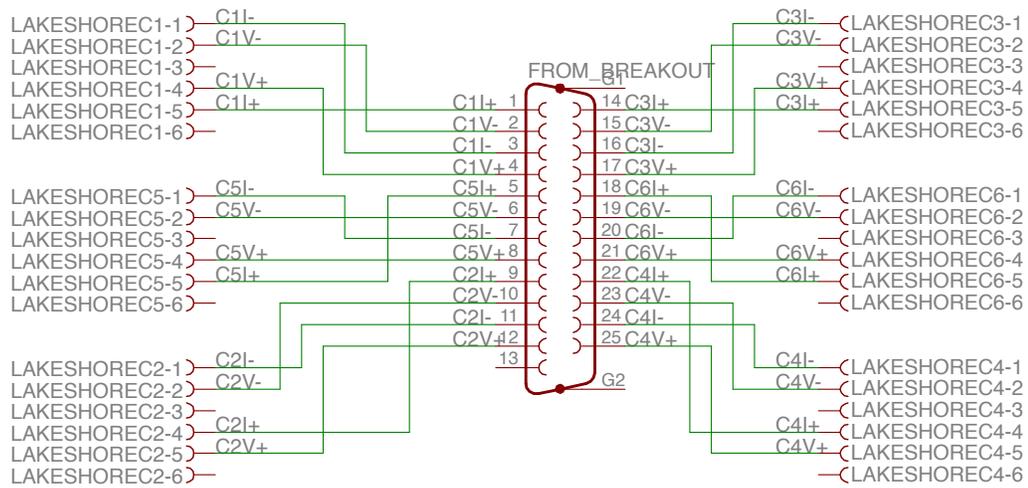


Figure A.18: Wiring for each of the two DB-25 to 6 DIN 45322 breakout cables.

A.1.4 Vacuum gauge and controller

The Varian Vacuum Gauge Controller (Agilent XGS-600) reads in all of the vacuum pressure data from the two pressure gauges attached to the motion boxes. It connects

ID	Qty.	Rack Loc.	Cryo. Loc.	Wires	Diameter	Max V, power
A	4	Control box	Motion box	50 @ 24 AWG	12.8 mm	5 V, low
B	4	Control box	MB Preamps	6 @ 24 AWG	5.9 mm	10 V, low
C	4	Control box	Motion box	18 @ 18 AWG	16.5 mm	24 V, 150 W
D	1	Control box	MB Vac. Gauge	8 @ 24 AWG	5.2 mm	30 V, low
E	4	Control box	Gate Valve	8 @ 24 AWG	5.2 mm	30 V, low
F	1	Breakout box	Fischer box	20 @ 24 AWG	9.5 mm	low, low
G	1	Breakout box	Linear Acts.	18 @ 18 AWG	16.5 mm	24 V, 150 W
H	1	Breakout box	Linear Acts.	25 @ 24 AWG	10.6 mm	5 V, low
J	2	Breakout box	IP Cameras	8 @ 24 AWG	5.2 mm	5 V, low
K	2	Lakeshore 224	Fischer box	24 @ 24 AWG	10.1 mm	low, low

Table A.1: Wiring from the rack to the cryostat.

the computer through a DE-9 cable. The male end plugs into the SERIAL COMM port on the Varian Pressure Gauge Controller, and the female end connects to the serial port on the computer.

The gauges themselves, Agilent Vacuum Gauge Full Range Pirani/Cold Cathode (FRG700CF35) gauges, are connected to special Ethernet to DE-9M cables (Agilent 656458203). A standard ethernet cord can be used to extend the length of this cord if necessary. There are 4 positions to read in a vacuum gauge on the Vacuum Gauge Controller, so up to four gauges can be used.

A.2 Rack-to-cryostat wiring

The wiring from the electronics rack to the cryostat is shown in [Table A.1](#). The connectors on these cables, and their sizes, are shown in [Table A.2](#). Each cable to the motion boxes is given a letter and a number; the letter identifies the type of cable, and the number the corresponding motion box. For wiring color schemes, see [Figure A.19](#) through [Figure A.26](#). Cables D and J are standard ethernet cables. All cables to the motion boxes (20 cables) and their connectors are threaded through a single rectangular hole in the Faraday cage.

