

A STUDY OF MAJORANA DEMONSTRATOR BACKGROUNDS WITH BAYESIAN  
STATISTICAL MODELING

Christopher R. Haufe

A dissertation submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Physics and Astronomy.

Chapel Hill  
2023

Approved by:

John F. Wilkerson

Reyco Henning

Matthew Green

Amy Nicholson

Fabian Heitsch

© 2023  
Christopher R. Haufe  
ALL RIGHTS RESERVED

## ABSTRACT

Christopher R. Haufe: A Study of MAJORANA DEMONSTRATOR Backgrounds with Bayesian Statistical Modeling  
(Under the direction of John F. Wilkerson)

The MAJORANA DEMONSTRATOR was a neutrinoless double beta decay ( $0\nu\beta\beta$ ) experiment consisting of around 30 kg of germanium detectors enriched to 88% in  $^{76}\text{Ge}$  and around 14 kg of natural germanium detectors. The detectors were divided between two cryostats and surrounded by a graded passive shield. The DEMONSTRATOR concluded in March 2021 and set a  $0\nu\beta\beta$  half-life limit of  $T_{1/2} = 8.3 \times 10^{25}$  years (90% C.L.) based on its full exposure of 64.5 kg·yr. The experiment achieved one of the lowest background rates in the region of the  $0\nu\beta\beta$  Q-value,  $(6.23_{-0.52}^{+0.55}) \times 10^{-3}$  cts/(keV·kg·yr). This background rate, however, was higher than the rate of  $(1.16 \pm 0.04) \times 10^{-3}$  cts/(keV·kg·yr) projected by material assays and simulations. This discrepancy arises from an excess of events from the  $^{232}\text{Th}$  decay chain. Background model fits presented in this work aim to understand the observed  $^{232}\text{Th}$  excess and other deviations from assay-based projections, as well as allow a precision measurement of the  $2\nu\beta\beta$  half-life. Comparisons of the data with simulations indicate the  $^{232}\text{Th}$  excess cannot arise from near-detector components, and likely originates from a component of the cryogenic system. This is an important finding related to the design and implementation of the LEGEND-200 experiment. The final results of the DEMONSTRATOR are presented along with this study of its backgrounds using Bayesian statistical inference methods. The efficacy of these inference methods is limited by the statistics of the DEMONSTRATOR dataset. However, these methods remain useful for supplementary background studies, and support the analysis to determine the location of the  $^{232}\text{Th}$  excess.

## **ACKNOWLEDGEMENTS**

This material is supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, the Particle Astrophysics and Nuclear Physics Programs of the National Science Foundation, and the Sanford Underground Research Facility.

## TABLE OF CONTENTS

<b>LIST OF TABLES</b> . . . . .	<b>ix</b>
<b>LIST OF FIGURES</b> . . . . .	<b>xi</b>
<b>LIST OF ABBREVIATIONS</b> . . . . .	<b>xvii</b>
<b>CHAPTER 1: Neutrinoless Double Beta Decay (<math>0\nu\beta\beta</math>)</b> . . . . .	<b>1</b>
<b>CHAPTER 2: The MAJORANA DEMONSTRATOR</b> . . . . .	<b>5</b>
Section 2.1: DEMONSTRATOR Goals . . . . .	5
Section 2.2: DEMONSTRATOR Overview . . . . .	7
2.2.1 Shielding Configuration . . . . .	7
2.2.2 Modular Detector Arrays . . . . .	8
2.2.3 P-Point Contact (PPC) HPGe Detectors . . . . .	8
2.2.4 Data Acquisition . . . . .	11
2.2.5 Calibration . . . . .	13
Section 2.3: Background Mitigation and Sensitivity to $0\nu\beta\beta$ Discovery . . . . .	13
Section 2.4: Sources of Background and Background Mitigation in the DEMONSTRATOR . . . . .	15
2.4.1 Cosmogenic Sources . . . . .	15
2.4.2 Primordial Radiation Sources . . . . .	16
Section 2.5: Recognizing and Mitigating Backgrounds through Offline Analysis . . . . .	19
2.5.1 Background Signatures in Data . . . . .	19
2.5.2 Pulse Shape Discrimination (PSD) . . . . .	21
<b>CHAPTER 3: Spectra of the DEMONSTRATOR</b> . . . . .	<b>25</b>
Section 3.1: Data Sets . . . . .	25

Section 3.2: Selection of Data for the Final $0\nu\beta\beta$ Result and Background Analysis . . . . .	26
Section 3.3: Final Results of the MAJORANA DEMONSTRATOR . . . . .	28
Section 3.4: Evidence in Module 1 Natural Detector Spectra for a $^{232}\text{Th}$ Excess . . . . .	30
<b>CHAPTER 4: Simulations of the DEMONSTRATOR . . . . .</b>	<b>32</b>
Section 4.1: Component Groupings . . . . .	32
Section 4.2: Processing with GAT and MPP and PDF Generation . . . . .	33
Section 4.3: Assay-Based Background Model . . . . .	34
Section 4.4: Analysis of $^{232}\text{Th}$ Excess through Simulations . . . . .	34
4.4.1 Module 1 Bulk Component Simulation Studies . . . . .	35
4.4.2 Point Source Studies: Peak Ratios . . . . .	39
4.4.3 Point Source Studies: Relative Rates . . . . .	42
4.4.4 Re-assay of the Thermosyphon Block . . . . .	50
<b>CHAPTER 5: Building a Bayesian Statistical Model . . . . .</b>	<b>51</b>
Section 5.1: Markov Chain Monte Carlo . . . . .	52
Section 5.2: Hamiltonian Monte Carlo, NUTS, and PyMC3 . . . . .	53
Section 5.3: A Statistical Model of the DEMONSTRATOR . . . . .	55
<b>CHAPTER 6: Implementation of Inference Methods . . . . .</b>	<b>57</b>
Section 6.1: Structure of Spectra and Component Groupings . . . . .	57
Section 6.2: Handling of Cosmogenic Decay Chains . . . . .	57
Section 6.3: Handling of $^{40}\text{K}$ and $^{60}\text{Co}$ Priors . . . . .	58
Section 6.4: Simulated Dataset Studies . . . . .	59
6.4.1 Overview and Approach . . . . .	62
6.4.2 Fits with a $^{232}\text{Th}$ excess in Various M1 Component Groupings . . . . .	63
6.4.3 Fits with a $^{232}\text{Th}$ Excess in the M1Cavity Group . . . . .	65
6.4.4 Fits with a Change to the Prior Uncertainties . . . . .	66

6.4.5	Fits with the M1Cavity Region in the Simulated Dataset, But Not in the Component Groupings List . . . . .	67
6.4.6	Fits with Flat Priors . . . . .	68
6.4.7	Simulated Dataset Study Conclusions . . . . .	69
Section 6.5:	Determining the Most Likely Model with Bayes Factors . . . . .	73
6.5.1	Methodology . . . . .	73
6.5.2	Results with Fits to DS0-7 Open and Blind Data . . . . .	75
6.5.3	Results with Fits to Module 2 DS3-6 Open and Blind Data . . . . .	75
6.5.4	Results with Fits to Module 1 DS0-7 Open and Blind Data . . . . .	76
6.5.5	Results when Changing Other Prior Uncertainty Values . . . . .	77
Section 6.6:	Final Fit to DS0-7 Open and Blind Data and Estimate of the $2\nu\beta\beta$ Rate . . . . .	79
<b>CHAPTER 7:</b>	<b>Conclusions . . . . .</b>	<b>82</b>
Section 7.1:	Future Studies . . . . .	83
<b>APPENDIX A:</b>	<b>A CABLE AND CONNECTOR UPGRADE FOR THE</b>	
	<b>MAJORANA DEMONSTRATOR . . . . .</b>	<b>86</b>
Section A.1:	Motivation . . . . .	86
Section A.2:	Overview of Cables . . . . .	86
A.2.1	Signal Cables . . . . .	86
A.2.2	HV Cables . . . . .	86
Section A.3:	Overview of Connectors . . . . .	89
A.3.1	Nano Twist-pin Connectors . . . . .	89
A.3.2	Signal Feedthrough Flange Connectors . . . . .	90
A.3.3	HV Feedthrough Flange Connectors . . . . .	91
A.3.4	HV Fork Connectors . . . . .	92
Section A.4:	Assembly and Fabrication at UNC . . . . .	92
A.4.1	Fabrication of Signal Cable Bundles . . . . .	92
A.4.2	Fabrication of HV Cable Bundles . . . . .	100

A.4.3 Baffle Plate Fabrication . . . . .	101
Section A.5: Installation at SURF . . . . .	101
Section A.6: Results and Impact . . . . .	104
<b>APPENDIX B: COMPONENT GROUPS . . . . .</b>	<b>106</b>
<b>APPENDIX C: DECAY CHAIN INFORMATION . . . . .</b>	<b>112</b>
Section C.1: Decay Chain Segments of $^{232}\text{Th}$ and $^{238}\text{U}$ . . . . .	112
Section C.2: Prominent $\gamma$ -decay Peak Energies . . . . .	113
<b>REFERENCES . . . . .</b>	<b>114</b>

## LIST OF TABLES

Table 3.1:	Table of MAJORANA DEMONSTRATOR dataset information. . . . .	27
Table 4.1:	Table of estimated required activities for candidate components to explain the $^{232}\text{Th}$ excess using the 238 keV peak scaling method as well as their assayed activities. . . .	37
Table 4.2:	Table of 238 keV and 2615 keV peak comparisons between Module 1 and Module 2. The second and third columns list the magnitude of a peak in Module 1 over the magnitude of a peak in Module 2 for a given set of enriched or natural detectors. The peaks are background subtracted. The uncertainties on the data ratios were calculated through the standard Poisson error $\sqrt{N}$ that was propagated through the sideband subtraction and ratio. . . . .	38
Table 4.3:	Table of theoretical background indices from components in Module 2, assuming all of the counts in the DS6 Module 1 BEW can be attributed to the component(s) in the first column. . . . .	39
Table 4.4:	Table of ratios of counts in the 238 keV gamma peak (from $^{212}\text{Pb}$ in the $^{232}\text{Th}$ chain) over the 2615 keV peak (from $^{208}\text{Tl}$ ). The sideband subtracted peak counts are taken from the hot detector that was primarily online throughout the DS3-6 configuration, C1P7D1. The error on the data value was calculated through the standard Poisson error $\sqrt{N}$ that was propagated through the sideband subtraction and ratio. . . . .	41
Table 4.5:	Table of ratios of counts in the 238 keV gamma peak (from $^{212}\text{Pb}$ in the $^{232}\text{Th}$ chain) over the 2615 keV peak (from $^{208}\text{Tl}$ ). The sideband subtracted peak counts are taken from the hot detector that was primarily online throughout the DS3-6 configuration, C1P7D1. The error on the data value was calculated through the standard Poisson error $\sqrt{N}$ that was propagated through the sideband subtraction and ratio. CPICB is an abbreviation for CPInterfaceCavityBottom. . . . .	43
Table 6.1:	Posterior activities of $^{40}\text{K}$ and $^{60}\text{Co}$ component groups in Bq that have upper limits in their assay measurements, following fits where a particular $^{232}\text{Th}$ group was allowed to float with a flat prior and other groups floated with truncated gaussian priors. $^{40}\text{K}$ groups with absolute assay measurements were fixed for these fits. The last column is the original upper limit on the assayed activity value. . . . .	60
Table 6.2:	Posterior activities of $^{40}\text{K}$ and $^{60}\text{Co}$ component groups in Bq that have upper limits in their assay measurements, following fits where a particular $^{232}\text{Th}$ group was allowed to float with a flat prior and other groups floated with truncated gaussian priors. $^{40}\text{K}$ groups with absolute assay measurements were fixed for these fits. The last column is the original upper limit on the assayed activity value. . . . .	61
Table 6.3:	Log marginal likelihoods and log Bayes Factors of several component groups from single component fits where one $^{232}\text{Th}$ group is allowed to float and all other remain fixed. All DS0-7 open and blind data are used in these fits. All log Bayes Factors are compared to the M1Cavity group. . . . .	76

Table 6.4:	Log marginal likelihoods and log Bayes Factors of several component groups from single component fits where one $^{232}\text{Th}$ group is allowed to float and all other remain fixed. Only M2 DS3-6 data is used in these fits. All log Bayes Factors are compared to the M2Cavity group. Note that these Bayes Factors are smaller in general than those in table 6.3. Only 4 submodels are used when fitting to Module 2 data, so the log marginal likelihood of each single-component fit is artificially reduced due to there being fewer submodels than the original set of 17, which leads to smaller log Bayes Factors. . . . .	77
Table 6.5:	Log marginal likelihoods and log Bayes Factors of several component groups from single component fits where one $^{232}\text{Th}$ group is allowed to float and all other remain fixed. Only M1 DS0-7 data is used in these fits. All log Bayes Factors are compared to the M1Cavity group. . . . .	78
Table 6.6:	Log marginal likelihoods and log Bayes Factors of several component groups from single component fits where one $^{232}\text{Th}$ group is allowed to float and all other remain fixed. Only M1 DS0-7 data is used in these fits. All log Bayes Factors are compared to the M1Cavity group. . . . .	78
Table 6.7:	Log marginal likelihoods and log Bayes Factors of the M1Cavity and M1LMFEs groups in a “worst-case scenario” configuration from single component fits where one $^{232}\text{Th}$ group is allowed to float and all other remain fixed. Only M1 DS0-7 data is used in these fits. The M1LMFEs log Bayes Factor is compared to the M1Cavity group. . . . .	78
Table 7.1:	Table of prominent $^{238}\text{U}$ chain peak counts in the assay-based model versus DS3-6 data. The detectors are divided into various groupings, including M1 only, M2 only, Enr only, and Nat only. . . . .	84
Table A.1:	Table of Axon’ signal cable properties. . . . .	87
Table A.2:	Table of Axon’ HV cable properties. . . . .	88
Table A.3:	Table of HV cable clamping test results. . . . .	89
Table A.4:	Table of Axon’ Nano Twist-Pin Connector assay results, compared to the previously used Vespel connectors. Assays performed at Pacific Northwest National Laboratory (PNNL). . . . .	90
Table A.5:	Table of changes made to detector and string positions during the upgrade. If a detector is not specified, then the change applies to the entire string with detector position along the vertical axis unchanged. The four ICPC detectors were installed to the strings in positions 4 and 5, with two detectors at the bottom of each. . . . .	105
Table B.1:	List of parts and their corresponding component groups. Note that if the component group begins with M1 or M2, it is assumed that the list of parts are actually pairs of parts - one for each module. . . . .	106
Table C.1:	Prominent peaks from $\gamma$ -decays observed by the DEMONSTRATOR. Intensity refers to the intra-isotope intensity. . . . .	113

## LIST OF FIGURES

Figure 1.1:	The effective Majorana mass as a function of the lightest neutrino mass. The spread in valid $\langle m_{\beta\beta} \rangle$ values assuming an inverted neutrino mass ordering is shown in green, whilst assuming a normal neutrino mass ordering is in red. These bands converge above around 0.015 eV. The darker shaded areas of the green and red bands are $3\sigma$ regions from the error propagation of the PMNS matrix element uncertainties. The gray band on top is the region of parameter space excluded from previous experiments, and the blue band is the most recently excluded region based on the results of the KamLAND-Zen 800 experiment [6]. The blue dotted line highlights the expected sensitivity of ton-scale $0\nu\beta\beta$ experiments. This figure is adapted from [5] and [7]. . . .	2
Figure 2.1:	The MAJORANA DEMONSTRATOR shield system layout. Additionally shown here is one of the crossarms which link each of the modules to its cooling system, electronic boxes, and high vacuum system. The assembly is nine feet high, for scale. [17] . . . .	7
Figure 2.2:	Configuration of the string position numbers in a cryostat. This is a top-down view of the “coldplate”, which sits at the top of the cryostat, with the crossarm running to it from the left. The coldplate is a copper plate from which the strings of detectors are suspended. . . . .	9
Figure 2.3:	Components of the MAJORANA DEMONSTRATOR detector unit and string assembly [17]. . . . .	10
Figure 2.4:	Weighting potential of a PPC detector, mapped by collaborator Benjamin Shanks [32]. The point contact is located at (0,0). Each black line is an “isochrone” - locations from which holes have an equal amount of time to drift to the point contact. Each isochrone from the point contact is an additional 100 ns of drift time. . . . .	11
Figure 2.5:	Waveform of a single-site gamma-ray physics event read out by a PPC detector and digitized. The initial shallow rise (imperceptible in this figure) from baseline accounts for the majority of the drift path the holes undertake. The sharp rise accounts for the final period of drift when the holes are within the innermost isochrone of Figure 2.4 relative to the point contact [32]. . . . .	12
Figure 2.6:	The sensitivity to a $0\nu\beta\beta$ decay signal in $^{76}\text{Ge}$ as a function of exposure and background for a $3\sigma$ (99.7% C.L.) discovery sensitivity (DS). Note, the background rates are normalized to a 2.5 keV FWHM energy resolution. [37] . . . . .	14
Figure 2.7:	$^{232}\text{Th}$ decay chain. Notable decays important to background studies include the alpha decay of $^{220}\text{Rn}$ and beta decay of $^{208}\text{Tl}$ [39]. . . . .	17
Figure 2.8:	$^{238}\text{U}$ decay chain. A notable decay important to background studies is the alpha decay of $^{222}\text{Rn}$ [40]. . . . .	18

Figure 2.9:	A graphic that illustrates the dominant gamma ray interaction in the parameter space defined by the energy of the incident gamma ray (in MeV) and the proton number of the absorbing material ( $Z=32$ for germanium). The black lines are points in the parameter space where the probabilities of two of the interactions are equivalent. [29]	20
Figure 2.10:	Comparison of a single-site gamma-ray waveform (top) and a multi-site waveform (bottom). The multi-site current waveform (bottom, red) clearly shows several peaks of lower amplitude which the AvsE cut can detect. [42]	22
Figure 2.11:	Illustration of the regions used to make a measurement for the LQ cut. The area in red determines whether the event will pass the LQ cut, with larger areas tending to fail the cut. Regions I and II are used to numerically estimate the area in red. [43]	23
Figure 2.12:	A waveform from a $\gamma$ event in blue and a waveform from an $\alpha$ event in red. $\alpha$ events tend to have a shallow rising tail compared to $\gamma$ events, which is tagged by the DCR cut. [44]	24
Figure 3.1:	The energy spectra, before and after cuts, and results from 64.5 kg-yr enriched detector operation of the DEMONSTRATOR. The background estimation window is highlighted in light blue and shown in the inset. The blue vertical band around 2039 keV in the inset is the 10-keV window around $Q_{\beta\beta}$ , while the gray vertical bands in the inset are the 10-keV windows around peaks from the decays of $^{208}\text{Tl}$ and $^{214}\text{Bi}$ . These four vertical bands are excluded from the BEW. Figure adapted from the data in Ref. [15]	28
Figure 3.2:	Predicted energy spectrum of the assay-based background model (blue curve) plotted against a 50.83 kg-yr subset of the enriched-detector data (black). The model has been broken down into its decay chain constituents, represented by other colors. Note that while $^{57}\text{Co}$ and $^{54}\text{Mn}$ are also contributors to the assay-based model, their spectra fall below the y-scale of this figure and were thus excluded.	29
Figure 3.3:	Direct evidence for a non-uniform $^{232}\text{Th}$ excess in Module 1 natural detectors, based on rates of $^{232}\text{Th}$ -chain $\gamma$ -ray peaks. <b>Left:</b> Rate of the 2615-keV peak in individual detectors through datasets before the removal of enriched detectors. Detectors C1P2D1 and C1P7D1 have the highest rates. Detectors in gray were not operational during this period. <b>Right:</b> Rate of the 238-keV peak in operable Module-1 natural detectors through datasets before the removal of enriched detectors. Detector C1P7D1 has a much higher rate than the other detectors. Detector C1P2D1 is excluded due to limited statistics as it was biased down early in the dataset.	30
Figure 3.4:	The two natural detectors which saw an excess of events between 100 and 3000 keV when positioned next to the Module 1 crossarm no longer had elevated rates after being moved to a new location in Module 2.	31
Figure 4.1:	The location of components identified as potential source for the $^{232}\text{Th}$ excess in Module 1. One of the two natural detectors that observed a high rate from the $^{232}\text{Th}$ chain, C1P7D1, is labeled.	36

Figure 4.2:	Locations of the point sources simulated to evaluate the possibility of a $^{232}\text{Th}$ excess from a point source. Left is a cutaway graphic of the module 1 cryostat (point sources are yellow circles). Right is a full 7-string graphic of Module 1 (point sources are the numbered rectangles). . . . .	41
Figure 4.3:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the point source <code>CPInterfaceCenter</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	42
Figure 4.4:	A top-down view of the thermosyphon block showing locations of the sources simulated on the bottom surface of the cavity to study the relative rates from the $^{232}\text{Th}$ chain in different detectors. The yellow box outlines the surface area where primaries were generated for the <code>CPInterfaceCavityBottomSurface</code> source. The detector positions are not to scale, and in reality C1P7D1 and C1P2D1 are partially beneath the thermosyphon block. . . . .	44
Figure 4.5:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the point source <code>CPInterfaceCavityBottomFront</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	45
Figure 4.6:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the point source <code>CPInterfaceCavityBottomBack</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	45
Figure 4.7:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the point source <code>BafflePlateBottom</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	46
Figure 4.8:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the point source <code>CPInterfaceCavityBottomTM2</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	46
Figure 4.9:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the point source <code>CPInterfaceBottom</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	47
Figure 4.10:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the point source <code>ClampPlateTop</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	47

Figure 4.11:	A side view of simulated primaries generated inside the copper bulk of the thermosyphon block ( <code>CPInterface</code> ) as well as the thermosyphon bolts ( <code>TSBoltsBulk</code> ) that hold the block in place. The components are outlined in green and blue respectively. . . . .	48
Figure 4.12:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the grouping <code>TSBoltsBulk</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	48
Figure 4.13:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the grouping <code>TSBoltsSurf</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	49
Figure 4.14:	Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. <b>Left:</b> Rates from the surface source <code>CPInterfaceCavityBottomSurface</code> . <b>Right:</b> Rates from the point source <code>CPInterfaceCavityBottomCenter</code> . . . . .	49
Figure 5.1:	Sampling results from different distributions using the Metropolis algorithm. Density refers to the density of samples, and each peak contains 1000 samples of a different distribution. The dimensionality of the distribution is shown above the peak. This demonstration was done by Richard McElreath for his book “Statistical rethinking : a Bayesian course with examples in R and Stan”. [51] . . . . .	53
Figure 5.2:	Sampling results after four samples ( $X$ 's) from a 2D Gaussian posterior using HMC and two different leapfrog step values ( $L=11$ and $L=28$ ) with the same step sizes. <b>Left:</b> With a good $L$ value, HMC explores the posterior efficiently. <b>Right:</b> With a poor $L$ value, the particle can return to the area near the origin at each sampling point, leading the sampler to getting stuck. This can be avoided by using NUTS. This demonstration was done by Richard McElreath for his book “Statistical rethinking : a Bayesian course with examples in R and Stan”. [51] . . . . .	54
Figure 6.1:	Fits to simulated datasets with a $10x$ $^{232}\text{Th}$ excess in the <code>M1LMFEs</code> group. . . . .	63
Figure 6.2:	Fits to simulated datasets with a $10x$ $^{232}\text{Th}$ excess in the <code>M1CrossarmAndCPCables</code> group. . . . .	64
Figure 6.3:	Fits to simulated datasets with a $10x$ $^{232}\text{Th}$ excess in the <code>RadShieldPb</code> group. . . . .	64
Figure 6.4:	Fits to a “MJD-stats” simulated dataset with some $^{232}\text{Th}$ excess in the <code>M1Cavity</code> group. . . . .	65

Figure 6.5:	Fit to a “MJD-stats” simulated dataset with a 3x <sup>232</sup> Th excess in the M1Cavity group with a 0.25x prior. . . . .	66
Figure 6.6:	Comparison of fits to a “MJD-stats” simulated dataset when prior uncertainties are changed. . . . .	67
Figure 6.7:	Fit with the M1Cavity group in the dataset but not in the list of components. Ignore the second line of the plot title. The number of counts in the simulated dataset for each component is equal to the prior counts, so the blue bars are not shown here. . .	68
Figure 6.8:	Fits with flat priors to a “MJD-stats” simulated dataset. . . . .	70
Figure 6.9:	Plots of the energy spectrum from the fitted model and data. The fitted model is broken down into its decay chain subsets, which are represented by the colored spectra. These are from fits using flat priors to a “MJD-stats” simulated dataset, with a factor of 10 <sup>232</sup> Th excess placed in one component group. Note that these fits severely underfit several gamma peaks, as well as the energy spectrum at low energies. The scatter plot of pulls at the bottom of each sub-figure show the extent of the discrepancies as a function of energy. . . . .	71
Figure 6.10:	Integrated count plots from various decay chains. These are from fits using flat priors to a “MJD-stats” simulated dataset, specifically with a 3x <sup>232</sup> Th excess in the M1Cavity group. . . . .	72
Figure 6.11:	Fit to DS0-7 open and blind data, enriched and natural detectors. These are the integrated <sup>232</sup> Th counts in all component groups, with the gray bars representing the number of counts as given by the prior, and the green bars representing the number of counts as given by the posterior. . . . .	80
Figure 6.12:	Fit to DS0-7 open and blind data, enriched and natural detectors. . . . .	81
Figure A.1:	Schematic of an MJD module, highlighting the locations of cables and connectors. . .	87
Figure A.2:	One half of an Axon’ Nano twist-pin connector. Axon’ delivered these connectors pre-terminated to signal cables. . . . .	91
Figure A.3:	Comparison of old and new HV Feedthrough Connector designs. . . . .	93
Figure A.4:	Cable and connector setup at the feedthrough flange, before and after the upgrade. Notice the improved cable bundling with plain dental floss in the background. . . . .	94
Figure A.5:	Comparison of old and new HV Fork designs. . . . .	95

Figure A.6: Overview of HV Fork and HV Ring placement. . . . .	96
Figure A.7: The cable stripping machine used for the upgrade and a stripped cable. . . . .	97
Figure A.8: Fabrication of a signal cable bundle. The ground and central conductor of the signal cable is passed through the back shell of the 50-pin D-sub before being crimped into sockets. Those sockets are then loaded into the main body of the D-sub. . . . .	98
Figure A.9: Termination of signal cable to LMFE terminals. Collaborators perform this work under a microscope with silver epoxy at LBNL. . . . .	99
Figure A.10: Configuration of the heat shrink plastic relative to the exposed ground shield wires, according to NASA specifications. [58] . . . . .	100
Figure A.11: Old and new baffle plate configurations. . . . .	102
Figure A.12: The 3D-printed Axon' connector holder, known colloquially as the "diving board" due to its position on the cryostat-end of the thermosyphon (see Figure A.1). The fabrication team at UNC securely positioned the nano twist-pin connectors on the board with kapton tape for the shipment to SURF. During installation, a team of experts carefully removed the tape and transferred five Axon' connectors at a time to the identical position on a copper structure of the same design. . . . .	103
Figure A.13: Configuration of the string position numbers in a cryostat. This is a top-down view of the cryostat with the crossarm running to it from the left. . . . .	104
Figure C.1: Four segments of the $^{232}\text{Th}$ chain. [63] . . . . .	112
Figure C.2: Ten segments of the $^{238}\text{U}$ chain. [63] . . . . .	113

## LIST OF ABBREVIATIONS

$0\nu\beta\beta$	Neutrinoless double beta decay
$2\nu\beta\beta$	Two neutrino double beta decay
BEW	Background estimation window
BF	Bayes factor
CCF	Cryostat copper far
CCJ	Cryostat hoop junction
CCN	Cryostat copper near
CCP	Crossarm and cold plate
CSA	Charge sensitive amplifier
DCR	Delayed charge recovery
FWHM	Full width half max
GAT	Germanium analysis toolkit
HMC	Hamiltonian Monte Carlo
HV	High voltage
ICPC	Inverted coaxial point contact detector
LMFE	Low mass front end board
LQ	Late charge
MCMC	Markov chain Monte Carlo
ML	Marginal likelihood
MPP	MaGe-Post-Processor
NUTS	No U-turn sampler
PDF	Probability density function

PPC P-Point contact detector  
PSD Pulse shape discrimination  
ROI Region of interest  
UGEFCu Underground electroformed copper

## CHAPTER 1: Neutrinoless Double Beta Decay ( $0\nu\beta\beta$ )

Out of Enrico Fermi's treatment of  $\beta$ -decay [1] came a postulation of double beta decay by Maria Goeppert Meyer in 1937 [2]. Two neutrino double beta decay ( $2\nu\beta\beta$ ) is a second-order weak process and the rarest observed nuclear process. In a nucleus with an even number of protons and neutrons and under the condition that  $\beta$ -decay is a processes that is energetically disfavored, two neutrons can decay into two protons, emitting in the process two electrons and two anti-neutrinos. Thus far it has only been observed to occur in 11 nuclei, including  $^{76}\text{Ge}$ , at half-lives ranging from  $10^{18}$  to  $10^{24}$  years [3].