Cables	Cryostat End	Rack End
A1 – A4	DD-50 M (23.1 mm × 66.8 mm)	Same
B1 – B4	DE-9 M (15.6 mm × 31.0 mm)	Same
C1 – C4	16-pin F Molex Mini-Fit Jr. (14.8 mm × 34.8 mm)	Same
D1 – D4	8P8C/RJ45 (11.5 mm × 11.8 mm)	Same
E1 – E4	DE-9 M (11.5 mm × 11.8 mm)	Same
F	Fischer S 105 a102-130+ (26.9 mm diameter)	DB-25 M (15.6 mm × 53.4 mm)
G	18-pin F Molex Mini-Fit Jr. (14.8 mm × 39.0 mm)	4 × 4-pin F Molex Mini-Fit Jr.
H	DA-15 M (15.5 mm × 39.7 mm)	DA-15 F
J1, J2	8P8C/RJ45 (11.5 mm × 11.8 mm)	Same
K1, K2	Fischer S 105 a102-130+ (26.9 mm diameter)	DB-25 M (15.6 mm × 53.4 mm)

Table A.2: Connectors on the rack-to-cryostat cables.

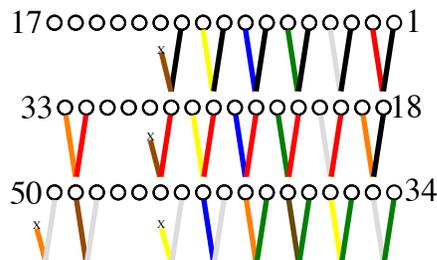


Figure A.19: Cable A coloring scheme. Twisted pairs are shown together.

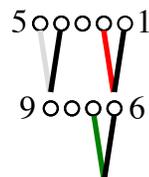


Figure A.20: Cable B coloring scheme. Twisted pairs are shown together.

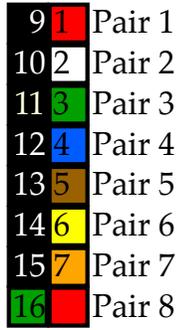


Figure A.21: Cable C coloring scheme. Twisted pairs are shown next to each other horizontally.

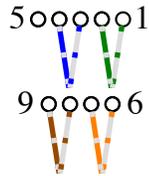


Figure A.22: Cable E coloring scheme. Twisted pairs are shown together.

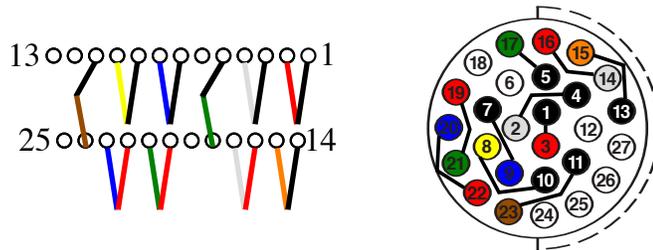


Figure A.23: Cable F coloring scheme. Twisted pairs are shown together in the DB-25 connector on the left, and are shown linked by a line in the Fischer connector on the right.

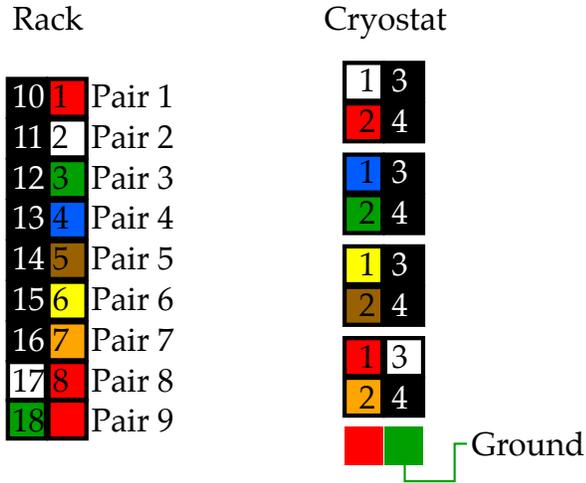


Figure A.24: Cable G coloring scheme. Twisted pairs are shown next to each other horizontally.

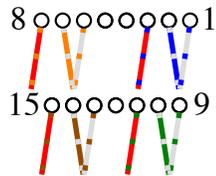


Figure A.25: Cable H coloring scheme. Twisted pairs are shown next to each other horizontally.

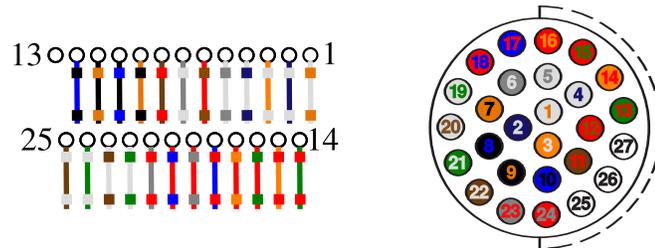


Figure A.26: Cable K coloring scheme for Cernoxes 7-12. Striped wires are shown on the Fischer connector with the pin number as the stripe color.

A.3 On the cryostat

A.3.1 Motion box

Each motion box contains 3 drive spools, a proximity sensor, and a gate valve, and 2 motion boxes contain pressure gauges. Coming to each motion box from its corresponding motion control box is a DD-50 cable, DE-9 cable, and 14-pin Molex power for the motors and preamps. The top of a motion box, where the cables arrive at the motion box from the rack, is shown in [Figure A.27](#).

The motor and preamp power is split from the 16-pin Molex to go the appropriate connectors: a 4-pin Molex to each motor and 2 pins shared and split to the three preamps. The ground (PE) wire (pin 16) is split into three wires and attached with a #20 hose clamp to the motors (one per motor) as shown in [Figure A.28](#).

The DE-9 cable is split into three cables, each carrying two signals from the preamp output: positive output and common.

The DD-50 cable is plugged into a breakout PCB, which connects to the three encoders, three motion box feedthroughs (containing the home switch and global stop signals), and the proximity sensor ([Figure A.29](#)). The encoders are connected with US Digital CA-FC10-SH-NC-5 cables soldered to DB-9 connectors on the PCB side. The pinout of the encoder cables is given in [Figure A.30](#).

Each sensor port on the PCB is connected to a motion box feedthrough and to a preamp (Honeywell 060-6827-04) with a custom cable ([Figure A.31](#)). Each preamp is calibrated to a specific load cell, identified by the drive spool number. The excitation voltage of the load cell is 5 volts, indicated by placing the jumpers in the preamp in the middle position. The output of the load cell is nominally 2 mV/V, so the DIP switches on the preamp are set with switch 1 on and all others off. The fine gain of the preamp and the zero point is adjusted with a screwdriver so that the full range of output is within the acceptable ADC range. The output/power side of the preamp

wiring is shown in [Figure A.32](#). The motion box wiring inside the feedthrough is shown in [Figure A.33](#).

The proximity sensor (Proxitron IKVS-010.23-G-S4) is connected to the motion box as shown in [Figure A.34](#), and the output is connected the PCB with a custom cable ([Figure A.35](#)).

The gate valve (VAT Valves 12136-PA44-X 2 1/2) is connected to the DB-9 connector coming from the gate valve control box ([Figure A.36](#)).



Figure A.27: A photo of the motion box PCB enclosure and wiring.

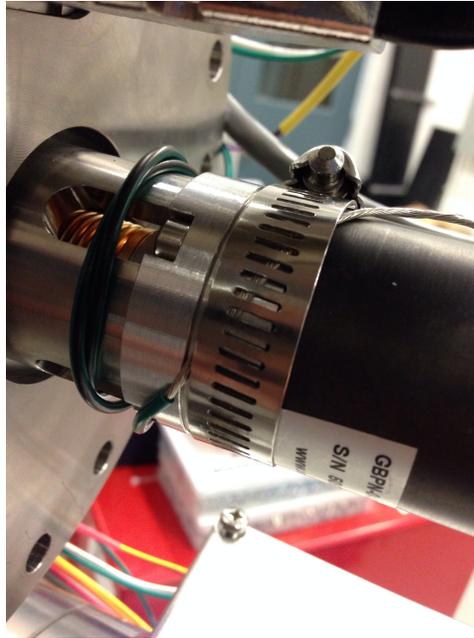


Figure A.28: The motor grounding scheme, with the ground wire stripped back several centimeters and attached to the motor body.