In 1937, Wendell Furry published a study on double beta decay transition probabilities [4]. It was in this paper that Furry identified Ettore Majorana's theory of the neutrino as an interesting source of an alternate scenario of double beta decay, one in which instead of emitting two electrons and two anti-neutrinos, a nucleus emits only two electrons. This process is only possible if the neutrino is its own anti-particle. In the present day such a neutrino is called a "Majorana particle", and this process is called neutrinoless double beta decay.

Neutrinoless double beta decay has yet to be observed by experiment, yet its discovery would have implications for physics beyond the standard model. The decay violates total lepton number conservation - an "accidental symmetry" that has yet to be observed to be violated by experiment but is not based on any underlying symmetry. Violation of this conserved quantity has implications in theories of leptogenesis, which possibly explain the baryon excess in our universe [5]. A positive result for  $0\nu\beta\beta$  would also definitively classify the neutrino as a Majorana particle and would provide information on the neutrino mass.

The half-life is the observable in experimental searches for  $0\nu\beta\beta$ . The half-life can be expressed as<sup>1</sup>,

$$T_{1/2} = G^{0\nu}(Q, Z)|M^{0\nu}|^2 \left( \frac{\langle m_{\beta\beta} \rangle}{m_e} \right)^2. \quad (1.1)$$

$G^{0\nu}$  is the lepton-phase space integral, which depends on the Q-value ( $Q_{\beta\beta}$ ) for  $0\nu\beta\beta$  (nearly all of the energy carried away by the beta electrons in  $0\nu\beta\beta$ ) as well as the charge (proton number) of the final state nucleus [3].  $M^{0\nu}$ , a dimensionless number, is the nuclear matrix element (NME) for  $0\nu\beta\beta$  and  $\langle m_{\beta\beta} \rangle$  is the effective Majorana mass of the neutrino, which is normalized by the electron mass  $m_e$ . Both the nuclear matrix element and effective Majorana mass are explained in more detail below.

---

<sup>1</sup>Assuming the most prevalent  $0\nu\beta\beta$  mechanism "light neutrino exchange", to be defined on page 3

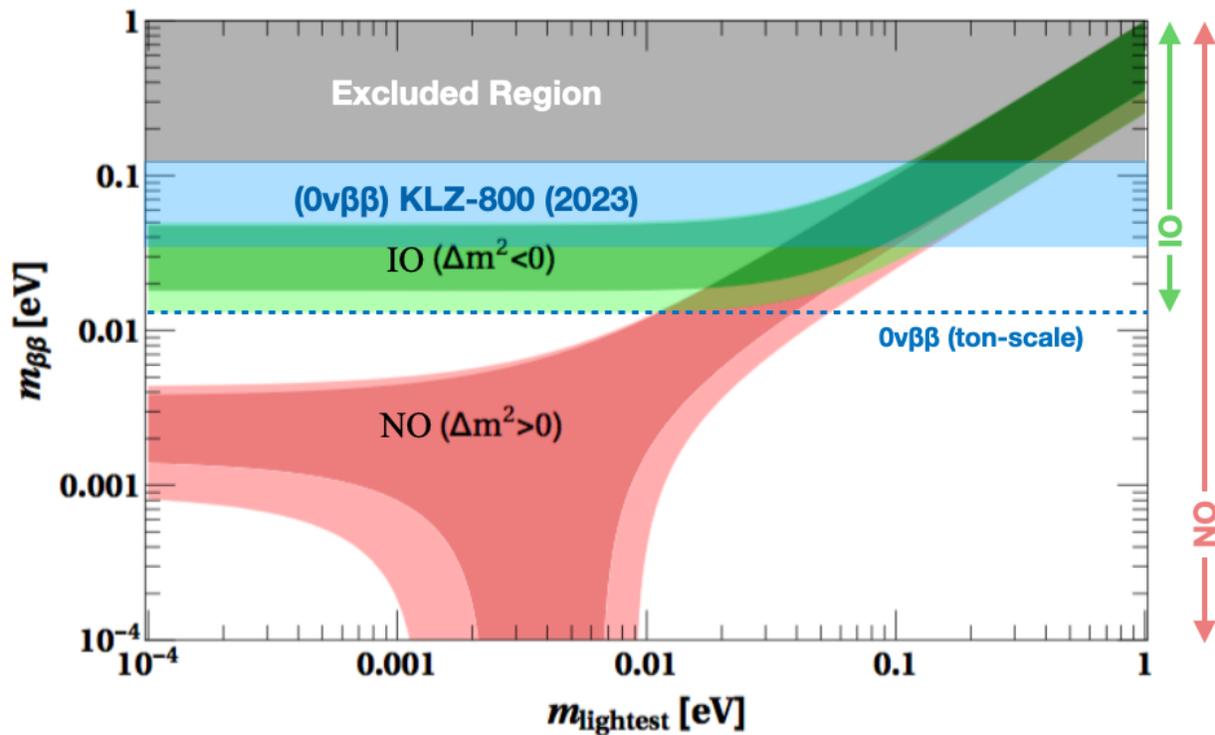


Figure 1.1: The effective Majorana mass as a function of the lightest neutrino mass. The spread in valid  $\langle m_{\beta\beta} \rangle$  values assuming an inverted neutrino mass ordering is shown in green, whilst assuming a normal neutrino mass ordering is in red. These bands converge above around 0.015 eV. The darker shaded areas of the green and red bands are  $3\sigma$  regions from the error propagation of the PMNS matrix element uncertainties. The gray band on top is the region of parameter space excluded from previous experiments, and the blue band is the most recently excluded region based on the results of the KamLAND-Zen 800 experiment [6]. The blue dotted line highlights the expected sensitivity of ton-scale  $0\nu\beta\beta$  experiments. This figure is adapted from [5] and [7].

There are several proposed mechanisms for  $0\nu\beta\beta$ , but the most prevalent of these is known as "light neutrino exchange", when an anti-neutrino is emitted from one vertex and absorbed by the other vertex. This theoretical mechanism contains a parameter known as the effective Majorana mass ( $\langle m_{\beta\beta} \rangle$ ) which serves to describe the physics of massive neutrinos. Unsurprisingly, this parameter is partially defined by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [8][9] which describes the mixing of neutrino mass and flavor eigenstates. This mixing also gives rise to neutrino oscillations - the observed phenomenon that proves neutrinos have mass. The full definition of  $\langle m_{\beta\beta} \rangle$  is:

$$\langle m_{\beta\beta} \rangle = |u_{e1}^2 m_1 + u_{e2}^2 e^{i\alpha_1} m_2 + u_{e3}^2 e^{i\alpha_2} m_3|. \quad (1.2)$$

$u_{ej}$  are the mixing amplitude elements of the PMNS matrix, the masses  $m_j$  are the three neutrino mass eigenstates, and  $\alpha_k$  are two phase parameters that are allowed if neutrinos are Majorana particles and are hence known as the "Majorana phases". Figure 1.1 illustrates the spread of possible  $\langle m_{\beta\beta} \rangle$  values that have yet to be excluded by observation as a function of the lightest neutrino mass. Such a spread exists for the following reasons:

- The Majorana phases are unknown.
- Elements of the PMNS matrix each have uncertainties associated with them from neutrino oscillation experiments.
- The neutrino mass ordering is unknown, resulting in two separate  $\langle m_{\beta\beta} \rangle$  ranges at lower values of the lightest neutrino mass.

The NME values for each  $0\nu\beta\beta$  isotope are an important contributor to the uncertainty in relating the half-life to the effective neutrino mass. A number of different models are used to calculate the NME, but even within the same isotope the various models output results that differ from each other by factors of 2-3 [3]. Therefore it seems that many of these models are missing important elements of the nuclear structure involved in  $0\nu\beta\beta$ . Furthermore, when these models are applied to two-neutrino double beta decay and compared with measurements, the resulting half-lives are systematically faster than measurement. As a result, theorists have inserted a quenching factor on the axial coupling  $g_A$  into the NME calculation that varies from element to element. This resolves the systematic underestimate of the half-life, but adds to the NME uncertainty. Recent work [10] using *ab initio* calculations of the many-body structures of various  $0\nu\beta\beta$  candidate nuclei avoid the perturbative shell-model effective interaction framework. These calculations, rooted in chiral effective field theory, have produced NME values on the lower end of the range of possible

values. However, these calculations disfavor extremely low NME values and are currently a work in progress, with refined calculations expected to be completed in the near future.

Given the current status of the half-life calculation including all of its uncertainties, theorists have been able to provide order of magnitude estimates of the size of  $\langle m_{\beta\beta} \rangle$  for a measured half-life: In the range of  $10^{26}$  to  $10^{28}$  years for both mass orderings and greater than  $10^{28}$  years for the normal ordering. It is expected that most next-generation  $0\nu\beta\beta$  experiments will be able to extend their sensitivity through  $10^{28}$  years.

## CHAPTER 2: The MAJORANA DEMONSTRATOR

Though there are up to eleven isotopes known to double beta decay, three have been used by the majority of  $0\nu\beta\beta$  experiments.  $^{136}\text{Xe}$  was the isotope of choice for the KamLand-Zen and EXO-200 experiments. KamLand-Zen uses a large volume of  $^{136}\text{Xe}$ -doped liquid scintillator enriched to 90.85%  $^{136}\text{Xe}$  as their source and detection medium [6], while EXO-200 used a liquid Xenon time projection chamber (TPC) enriched to 80.6% in  $^{136}\text{Xe}$ [11]. These experiments benefited from a high active mass of  $0\nu\beta\beta$  isotope, but have modest energy resolution.  $^{130}\text{Te}$  is used by the CUORE experiment [12] and will be used by the SNO+ experiment [13]. SNO+ plans to deploy their  $^{130}\text{Te}$  through liquid scintillator, while CUORE has an array of crystal tellurium bolometers housed in a dilution refrigerator. It is through this implementation that CUORE benefited from a much better energy resolution, though the collaboration faced greater difficulty with background discrimination.

$^{76}\text{Ge}$  is the third of these isotopes, used by both the GERDA [14] and MAJORANA DEMONSTRATOR [15] experiments. Both experiments deployed crystals enriched in  $^{76}\text{Ge}$ , but in different ways. GERDA operated their array of crystals in liquid argon which served to cool the detectors as well as passively shield them from backgrounds. The liquid argon also served as an active Compton scattering veto to mitigate gamma backgrounds [14]. The MAJORANA DEMONSTRATOR operated crystals in a vacuum cryostat surrounded by ultrapure electroformed copper [15]. Both experiments demonstrated excellent energy resolution as well as low backgrounds. Though fabricating individual large mass crystals ( $>2\text{kg}$ ) of  $^{76}\text{Ge}$  has its own set of unique challenges, LEGEND, which combines the strengths of GERDA and the MAJORANA DEMONSTRATOR, is scaling up with such crystals, aiming for 1000 kg of active mass as LEGEND-1000, a next-generation  $0\nu\beta\beta$  experiment based in germanium [16].

An overview of all of these experiments and their next-generation counterparts can be found in the referenced publication by Cardani [3]. The remainder of this chapter will discuss the MAJORANA DEMONSTRATOR with a focus on germanium detector technology and background mitigation.

### Section 2.1: DEMONSTRATOR Goals

The MAJORANA DEMONSTRATOR was a  $^{76}\text{Ge}$   $0\nu\beta\beta$  experiment located at the 4850' (1.48 km) level of the Sanford Underground Research Facility (SURF) in Lead, South Dakota. It deployed an array of enriched

germanium semiconductor detectors that also served as the  $0\nu\beta\beta$  source. The DEMONSTRATOR aimed to achieve a list of goals related to its  $0\nu\beta\beta$  search, energy resolution, and background mitigation techniques [17]:

1. Measure the  $0\nu\beta\beta$  rate, likely as a lower-limit with 65 kg.yrs of exposure.
2. Demonstrate the feasibility of building a tonne-scale, modular detector array.
3. Demonstrate the ability to achieve a background rate at or below 2 cts./(FWHM-t-yr) at  $Q_{\beta\beta}$ .
4. Achieve an energy resolution at or better than 2.5 keV FWHM at  $Q_{\beta\beta}$ .
5. Perform searches for physics beyond the standard model, other than for  $0\nu\beta\beta$ .

The DEMONSTRATOR achieved all of these goals with the exception of goal 3. The final  $0\nu\beta\beta$  result that satisfies goal 1 is discussed in section 3.3. The achievement of goal 2 relied on the arrangement and operation of the modular germanium detector arrays. Since the detection efficiency of  $0\nu\beta\beta$  does not appreciably change with individual detector size or number of detectors, the number of modules can be multiplied to reach the required mass of germanium for next-generation experiments like LEGEND-1000 while also maintaining the excellent energy resolution of the DEMONSTRATOR.

Motivation for goal 3 originated from sensitivity calculations for the two phase LEGEND experiment. The first phase, named LEGEND-200 with a 200 kg germanium crystal payload, aims to be sensitive to  $0\nu\beta\beta$  half-lives above  $10^{27}$  years. This requires a background rate a bit below 1 ct./(FWHM-t-yr) at  $Q_{\beta\beta}$ . The second phase, LEGEND-1000, will aim to probe  $0\nu\beta\beta$  half-lives up to  $10^{28}$  years, and thus will require a background rate lower than 0.025 cts./(FWHM-t-yr). Later chapters will describe that, while goal 3 was ultimately missed, the potential for achieving such a background level is likely by combining the mitigation techniques used for both the DEMONSTRATOR and GERDA. LEGEND-200 is using and has improved upon these techniques. Further steps are planned to achieve the background rate goal for LEGEND-1000.

Unique properties of the detectors and electronic readout, described in section 2.2.3, allowed the DEMONSTRATOR to satisfy goal 4. Finally, goal 5 was motivated by the low energy detection threshold the DEMONSTRATOR was able to achieve. An excellent energy resolution and low backgrounds combined to yield a rich and broad physics program that included searches for dark matter [18][19], violations of fundamental symmetries [20], searches for solar axions [21], lightly ionizing particles [22], monopoles, baryon decay [23], charge non-conservation [20], exotic fermions and majoron emission during  $\beta\beta$  decays, and tests of quantum mechanics [24].

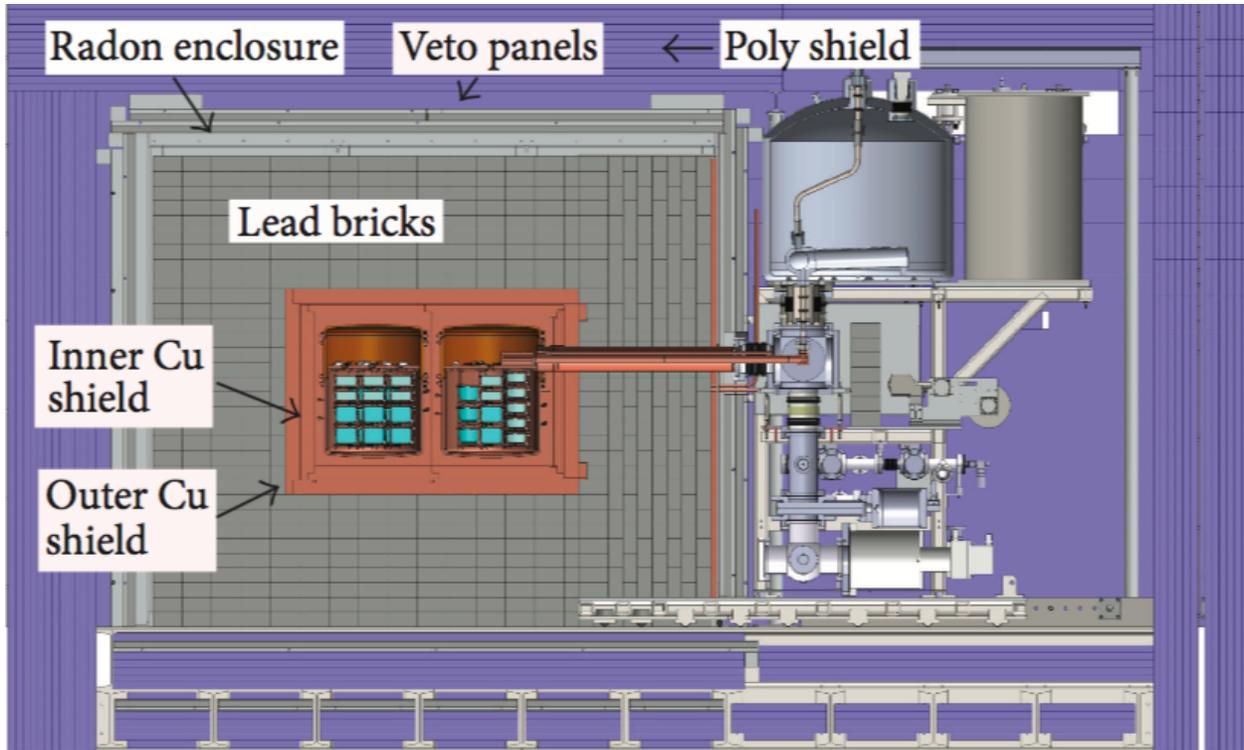


Figure 2.1: The MAJORANA DEMONSTRATOR shield system layout. Additionally shown here is one of the crossarms which link each of the modules to its cooling system, electronic boxes, and high vacuum system. The assembly is nine feet high, for scale. [17]

## Section 2.2: DEMONSTRATOR Overview

### 2.2.1: Shielding Configuration

Figure 2.1 illustrates the shielding arrangement of the DEMONSTRATOR. The two cryostats, made from underground electroformed copper (UGEFCu) were surrounded by an inner copper shield comprised of the same material. An outer copper shield of commercially obtained oxygen free high conductivity (OFHC) C10100 copper surrounded the inner shield, followed by a layer of high-purity lead shielding. These components were enclosed within a thin radon exclusion box purged with liquid nitrogen boil off. Installed around the radon enclosure were plastic scintillator panels that formed an active muon veto system. Layers of borated polyethylene and standard polyethylene completed the shielding structure. Additionally, at its location of 4850 ft. under the earth's surface, the DEMONSTRATOR benefited from a 4300 meters water equivalent (m.w.e.) rock overburden to shield the experiment from cosmogenic muons. Hollow copper crossarms, one running to each cryostat, penetrated these shielding layers to allow passage between the cryostats for the high-vacuum pumping and the signal and high voltage (HV) cables. The crossarm also contains the thermosyphon, a long copper tube that contains nitrogen in liquid and gaseous forms. Liquid nitrogen flows the

length of the tube to where it makes contact with the cryostat, cooling the vessel. The heat transfer from the cryostat boils the liquid nitrogen, and the resulting gas flows backward through the tube to be condensed via a heat exchanger at the other end. This closed loop forms the backbone of the liquid nitrogen cooling system.

### 2.2.2: Modular Detector Arrays

The high-purity germanium (HPGe) detectors of the DEMONSTRATOR were housed in the cryostats, referred to as “Module 1” (M1) and “Module 2” (M2). The modules contained a total of 35 detectors that were  $88.1 \pm 0.7\%$  enriched in  $^{76}\text{Ge}$ , totaling 29.7 kg [25]. The DEMONSTRATOR additionally contained 23 HPGe detectors fabricated from natural material with an abundance of  $^{76}\text{Ge}$  (7.8%) - totalling 14.4kg. These detectors had individual masses between 0.6 and 1.1kg [25].

Only UGEFCu machined in-house and NXT-85 [26] - a form of teflon manufactured in a clean room environment - were used in the support structure for the individual detectors [17]. This support structure consisted of individual frames, referred to as “detector units”, that housed the crystal, HV components, and signal components including the Low-Mass Front End board (LMFE). Figure 2.3a illustrates and labels the parts that comprised the detector units. The HV ring was the only copper component of the detector unit that the crystal was in contact with. This component applied high voltage from the HV fork to the n-contact of the detector. Otherwise, the crystal was kept electrically isolated from the other copper components by the use of crystal insulators made of NXT-85. Up to five detector units were threaded along copper rods to form a vertical stack of crystals, referred to as a “string”. Each module contained seven strings of detectors, with one string located in the center of the cryostat and the others forming a circle around the central string (see Figure 2.2). Figure 2.3b illustrates and labels the parts that comprised the strings.

### 2.2.3: P-Point Contact (PPC) HPGe Detectors

The enriched P-Point Contact HPGe detector [27][28] has all of the advantages of the traditional Coaxial HPGe detector, but with added benefits due to its geometry and signal properties. In particular, the PPC detector offers superb multisite discrimination capabilities due to the variation in drift times for holes (empty spaces in the crystal lattice) produced by ionizing radiation in the crystal. Additionally, these detectors feature a low capacitance that reduces electronic noise. This means the PPC detector has excellent energy resolution and the ability to read out characteristic energy signals down to below 1 keV, which as mentioned previously allowed the DEMONSTRATOR to participate in searches for physics beyond the standard model. Finally, germanium detectors are intrinsically low background, and have the ability to play the role as both

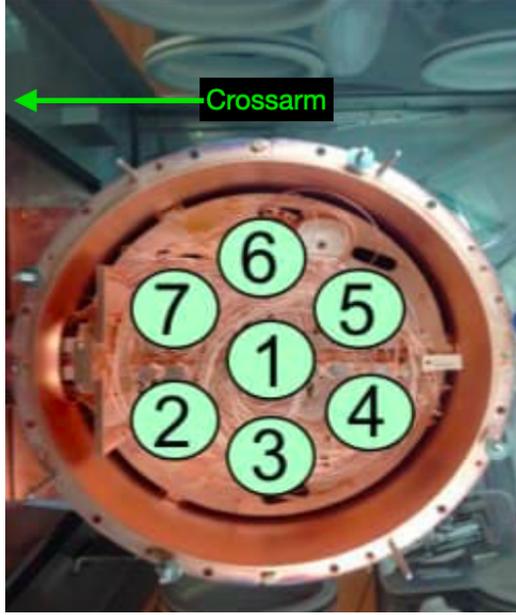


Figure 2.2: Configuration of the string position numbers in a cryostat. This is a top-down view of the “coldplate”, which sits at the top of the cryostat, with the crossarm running to it from the left. The coldplate is a copper plate from which the strings of detectors are suspended.

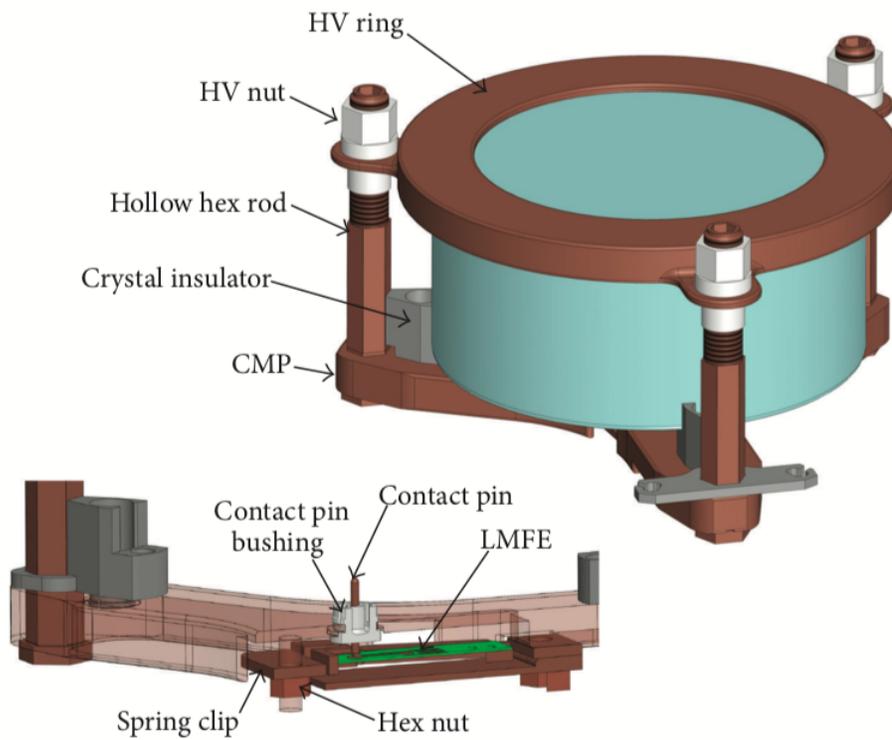
$0\nu\beta\beta$  detector and source.

The PPC HPGe detectors in the array operate with a high voltage reverse bias. As semiconductors with a largely p-type bulk and heavily doped n-type surface, a “p-n junction” exists in each detector - an effective border between regions with an excess of electrons and regions with an excess of holes. This junction creates a region in which electrons from the n-type side have diffused over the border and recombined with holes, and vice-versa. In equilibrium this region is called the “depletion region” and is characterized by low concentrations of mobile majority carriers (electrons from n-type regions and holes from p-type regions). Thus, ionizing radiation (like from  $0\nu\beta\beta$ ) can free electron-hole pairs in this depletion region. A reverse bias on the detector greatly enhances this effect, extending the depletion region through the p-type bulk and mitigating the movement of majority carriers [29].

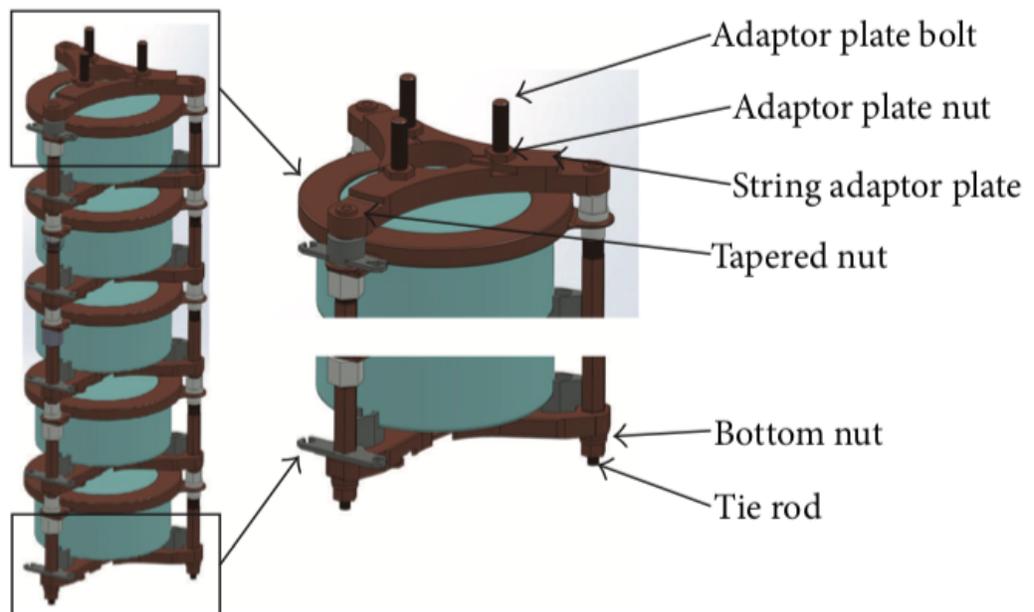
The movement of the electron-hole pairs through the crystal induces a charge at the readout electrode, referred to as the point contact. The magnitude of the induced charge is proportional to the movement of the holes and electrons through the biased crystal bulk, as given by the Shockley-Ramo theorem [30][31]:

$$Q(t) = q\Delta\phi_0(t) \tag{2.1}$$

$Q(t)$  is the total charge induced at the point contact,  $q$  is the charge of the electron or hole, and  $\phi_0$  is the “weighting potential” of the detector. The weighting potential (Figure 2.4) is defined as the solution to



(a) Components of the MAJORANA DEMONSTRATOR detector unit. CMP stands for Crystal Mounting Plate.



(b) Components of the MAJORANA DEMONSTRATOR string assembly.

Figure 2.3: Components of the MAJORANA DEMONSTRATOR detector unit and string assembly [17].