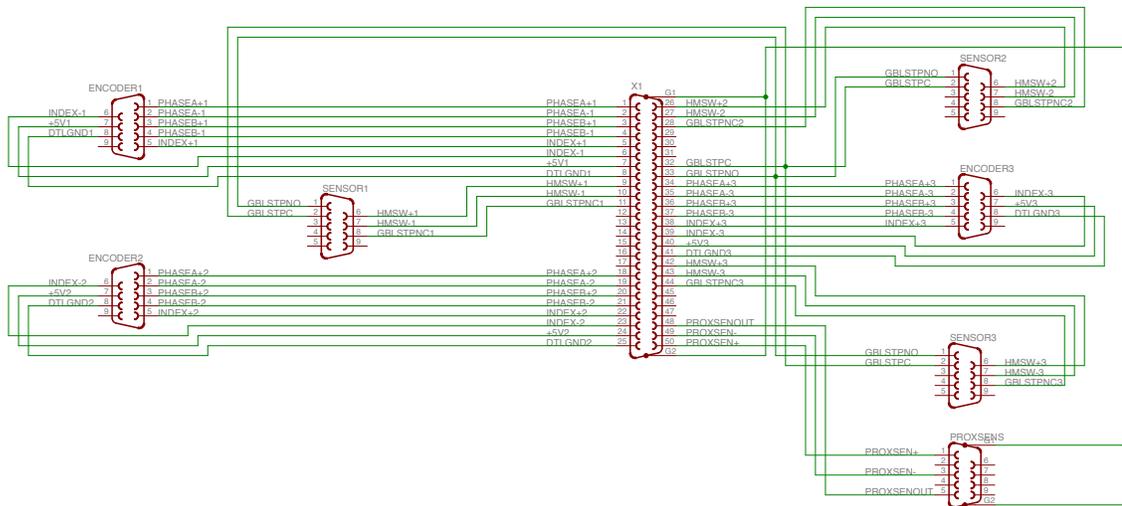
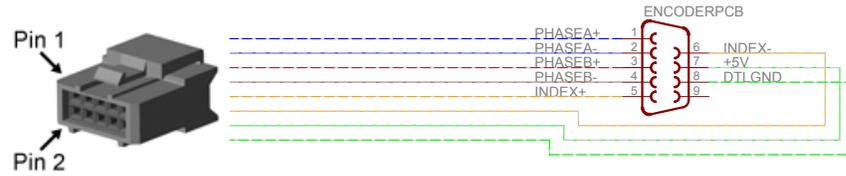


Figure A.29: The schematic of the PCB breakout on each motion box.

Pin-out



Pin	Description	Color
1	-	-
2	Ground	Green / White Stripe
3	I- channel	White / Orange Stripe
4	I+ channel	Orange / White Stripe
5	A- channel	White / Blue Stripe
6	A+ channel	Blue / White Stripe
7	Power	White / Green Stripe
8	-	-
9	B- channel	White / Brown Stripe
10	B+ channel	Brown / White Stripe

Figure A.30: The pinout of the encoder cable.

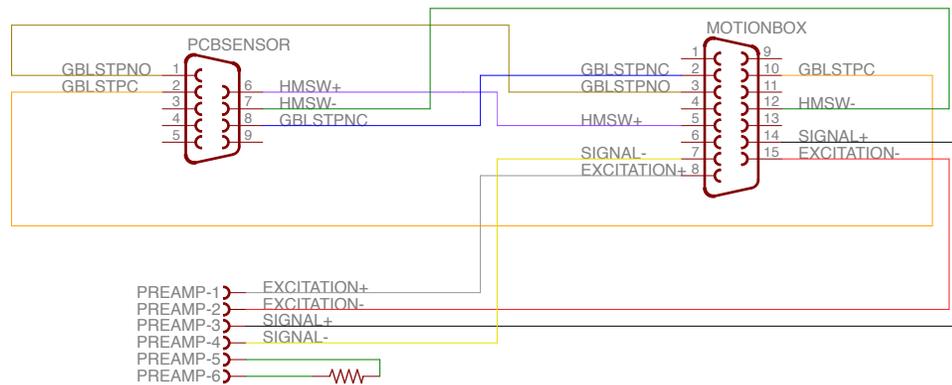


Figure A.31: The custom cable connecting the motion box breakout PCB to the motion box and to the preamp.

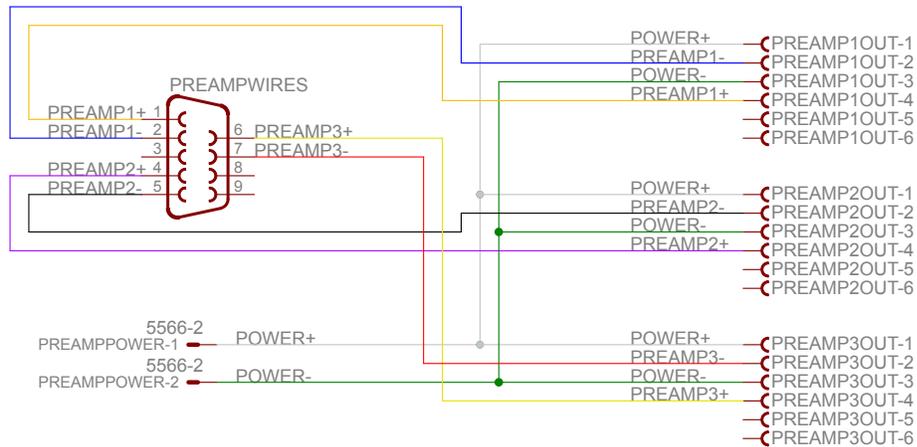


Figure A.32: The wire color code for the connection from the DE-9 cable from the motion control box to the preamps. The power wires are AWG 20 and the signal wires are AWG 24.

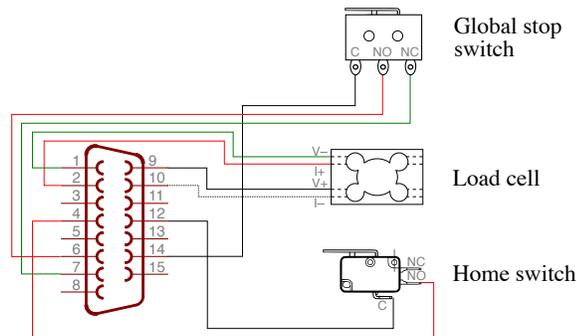


Figure A.33: The in-vacuum connections for each drive spool. The female connector here is a mirror image of the female connector in [Figure A.31](#) because of the male-to-male feedthrough welded into the motion box.

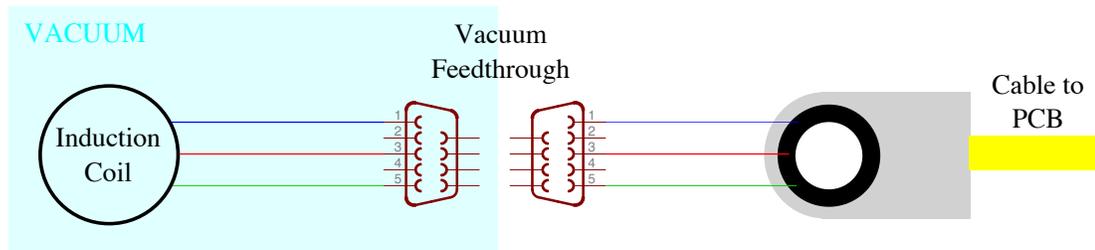


Figure A.34: The connection from the motion box vacuum to the proximity sensor body.

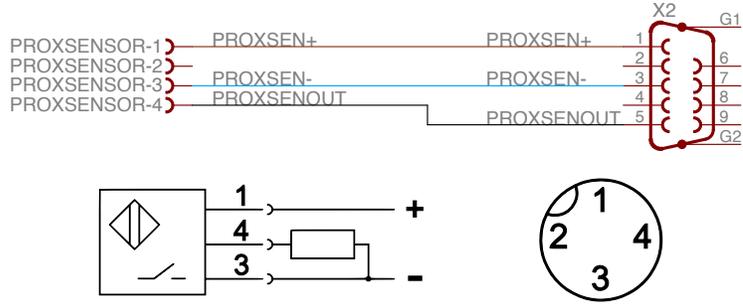


Figure A.35: The custom cable connecting the motion box breakout PCB to the output of the proximity sensor.