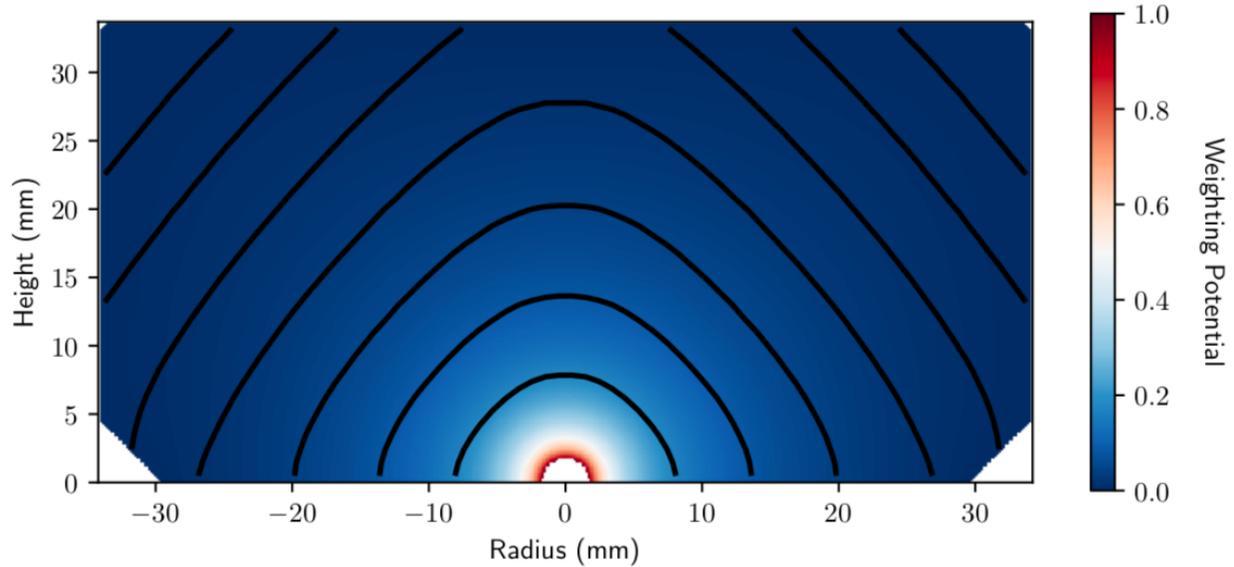


Figure 2.4: Weighting potential of a PPC detector, mapped by collaborator Benjamin Shanks [32]. The point contact is located at (0,0). Each black line is an “isochrone” - locations from which holes have an equal amount of time to drift to the point contact. Each isochrone from the point contact is an additional 100 ns of drift time.

the Laplace equation in which the point contact is fixed at 1 and the heavily n-doped surface is fixed at 0.

Due to the geometry of the PPC detector, the weighting potential increases sharply very close to the point contact. Therefore, most of the charge induced by the movement of holes towards the point contact will occur in the short period of time at the end of the drift path. This manifests itself in waveforms with sharp rise-times (see Figure 2.5). Since the waveform of a physics event has such a characteristic shape, it becomes easier to identify background events through pulse shape discrimination, to be discussed in section 2.5.2.

#### 2.2.4: Data Acquisition

The induced charge was read out of the point contact and passed through a charge sensitive amplifier (CSA) [33]. A radiopure, custom circuit board installed near the detector called a “Low Mass Front End” (LMFE) board contained part of this CSA, namely a JFET and feedback resistor. The signal traveled the length of an approximately 2m long miniature  $50 \Omega$  coaxial cable to reach the pre-amplifier part of the CSA. The pre-amplifiers sat on motherboards that were organized by string and output detector signals in two channels, one at high-gain (up to  $\sim 3$  MeV) and one at low-gain (up to  $\sim 10$  MeV). These signals were transmitted to and digitized by GRETINA digitizer boards (the same boards used in the GRETINA experiment) which were housed in VME crates [34]. The digitized signals were stored locally, and then forwarded to the a computing cluster at the National Energy Research Scientific Computing Center (NERSC) to be processed for offline data analysis.

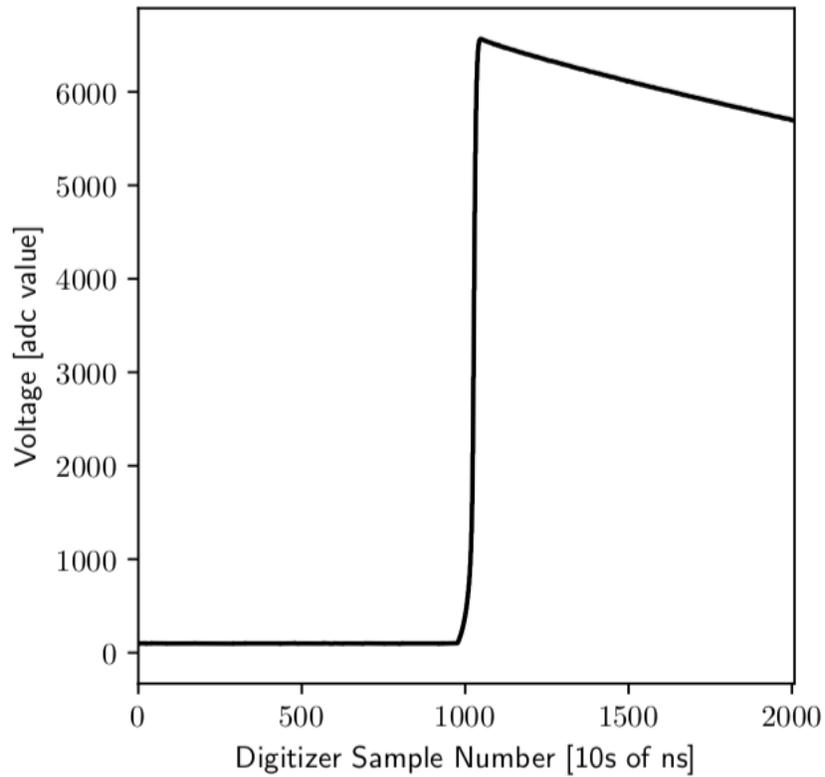


Figure 2.5: Waveform of a single-site gamma-ray physics event read out by a PPC detector and digitized. The initial shallow rise (imperceptible in this figure) from baseline accounts for the majority of the drift path the holes undertake. The sharp rise accounts for the final period of drift when the holes are within the innermost isochrone of Figure 2.4 relative to the point contact [32].

The DAQ system of the DEMONSTRATOR utilized an object-oriented real-time control and acquisition software application (ORCA) [35]. Written in objective C, ORCA is a highly-modular, general purpose, acquisition and control system that can be configured at run time to represent different hardware schema for any detection experiment [17]. Each object in ORCA maintains its own detailed data structures as well as its own set of diagnostic code that allows for the easy control of many DAQ system setups, including the DEMONSTRATOR's.

### 2.2.5: Calibration

The detectors were calibrated on a weekly basis by hour long  $^{228}\text{Th}$  line source runs and on a monthly basis by 24 hour long runs of the same source. The hour long calibrations measured gain stability in the detectors as well as the efficiency of Pulse Shape Discrimination (PSD) cuts that are discussed in section 2.5.2. The 24 hour long calibrations measured the time stability of these PSD cuts. The  $^{228}\text{Th}$  source was a line source that was deployed into the DEMONSTRATOR via PTFE tubing that wrapped in a helix around the cryostat and entered the module through a small cut-out in the shielding layers [36]. The helical position of the line source ensured that the majority of the solid angle with respect to the detectors was covered by the source radiation. When not in use, the line source was stored outside of the module on a mechanical apparatus controlled by the DAQ system. During deployment, a DAQ user activated a motor in the apparatus which inserted and retracted the source to a predefined stopping point.

Other calibrations took place throughout the DEMONSTRATOR's operation for a variety of studies. For example, a  $^{60}\text{Co}$  button source was placed in multiple locations to test the detector efficiency from hardware components distant from the cryostats.

## Section 2.3: Background Mitigation and Sensitivity to $0\nu\beta\beta$ Discovery

In order to be sensitive to the effective Majorana mass,  $0\nu\beta\beta$  experiments must limit background radiation to as low as possible levels. While the discovery sensitivity scales linearly with exposure in a background free environment, this relationship changes to a square root dependence when backgrounds are present. Therefore, any amount of background will lead to a diminished sensitivity.

This is shown clearly in Figure 2.6. The  $3\sigma$  discovery sensitivity of  $0\nu\beta\beta$  in  $^{76}\text{Ge}$  scales linearly with exposure as long as there is no background in the region of  $Q_{\beta\beta}$ . However, at even the background goal of the DEMONSTRATOR, just 1.0 count/FWHM $\cdot$ t $\cdot$ yr, this relationship diverges and requires around 100 ton $\cdot$ years of exposure to achieve a  $3\sigma$  discovery sensitivity that would only require 2 ton $\cdot$ years exposure in a background free case. This problem worsens when attempting to reach discovery sensitivities of higher half lives. Therefore,

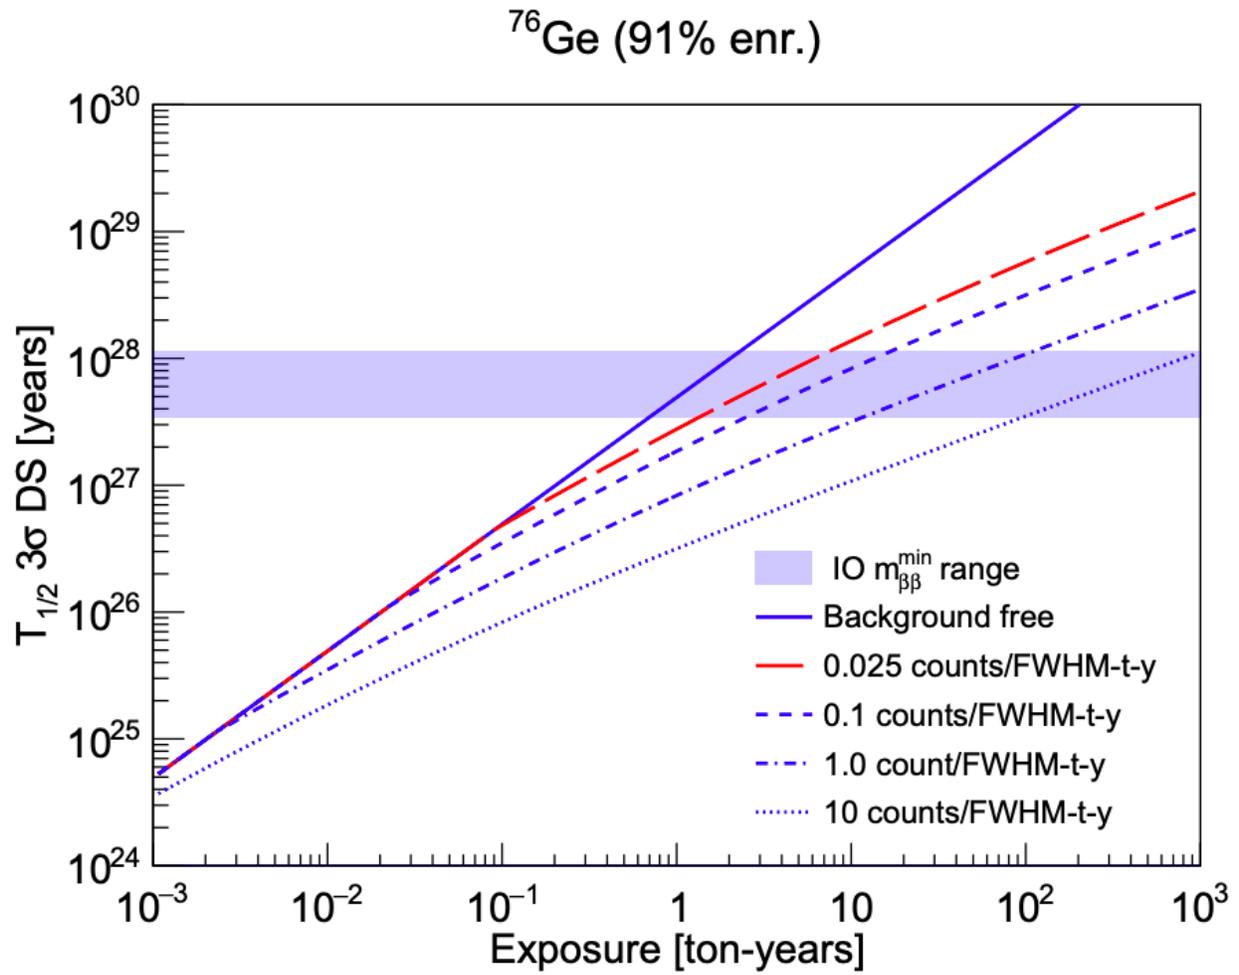


Figure 2.6: The sensitivity to a  $0\nu\beta\beta$  decay signal in  $^{76}\text{Ge}$  as a function of exposure and background for a  $3\sigma$  (99.7% C.L.) discovery sensitivity (DS). Note, the background rates are normalized to a 2.5 keV FWHM energy resolution. [37]

it was important for experiments like the MAJORANA DEMONSTRATOR and GERDA to develop background mitigation techniques prior to embarking on large scale  $0\nu\beta\beta$  searches. Based on what has been learned from the MAJORANA DEMONSTRATOR and GERDA, as well as from additional R&D for LEGEND-200, a background of  $<0.025$  counts/FWHM $\cdot$ t $\cdot$ yr is the goal for LEGEND-1000 [16].

## Section 2.4: Sources of Background and Background Mitigation in the DEMONSTRATOR

A signal from  $0\nu\beta\beta$  will be found at  $Q_{\beta\beta}$ , 2039 keV. The MAJORANA collaboration defines a “region of interest” (ROI) around this energy in which any count could potentially be identified as signal. The ROI has a width of 3.8 keV centered on  $Q_{\beta\beta}$  [15] which accounts for the energy resolution (2.52 keV FWHM) and the uncertainty in the  $0\nu\beta\beta$  peak shape. In order to correctly quantify the background index in the ROI, a larger window of 360 keV is used. This is done to account for events from the Compton continuum and degraded alphas that could contaminate the ROI (see section 2.5.1). This background estimation window (BEW) stretches from 1950 keV to 2350 keV, but excludes a 10 keV window centered on  $Q_{\beta\beta}$ , as well as three 10 keV windows around known gamma backgrounds (one from  $^{208}\text{Tl}$  and two from  $^{214}\text{Bi}$ )<sup>1</sup>. The MAJORANA DEMONSTRATOR implemented mitigation techniques to reduce common sources of background in this window, including cosmogenic muons, neutrons from muon spallation, alphas from radon plate-out, radiation from the decay of cosmogenic nuclides like  $^{60}\text{Co}$  and  $^{68}\text{Ge}$ , and radiation from the  $^{238}\text{U}$ ,  $^{232}\text{Th}$ , and  $^{40}\text{K}$  decay chains.

### 2.4.1: Cosmogenic Sources

Cosmic rays that strike the atmosphere produce a number of secondaries, including neutrons, charged pions that decay into muons, and neutrinos. Atmospheric neutrinos hardly interact with matter, and neutron secondaries do not have a sufficient energy and lifetime to penetrate the Earth’s surface deeper than a few meters. However, cosmogenic muons have a sufficient energy and lifetime to penetrate the Earth’s surface substantially. To reduce this muon flux by several orders of magnitude, the DEMONSTRATOR utilized a 4300 m.w.e. rock overburden. Even at this depth, there is still a measurable flux of muons that are energetic enough to produce counts near  $Q_{\beta\beta}$ . The DEMONSTRATOR accounted for this flux using an active muon veto system made of scintillator panels. A passing muon leads to scintillation photons in the panel that are detected by PMTs. The timing of the signal is recorded by the DAQ system and an exclusion window is placed in the data around the time of the muon passage.