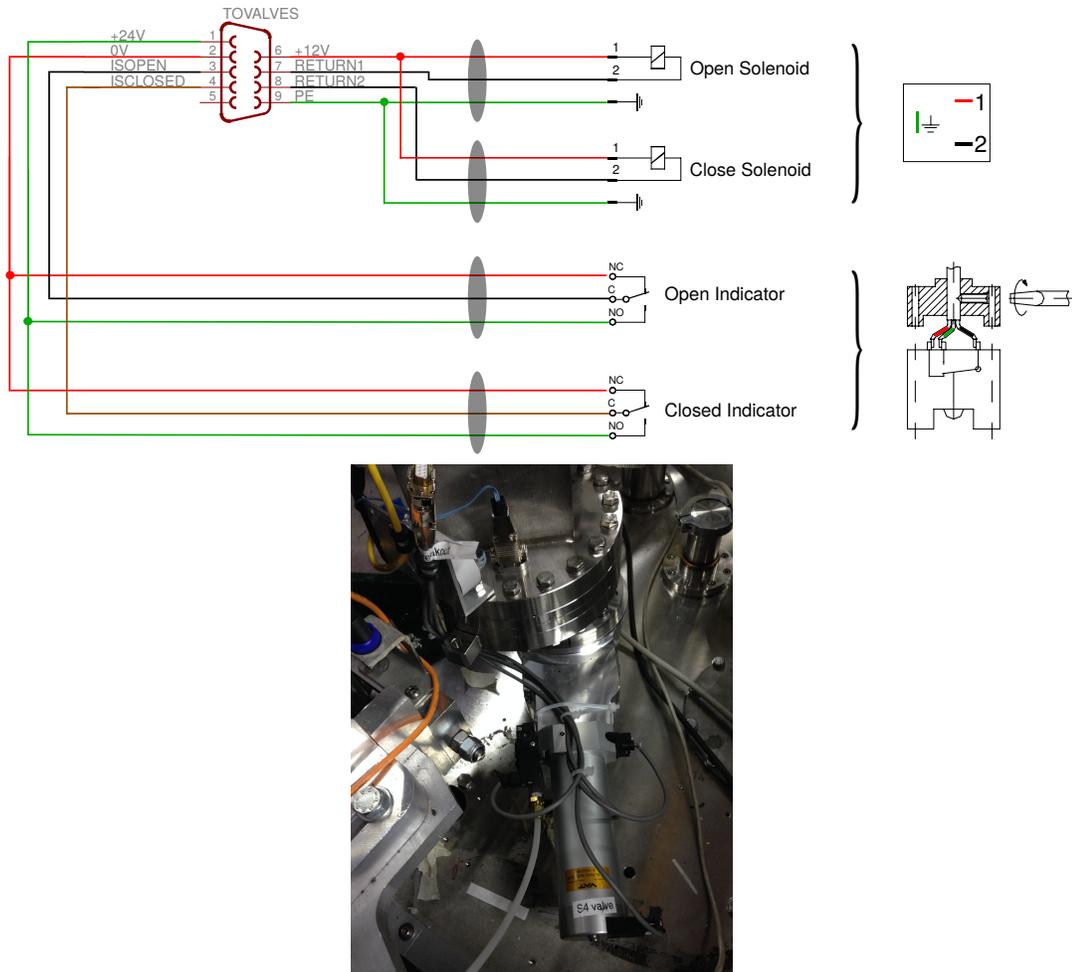


Figure A.36: The cable coming from the motion control box breaks out into the four cables going to the gate valve.

A.3.2 Temperature measurement

Temperature measurements come through vacuum Fischer feedthroughs (DBPE 105 A102-130). There are three Fischer connectors, one for the four thermalizer thermometers (including the thermalizer contact signal) and two for the twelve 600-mK thermometers. The Fischer connection wiring diagram for the thermalizers is in [Figure A.37](#), and for the 600-mK thermometers in [Figure A.38](#). The in-vacuum side of each Fischer feedthrough is soldered to a PCB. The in-vacuum PCB wiring is in [Figure A.39](#). The PCB connects to twisted-pair constantan ribbon cables that run down the Teflon spiral and into the cryostat. At the bottom of the spiral, there are DIP connectors connected to the ribbon cable, which mate with SIP connectors on wires that go to the thermalizers and stainless steel tubes below the thermalizers ([Figure A.40](#)).

Mounted to each of the thermalizers is a CX-1010-SD Cernox thermometer calibrated to 1.4 K. Mounted to each of the stainless steel tubes below the thermalizers are CX-1030-SD Cernox thermometers. Ten of these tubes have thermometers were factory calibrated to 1.4 K, while the others are calibrated to 300 mK in order to allow us to roughly calibrate the other thermometers to that range.

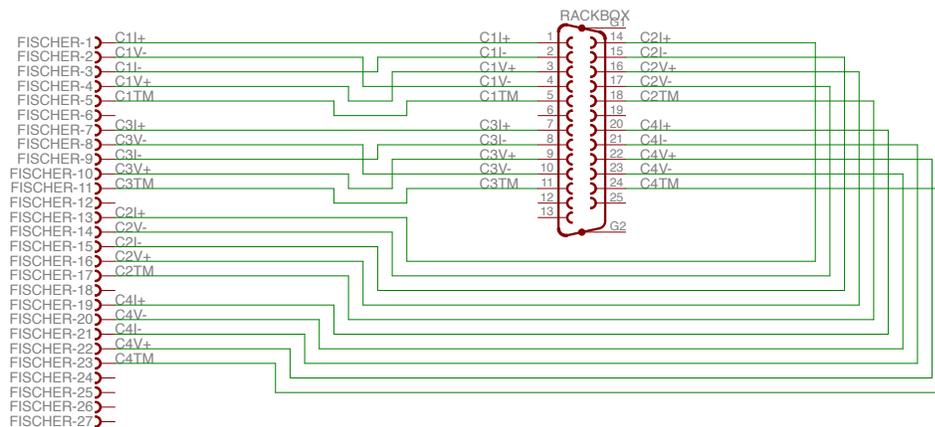


Figure A.37: The outside-vacuum part of the Fischer vacuum feedthrough for the thermalizer thermometers. For the color coding of this cable, see [Figure A.23](#).

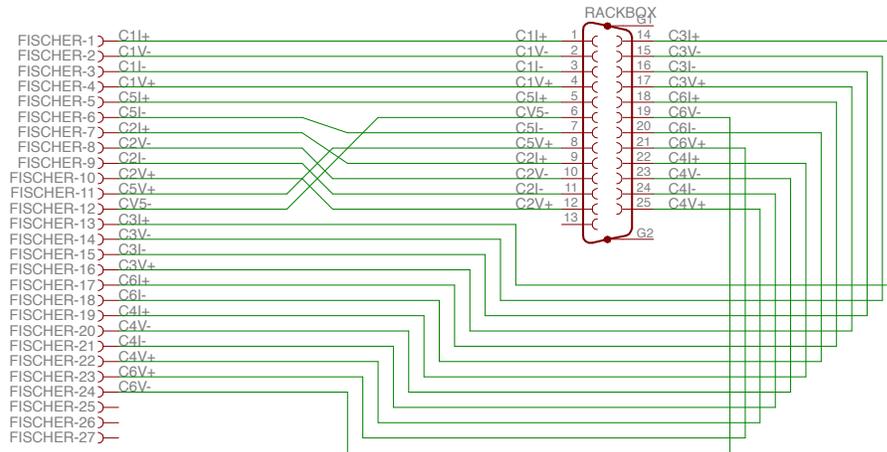


Figure A.38: The outside-vacuum part of the Fischer vacuum feedthrough for the 600-mK thermometers.

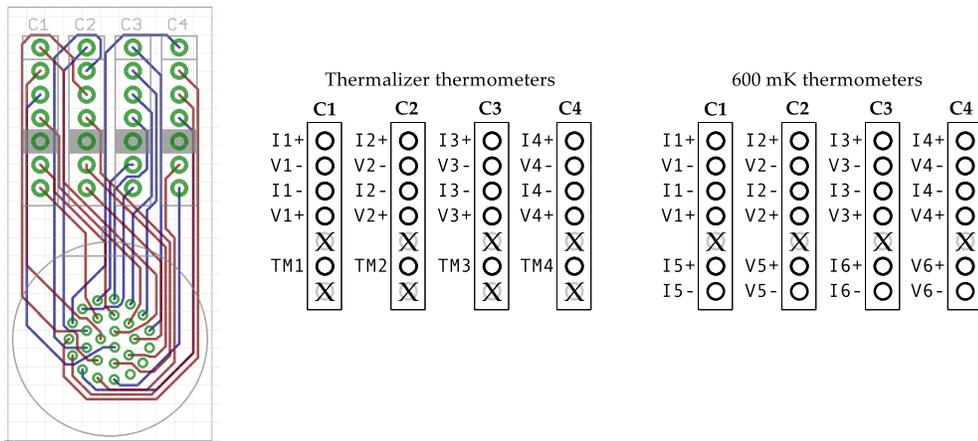


Figure A.39: The pinout of the PCB that connects to the Fischer vacuum feedthrough inside the vacuum.