<sup>1</sup>The ROI was used specifically with the Feldman-Cousins technique in combination with the BEW to produce a lower limit on the  $0\nu\beta\beta$  half-life. Other statistical techniques were used that did not include this 3.8 keV window, see [15].

Cosmic ray secondaries can also have an impact on HPGe detectors as they are transported and stored on the surface. Secondary neutrons and muons can interact with  $^{76}\text{Ge}$  to produce other isotopes including  $^{68}\text{Ge}$  and  $^{60}\text{Co}$ . These two isotopes decay and produce gammas with energies above  $Q_{\beta\beta}$ , thus potentially contributing to counts in the ROI [38]. The DEMONSTRATOR mitigated this process, known as “cosmogenic activation”, by limiting the amount of time the detectors spent on the surface. This involved storing them underneath some overburden after fabrication and quickly transporting the detectors to SURF.

Muons that reach the level of the DEMONSTRATOR have the potential to interact with the rock surrounding the lab, and muon spallation can occur. These interactions can lead to fast neutrons incident on the array with energies around  $Q_{\beta\beta}$ . To prevent these neutrons from reaching the detectors, polyethylene shielding moderated neutrons to lower energies. A layer of borated polyethylene underneath - polyethylene impregnated with boron - allowed for neutron capture of the moderated neutrons. Lead bricks installed behind the polyethylene layers blocked de-excitation gammas from this capture as well as gammas from the rock walls of the lab.

#### 2.4.2: Primordial Radiation Sources

$^{238}\text{U}$  and  $^{232}\text{Th}$  are extremely long-lived radioactive isotopes that are responsible for two prevalent decay chains (see figures 2.7 and 2.8). As primordial elements produced through stellar astrophysical processes,  $^{238}\text{U}$ ,  $^{232}\text{Th}$ , and  $^{40}\text{K}$  are abundant on Earth and consequently can make up part of the composition of everyday materials. Though radioactivity from these nuclei and their daughters is typically inconsequential in everyday life, it can form a significant background source for rare-event searches including the MAJORANA DEMONSTRATOR. To mitigate this risk, ultra-clean materials were used in components placed near the detector array. This included UGEFCu, which was electroformed in a separate laboratory space underground to eliminate cosmogenic exposure, and commercially produced plastic (PTFE/NXT-85) for plastic inserts involved in the detector support structure. Also, both decay chains feature short-lived isotopes of the noble gas radon ( $^{222}\text{Rn}$  via  $^{238}\text{U}$ ,  $^{220}\text{Rn}$  via  $^{232}\text{Th}$ ), which alpha decay. Isotopes above radon in the decay chains are abundant in the rock surrounding the lab, and so without proper consideration, radon daughters can accumulate on detector surfaces and emit alpha particles. Though such alphas have energies much higher than  $Q_{\beta\beta}$ , incomplete charge collection from them can potentially lead to counts in the ROI. To mitigate this alpha background, the DEMONSTRATOR featured a radon-exclusion box that fit between the muon veto and lead brick layers. The box was purged with low Rn nitrogen gas from a supply of liquid nitrogen, blocking passage of the radon nuclei.

Ultimately it is impossible to completely eliminate these isotopes and their progeny from materials.

## The Thorium-232 Decay Chain

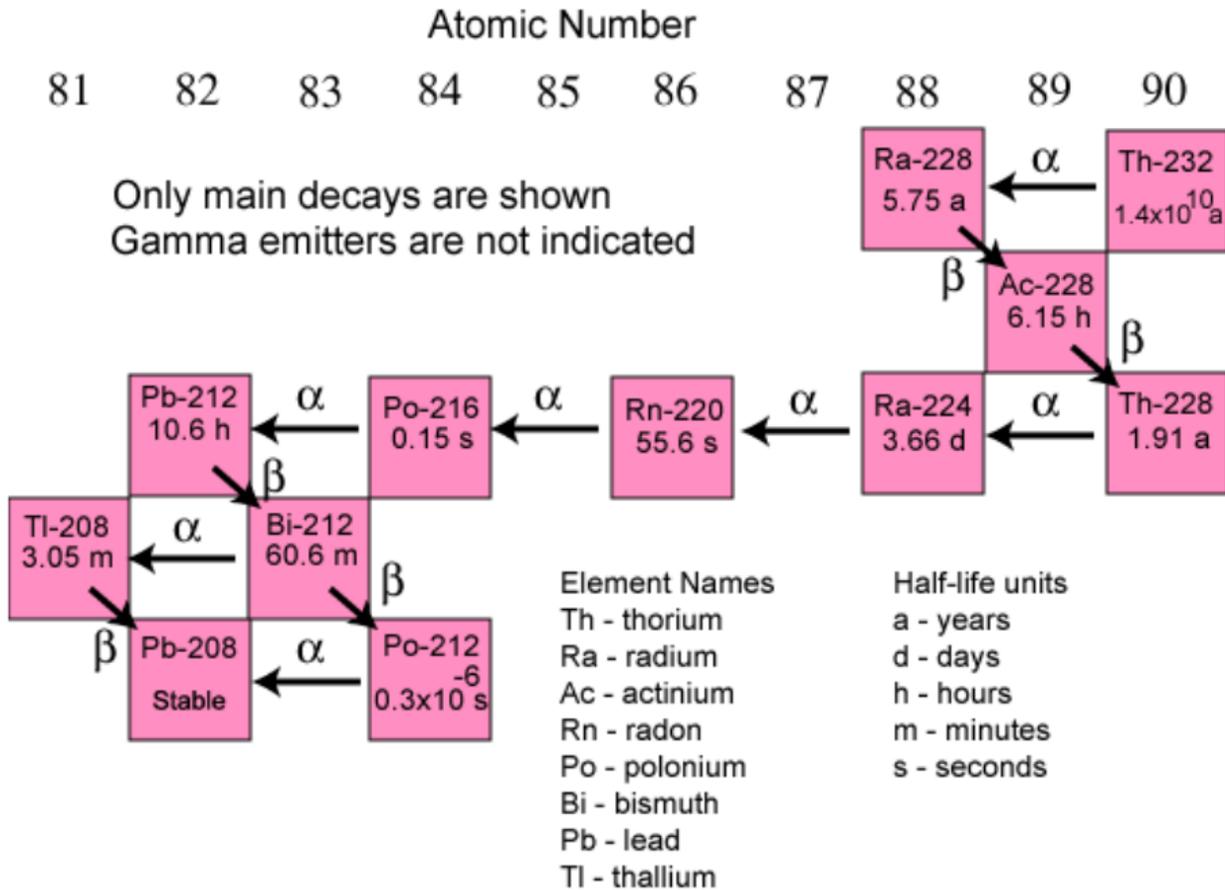


Figure 2.7: <sup>232</sup>Th decay chain. Notable decays important to background studies include the alpha decay of <sup>220</sup>Rn and beta decay of <sup>208</sup>Tl[39].

## The Uranium-238 Decay Chain

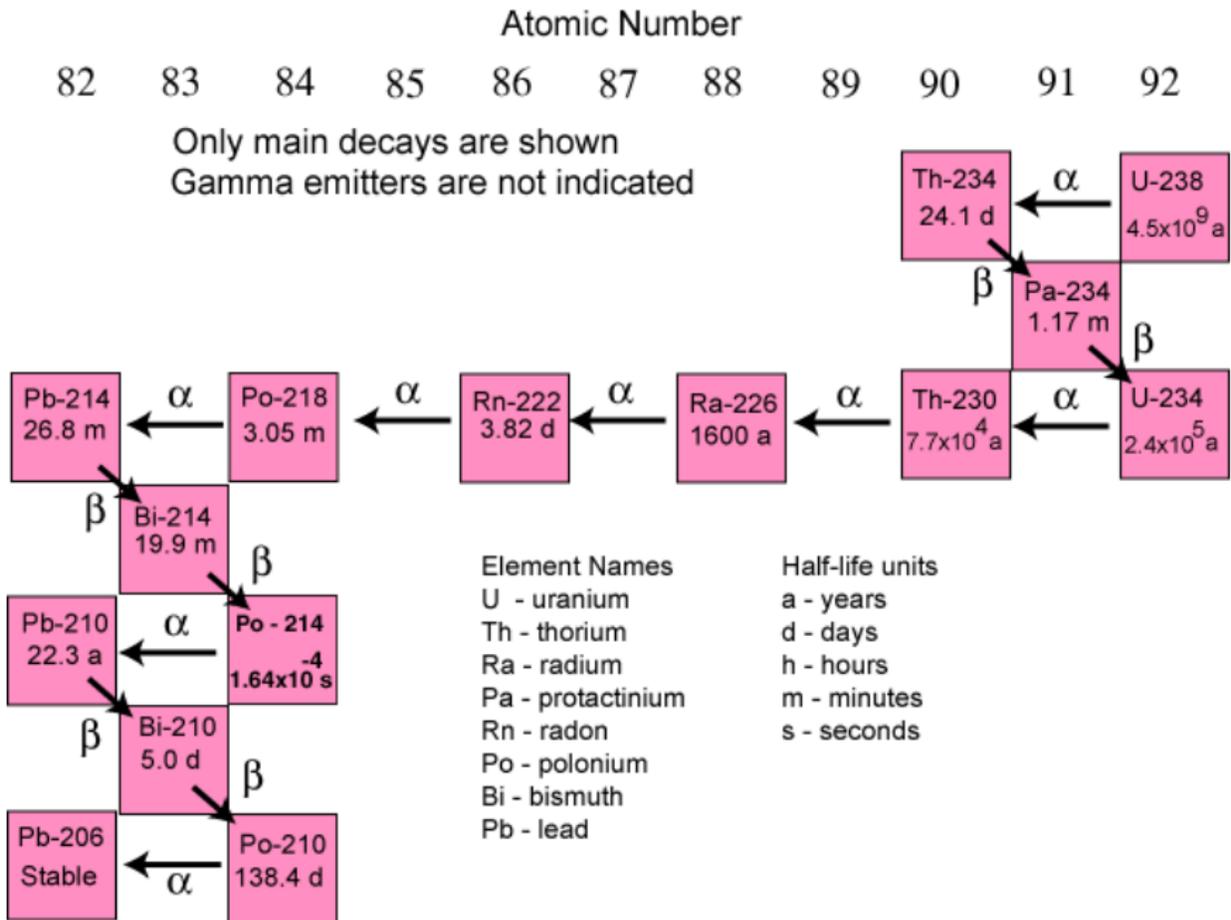


Figure 2.8: <sup>238</sup>U decay chain. A notable decay important to background studies is the alpha decay of <sup>222</sup>Rn[40].

In order to account for these radioactive sources in the DEMONSTRATOR, the collaboration undertook a rigorous assay program [41]. Samples of each material were analyzed through one of four methods -  $\gamma$ -ray counting, Glow Discharge Mass Spectrometry (GDMS), Inductively Coupled Plasma Mass Spectrometry (ICPMS), and Neutron Activation Analysis (NAA). The advantages and disadvantages of each method are described in [41]. Results for most assayed materials provided a value or upper-limit on their  $^{238}\text{U}$  and  $^{232}\text{Th}$  concentrations, and some also returned  $^{40}\text{K}$  concentration information. These results were combined to form the MAJORANA DEMONSTRATOR assay-based background model, to be described in chapter 4. In addition to the assay, strict procedures were written for the handling and installation of materials in order to reduce the chance of contamination.

## Section 2.5: Recognizing and Mitigating Backgrounds through Offline Analysis

### 2.5.1: Background Signatures in Data

The background distribution in the energy spectrum of the DEMONSTRATOR is constructed from gamma, alpha, and beta signatures that result from the sources discussed in section 2.4.2.

Gammas can interact in the germanium material in three different ways, depending on their incident energy [29]. At lower energies, an incident gamma will be absorbed by a germanium atom and a photoelectron will be emitted. This photoelectron will have energy

$$E_{e^-} = h\nu - E_b \tag{2.2}$$

where  $h\nu$  is the energy of the incident gamma ray and  $E_b$  is the binding energy of the electron in the germanium atom. Since the interaction requires the incident gamma to have a higher energy than the electron binding energy of a particular electron shell for the interaction to take place, the energy spectrum will exhibit edges at which the frequency of gamma events will decrease dramatically below the edge. As shown in Figure 2.9, this interaction is only dominant in germanium below 200 keV and is typically not a danger regarding counts in the ROI. However, its signature is important for low energy BSM physics searches that the DEMONSTRATOR carried out.

Another important contributor to the spectrum is pair production, which is only possible with gammas of the property  $h\nu > 1.02$  MeV. This interaction is rare for most identified background sources in the DEMONSTRATOR, as pair production does not become the dominant process in germanium until above 8 MeV. In pair production, a gamma produces an electron and positron pair. The electron slows and is absorbed by the germanium, and the positron re-combines with another electron to produce two gammas of

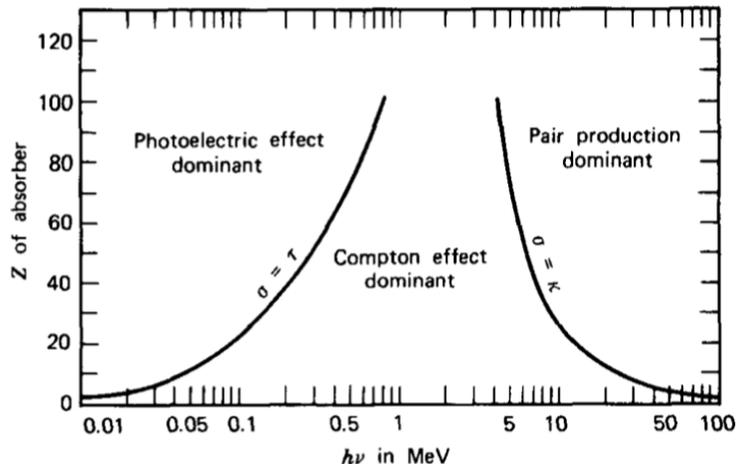


Figure 2.9: A graphic that illustrates the dominant gamma ray interaction in the parameter space defined by the energy of the incident gamma ray (in MeV) and the proton number of the absorbing material ( $Z=32$  for germanium). The black lines are points in the parameter space where the probabilities of two of the interactions are equivalent. [29]

the property  $h\nu = 0.511$  MeV. One, both, or neither of these gammas is absorbed by the germanium, and these cases can produce a single-escape peak (SEP), double-escape peak (DEP), or a full-energy peak (FEP). The most noteworthy gamma that pair produces in the DEMONSTRATOR is that from the  $^{208}\text{Tl}$  decay, which produces a FEP at 2615 keV and the appropriate relative DEP and SEP. These three peaks are used to calibrate the multi-site event (AvsE) cut, discussed in section 2.5.2.