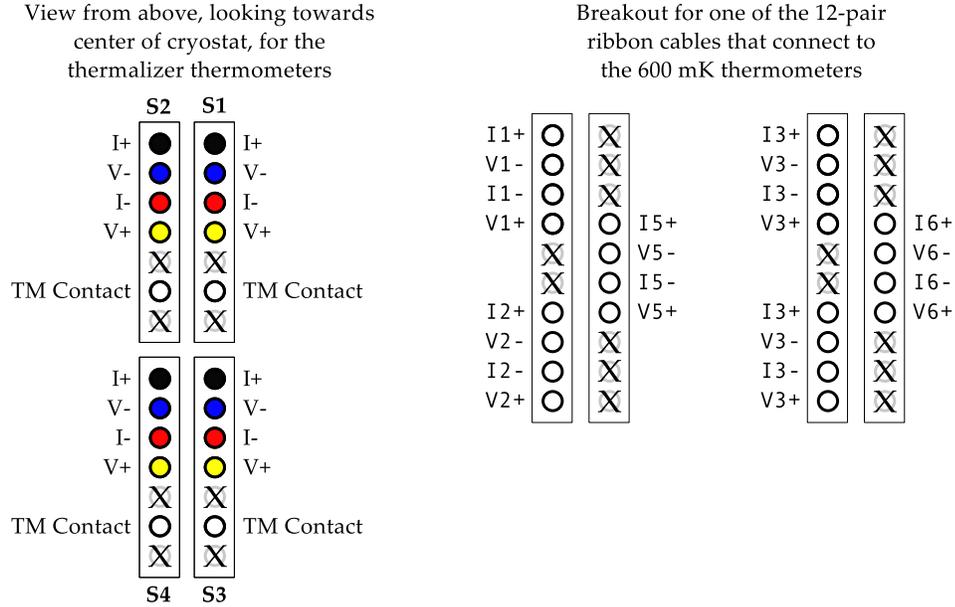


Figure A.40: The details of the breakouts at the bottom of the teflon spiral.

A.3.3 Linear actuators

Each of the four linear actuator has a potentiometer (DE-9) and power (4-pin Molex) connector.

The potentiometer signals for the four linear actuators travel together in one cable from the rack breakout box, which has a female DE-9 connector on the cryostat end. There is a one-to-four cable custom breakout (1 male DE-9 to 4 male DE-9 connectors), the wiring for which is shown in [Figure A.41](#). There are then four DE-9 F to M cables from the breakout to the potentiometers, which have DE-9 F connectors.

Power for the linear actuators is carried to the cryostat in a single 18-wire cable. This cable has an 18-pin F Molex connector that mates to the rack breakout box on one end, and four 4-pin M Molex connectors on the other end, plus a ground wire that can be connected to the linear actuators if needed. Each 4-pin Molex connector provides power to the linear actuators, as shown in [Figure A.42](#).

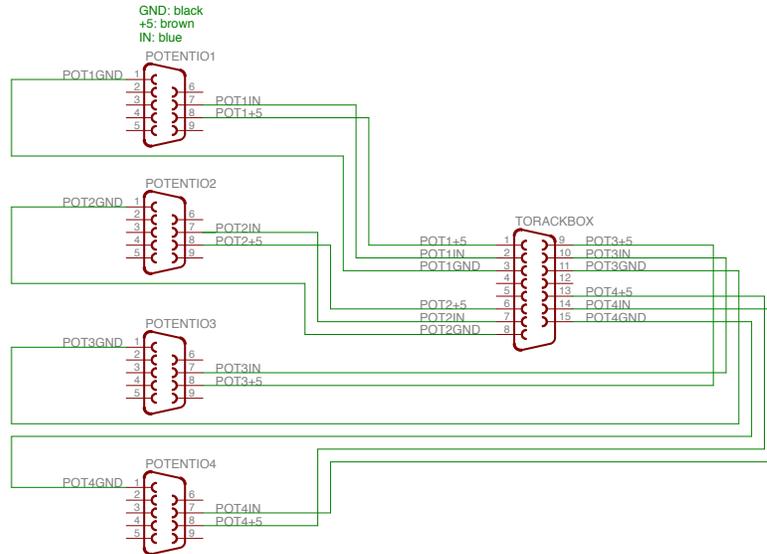


Figure A.41: The wiring for the 1 male DE-9 to 4 male DE-9 custom breakout.

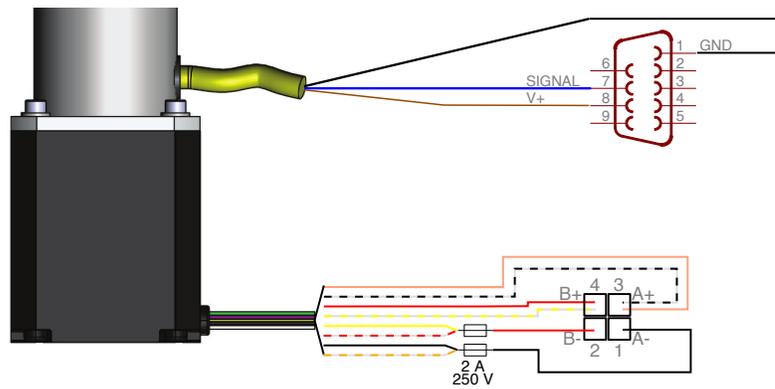


Figure A.42: The linear actuator gets power through a 4-pin Molex connector and outputs potentiometer signals in a DB-9 connector. 2 A, 250 V fuses are in place to protect the linear actuators.