The most notable gamma interaction in the DEMONSTRATOR and the one with the highest potential to contribute to counts in the ROI is Compton scattering. In this interaction, a gamma scatters off an electron in the germanium and transfers some energy. The energy of the departing gamma is given by:

$$h\nu' = \frac{h\nu}{1 + \frac{h\nu}{m_0c^2}(1 - \cos\theta)} \quad (2.3)$$

where  $m_0c^2$  is the rest mass energy of the electron, and  $\theta$  is the scattering angle of the departing gamma. The energy of the recoil electron can be deposited into the germanium and show up in the spectrum. We can obtain this energy by rearranging equation 2.3:

$$E_{e^-} = h\nu - h\nu' = h\nu \left( \frac{\frac{h\nu}{m_0c^2}(1 - \cos\theta)}{1 + \frac{h\nu}{m_0c^2}(1 - \cos\theta)} \right) \quad (2.4)$$

Background counts from the recoil electron form the Compton continuum that spans from  $E_{e^-} = 0$  to the Compton edge, where  $E_{e^-}$  is maximized at  $\theta = \pi$  and the gamma is backscattered. Compton scattering is a particular nuisance because gammas at energies above the Q-value can produce recoil electrons that potentially deposit energy into the germanium within the ROI energy range.

Coincident gammas from particular isotope decays can deposit energy in the detectors within the event time window, leading to energy peaks that are really the sum of two incident gammas. The energy peaks are known as “summed peaks” and are important to consider when interpreting the energy spectrum.

Alpha particles are a concern if they are incident on the passivated surface of the germanium detector - the surface shared with the point contact which is not lithiated. It is here that holes from an alpha energy deposition can become trapped and induce a charge at the point contact at a later time, leading to incomplete charge collection and a possible count in the ROI. This possibility is mitigated by the Delayed Charge Recovery (DCR) cut, discussed in the section 2.5.2.

Finally, beta particles cannot penetrate deeply into most materials, and thus they are only a concern when emitted very close to or within the detectors. They appear in the energy spectra with the characteristic beta decay distribution where the energy of the decay is shared between the antineutrino and electron. Beta particles can also indirectly effect the energy spectrum through bremsstrahlung, especially high energy betas in high-Z materials such as the lead shield.

The most prominent beta decay distribution in the energy spectrum of the DEMONSTRATOR is naturally that of  $2\nu\beta\beta$  from the enriched detectors. Due to the low background the DEMONSTRATOR was able to achieve up to  $Q_{\beta\beta}$ , this irreducible background is clearly visible and its spectral shape can be used to extract an updated estimate of the  $2\nu\beta\beta$  rate (see section 6.6). Counts from  $2\nu\beta\beta$  dominate the spectrum in the 1000-1400 keV range. Above this range, the number of counts from  $2\nu\beta\beta$  decreases up to the lower bound of the BEW, where the overall background becomes flat <sup>2</sup>. Therefore, counts from the  $2\nu\beta\beta$  spectrum make a negligible contribution to the BEW.

### 2.5.2: Pulse Shape Discrimination (PSD)

Offline analysis cuts and Pulse Shape Discrimination (PSD) are used to mitigate backgrounds that are not otherwise addressed by the selection of materials and the shielding configuration. Offline analysis cuts refer to timing cuts that track the refilling of liquid nitrogen (which tends to produce low-energy noise in detectors), multiplicity cuts that remove any multi-detector events from the analysis (as these events are not possible  $0\nu\beta\beta$  candidates), and data-cleaning cuts that search for and remove waveforms with non- $0\nu\beta\beta$  characteristics.

PSD is used in the form of three cuts - the “AvsE” cut that rejects multi-site events, the Late Charge (“LQ”) cut that rejects a population of multi-site events the AvsE cut tends to miss, and the Delayed Charge Recovery (“DCR”) cut that rejects alpha events. As mentioned in section 2.2.3, HPGGe detectors

---

<sup>2</sup>Counts from the  $2\nu\beta\beta$  spectrum are still decreasing up to  $Q_{\beta\beta}$ , but the amount is negligible compared to other sources.

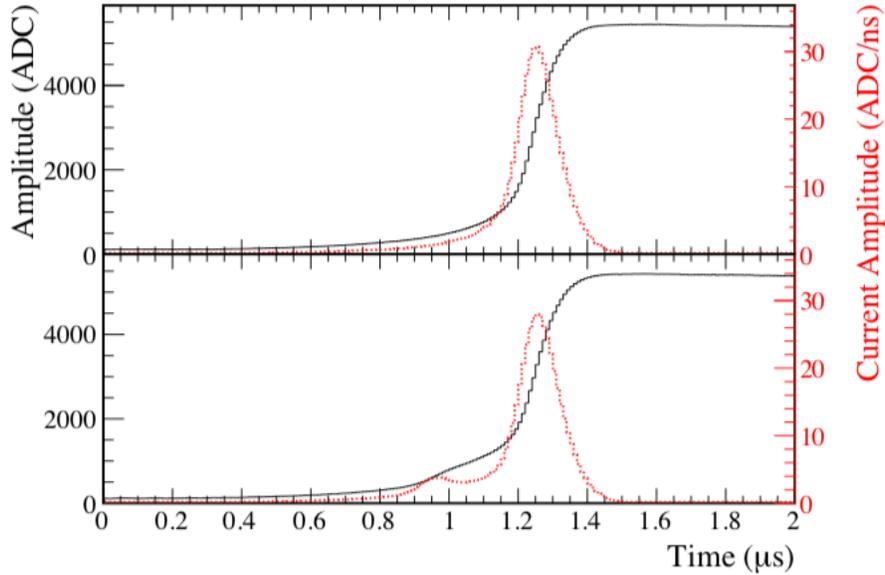


Figure 2.10: Comparison of a single-site gamma-ray waveform (top) and a multi-site waveform (bottom). The multi-site current waveform (bottom, red) clearly shows several peaks of lower amplitude which the AvsE cut can detect. [42]

have excellent characteristics for PSD. The variant strength of the weighting potential in different parts of the crystal means that a multi-site event (multiple energy depositions in the crystal) will appear in a current vs. time waveform as multiple peaks of smaller amplitudes (A) (see Figure 2.10). By comparing the maximum current waveform amplitude A to the energy of the event (E), a waveform can be categorized as single-site ( $0\nu\beta\beta$ -like) or multi-site. The algorithm that does this, the AvsE cut, is tuned through  $^{228}\text{Th}$  calibration runs to accept 90% of the DEP, which contains single-site events. The specific technique which compares A to E and additional properties of this cut are discussed thoroughly in [42].

The AvsE cut unfortunately has trouble cutting multi-site event populations where one event occurs close to the point contact. In this case, the maximum current amplitude appears large due to the proximity of the event and passes the cut, even though the charge vs. time waveform has the characteristic shape of a multi-site waveform. These events are handled by the complimentary LQ cut, which measures the area bounded by the waveform itself (in the charge domain), the rising edge of the waveform beginning at 80% of the waveform’s amplitude, and the mean value of the flat-top of the waveform (see Figure 2.11). A larger area corresponds to a multi-site event and is subsequently cut.

The DCR cut is sensitive to alphas plated-out on the passivated surface of detectors through radon contamination. Such waveforms have a characteristically shallow tail due to the slow drift time of charges as they become trapped and released near the detector surface (see Figure 2.12). Analysts tag and reject these waveforms using the appropriately named “Delayed Charge Recovery” (DCR) cut. A detailed analysis

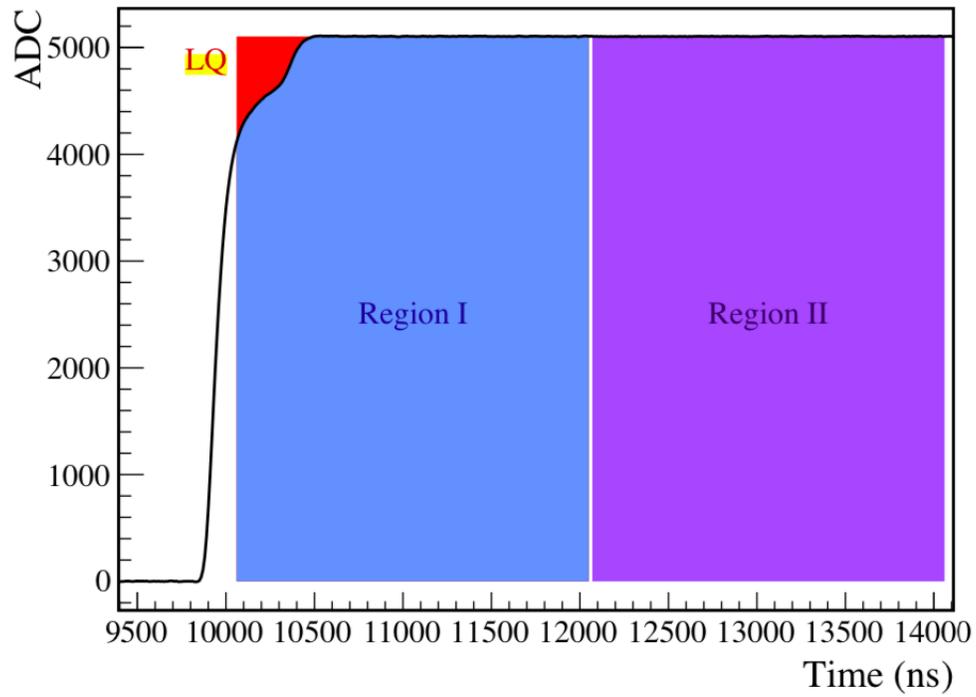


Figure 2.11: Illustration of the regions used to make a measurement for the LQ cut. The area in red determines whether the event will pass the LQ cut, with larger areas tending to fail the cut. Regions I and II are used to numerically estimate the area in red. [43]

of this cut is thoroughly discussed in [44].

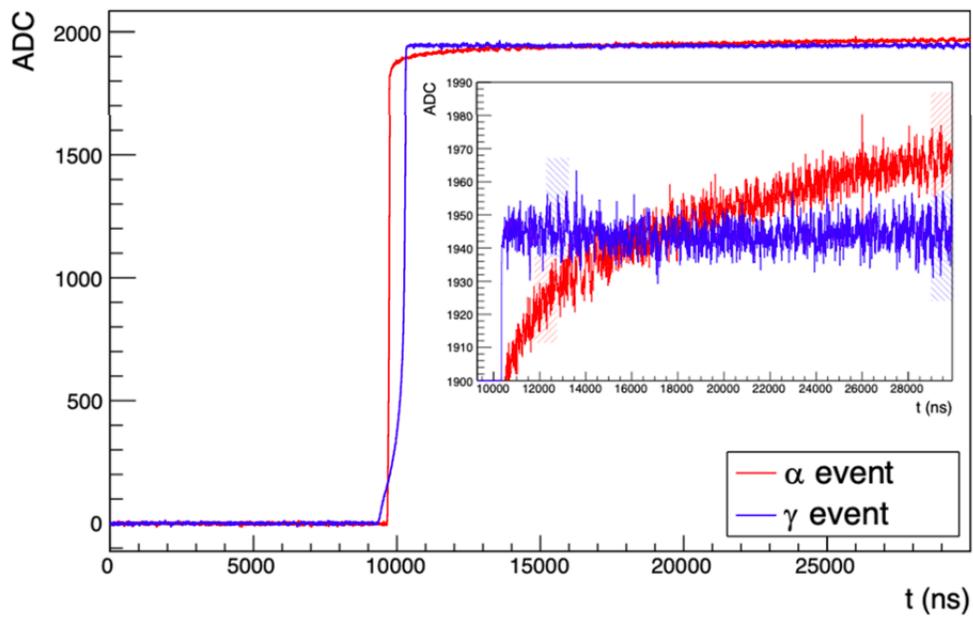


Figure 2.12: A waveform from a  $\gamma$  event in blue and a waveform from an  $\alpha$  event in red.  $\alpha$  events tend to have a shallow rising tail compared to  $\gamma$  events, which is tagged by the DCR cut. [44]

## CHAPTER 3: Spectra of the DEMONSTRATOR

The data and resulting energy spectra collected by the DEMONSTRATOR provide a window into the sources present in materials, with signatures discussed in section 2.5.1. The data is organized into sets that correspond to various configurations of shielding, detectors, and data collection, as well as time since detectors were moved underground. Partitioning the data in this way provides additional information that can be used to track decaying sources of radiation (cosmogenics), measure contaminant radiation in configurations with reduced shielding, and pinpoint sources of excess radiation.

### Section 3.1: Data Sets

In June 2015, commissioning data were taken with only Module 1 installed in the south shield position. This first dataset, DS0, was collected before the installation of the inner copper shield. A blank monolith was installed in the east shield position, the eventual location of Module 2. Following rework of the M1 strings and installation of the inner copper shield, DS1 started in December 2015. In May 2016, a test was conducted on the GRETINA digitizers to allow for “pre-summing”, which extends the digitization window of the waveform baseline and tail. Data collected during this test formed the entirety of DS2, and pre-summing was disabled at the conclusion of the dataset.

In August 2016, Module 2 replaced the blank monolith in the east shield position. Module 1 and Module 2 initially ran with two separate DAQ systems, producing two separate datasets: Module 1 data for DS3 and Module 2 data for DS4. In October 2016 these two separate DAQs were merged together, initiating DS5.

DS5 was divided into three sub-datasets. DS5a accounted for the first few months of combined DAQ data collection and was characterized by higher electronic noise and the final installation of the polyethylene shield during this period. Adjustment of the module and grounding configurations brought a period of low-noise, DS5b, which ran for a few months after DS5a. DS5c comprised the final two months of DS5, in which a blindness scheme was formally implemented, to be described in section 3.2.

DS6 commenced in May 2017 at which time the GRETINA digitizer pre-summing feature was re-enabled. DS6 was split again into three subsets. DS6a and DS6b had identical configurations, running for 11 months and 7 months respectively. In November 2018, a failed GRETINA card was removed from Module 2. Signal cables that were going to this card were transferred to a different, functioning GRETINA card. Data taken

following this rearrangement comprised DS6c.

In November 2019, a cable and connector upgrade commenced on Module 2, while Module 1 continued to operate. Data from Module 1 during this period comprised DS7. Completion of the upgrade in August 2020 reintroduced Module 2 with a larger total active detector mass and four new ICPC detectors (see Appendix A). These data formed DS8, which was split into “DS8 (PPC)” and “DS8 (ICPC)” to differentiate the exposure contributions from the two different types of detectors.

March 2021 signaled the end of the DEMONSTRATOR  $0\nu\beta\beta$  program. All enriched detectors were removed and shipped to Laboratori Nazionali del Gran Sasso (LNGS) in Italy for use in LEGEND-200. The remaining natural detectors were re-stringed and consolidated into Module 2 to run as DS9 for additional measurements and tests related to background estimation.

In April 2022, tantalum disks were installed in M2 in place of several natural detectors in an effort to measure the decay rate of  $^{180m}\text{Ta}$ . This experiment, outside of the scope of this dissertation, forms the basis of the currently running DS10.

A summary of each dataset, date range, enriched exposure, and change notes is shown in table 3.1.

### **Section 3.2: Selection of Data for the Final $0\nu\beta\beta$ Result and Background Analysis**

A large portion of the data were concealed by a blindness scheme in order to reduce the possibility of a biased result. This blindness scheme left 25% of the data “open” in order to monitor data stability and provide data for preliminary analyses. The “open”/“blind” label was assigned to runs on a 124 hour cycle which ran for 31 hours in “open” mode and 93 hours in “blind” mode. Calibration data and non-standard run data were labeled as “open”. Data below 100 keV and events with a multiplicity greater than 1 were labeled “blind” such that they could be handled differently for the low energy and  $0\nu\beta\beta$  excited state decay analyses [45].

When presenting a background spectrum in the context of a final  $0\nu\beta\beta$  result, the spectrum is presented in the lowest background configuration with only enriched detectors. This configuration excludes DS0, during which the inner copper shield was not installed, and previously excluded DS5a, which featured excess electronic noise. Analysis efforts to clean the data in DS5a were successful, and this data is now included in the final background configuration of DS1-DS8.

While not used in the final background spectrum for the  $0\nu\beta\beta$  result, DS0 is important for background studies because of the reduced shielding between the lead layer and the detectors. The data from this time period constrains the lead contribution to the background model. Similarly, energy spectra from natural detectors are critical for background studies as they can account for contributions in the  $2\nu\beta\beta$

Data Set	Start Date	End Date	Enr. Exposure (kg y)	Notes
DS0	06/15	10/15	$1.12 \pm 0.02$	M1 commissioning, incomplete shield
DS1	12/15	05/16	$2.26 \pm 0.03$	M1 only, inner shield installed
DS2	05/16	07/16	$1.12 \pm 0.02$	M1 only, pre-summing test
DS3	08/16	09/16	$0.97 \pm 0.01$	M1+M2 in shield, M1 DAQ
DS4	08/16	09/16	$0.26 \pm 0.00$	M1+M2 in shield, M2 DAQ
DS5a	10/16	01/17	$2.72 \pm 0.04$	DAQs combined, high noise, poly shield installed
DS5b	01/17	03/17	$1.77 \pm 0.03$	Grounding optimized to reduce noise, poly shield complete
DS5c	03/17	05/17	$2.17 \pm 0.03$	Blindness scheme implemented
DS6a	05/17	04/18	$14.76 \pm 0.21$	Blindness, pre-summing
DS6b	04/18	11/18	$9.78 \pm 0.14$	Blindness, pre-summing
DS6c	11/18	11/19	$13.26 \pm 0.19$	Blindness, pre-summing, GRETINA digitizer swap
DS7	11/19	08/20	$4.47 \pm 0.07$	Blindness, pre-summing, M1 only
DS8(PPC)	08/20	03/21	$6.99 \pm 0.09$	Blindness, pre-summing, M2 upgraded, PPCs only
DS8(ICPC)	08/20	03/21	$2.82 \pm 0.04$	Blindness, pre-summing, M2 upgraded, ICPCs only
DS9	03/21	04/22	n/a	M2 only, natural detectors only, <b>post <math>0\nu\beta\beta</math> program</b>
DS10	04/22	present	n/a	M2 only, tantalum disks installed, <b>post <math>0\nu\beta\beta</math> program</b>
<b>TOTAL ENR. EXPOSURE (kg y)</b>			$64.46 \pm 0.92$	DS0-8 (2015-2021)

Table 3.1: Table of MAJORANA DEMONSTRATOR dataset information.

region that would otherwise be undetectable by only enriched detectors. Natural detector spectra from DS9 are particularly important, as they can be compared to spectra of earlier datasets where detectors were in alternate positions, providing additional insight into particular sources of radioactivity.

Events in the lowest background configuration are generally from high-gain channels, though if the event only has a detection in the low-gain channel, that channel is used. The lowest background configuration additionally incorporates channel selection, which excludes events from detectors which do not pass a set of data quality criteria in a given period of time. This selection is done on a detector-by-detector basis within a dataset, meaning a high-gain channel and low-gain channel for the same detector will be assigned the same tag. These criteria eliminate detectors with issues in  $^{228}\text{Th}$  calibrations, event building, data cleaning, pulser

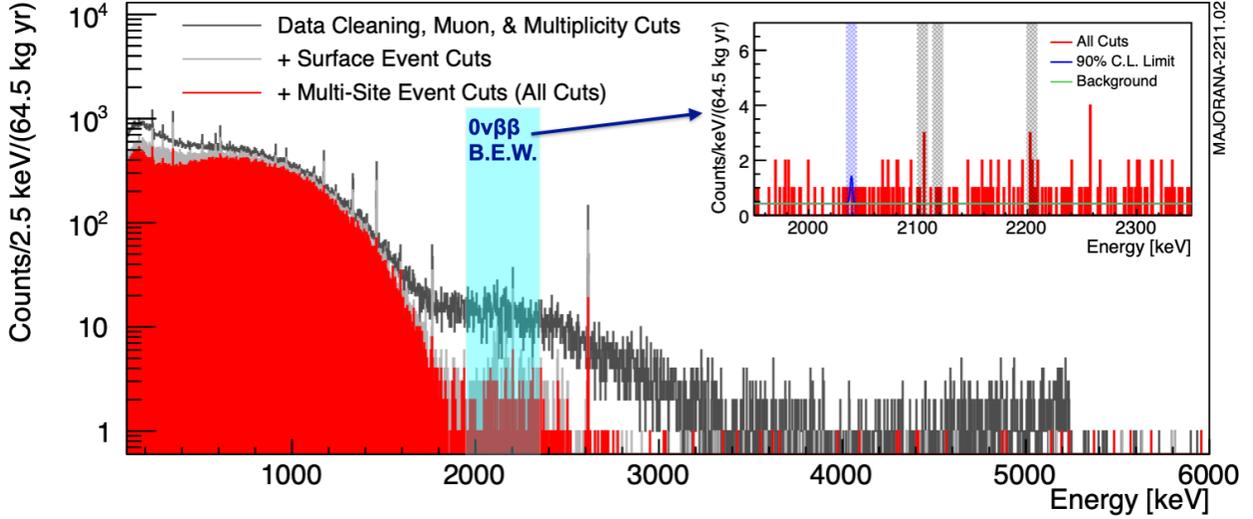


Figure 3.1: The energy spectra, before and after cuts, and results from 64.5 kg-yrs enriched detector operation of the DEMONSTRATOR. The background estimation window is highlighted in light blue and shown in the inset. The blue vertical band around 2039 keV in the inset is the 10-keV window around  $Q_{\beta\beta}$ , while the gray vertical bands in the inset are the 10-keV windows around peaks from the decays of  $^{208}\text{Tl}$  and  $^{214}\text{Bi}$ . These four vertical bands are excluded from the BEW. Figure adapted from the data in Ref. [15]

timing, and poor PSD performance or stability.

### Section 3.3: Final Results of the MAJORANA DEMONSTRATOR

As mentioned in section 2.4, a window surrounding  $Q_{\beta\beta}$  is used as a proxy to estimate the background rate in the lowest background configuration. The 360-keV-wide BEW ranges from 1950 keV to 2350 keV and excludes four 10-keV-wide regions: one around  $Q_{\beta\beta}$ , and three around prominent  $\gamma$ -peaks from the decays of  $^{208}\text{Tl}$  and  $^{214}\text{Bi}$  (see Figure 3.1).

An assay-based background model, produced through Monte-Carlo simulations of the DEMONSTRATOR and the radioassay program discussed in section 2.4.2, initially predicted a rate of  $8.75 \times 10^{-4}$  cts/(keV·kg·yr) in the BEW [41]. An updated Bayesian framework to produce an improved assay-based background model using Monte-Carlo uncertainty propagation returned a predicted rate of  $(1.16 \pm 0.04) \times 10^{-3}$  cts/(keV·kg·yr) in the BEW [46]. The construction of this assay-based model will be detailed in chapter 4.

Analysis of data from the enriched detectors produced a final background rate in the BEW of  $(6.23^{+0.55}_{-0.52}) \times 10^{-3}$  cts/(keV kg y) from the lowest background configuration spanning 63.3 of the  $64.5 \pm 0.9$  kg-yrs active exposure (see Figure 3.1) [15]. A  $0\nu\beta\beta$  signal was not observed, and a lower limit on the  $0\nu\beta\beta$  half-life was derived using an unbinned, extended profile likelihood method, producing a limit of  $T_{1/2} > 8.3 \times 10^{25}$

yrs<sup>1</sup>.

There is more than a factor of five difference between the final background rate and the updated predicted rate from the assay-based model. The answer for this discrepancy lies in the <sup>232</sup>Th decay chain, as shown in Figure 3.2. The measured rate in the prominent <sup>232</sup>Th chain  $\gamma$ -ray peaks, particularly at 238 keV and 2615 keV, exceeds the predicted values by about a factor of two and five respectively. Some of the discrepancy could also be contributed to the <sup>238</sup>U chain, as many of the prominent <sup>238</sup>U  $\gamma$ -ray peaks have a factor of two to three more counts in data than predicted (see table 7.1).

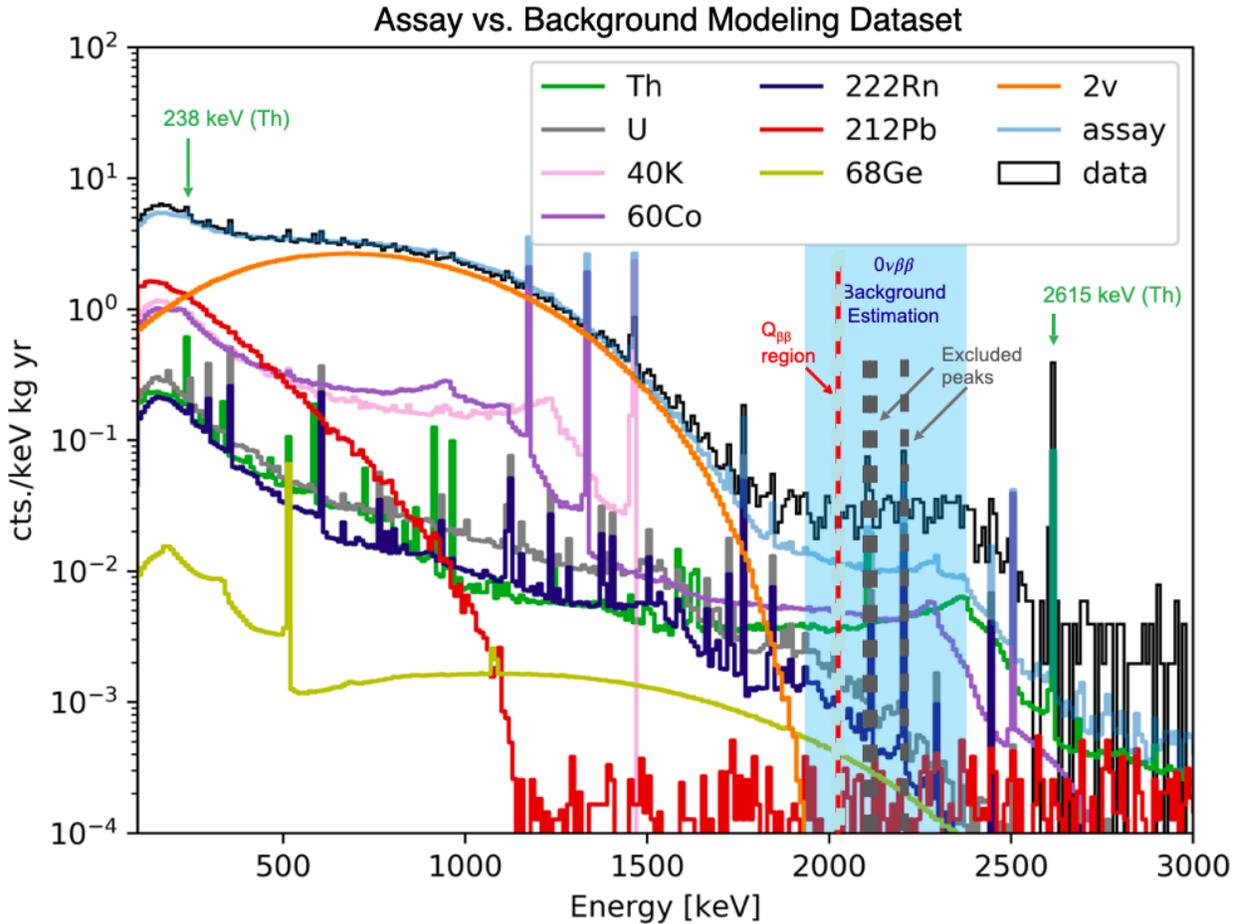


Figure 3.2: Predicted energy spectrum of the assay-based background model (blue curve) plotted against a 50.83 kg-yr subset of the enriched-detector data (black). The model has been broken down into its decay chain constituents, represented by other colors. Note that while <sup>57</sup>Co and <sup>54</sup>Mn are also contributors to the assay-based model, their spectra fall below the y-scale of this figure and were thus excluded.

In Module 1, the rate in the lowest background configuration is  $18.6 \pm 1.8$  cts/(FWHM t y) (or  $7.38 \pm 0.71 \times 10^{-3}$  cts/(keV kg y)), whereas in Module 2 the rate is only  $8.4_{-1.7}^{+1.9}$  cts/(FWHM t y) (or  $3.33_{-0.7}^{+0.8} \times 10^{-3}$  cts/(keV kg y)).

<sup>1</sup>A  $0\nu\beta\beta$  half-life lower limit was alternatively derived through the Feldman-Cousins technique using the 3.8 keV ROI window, producing a result of  $T_{1/2} > 7.2 \times 10^{25}$  yrs

cts/(keV kg y)). This difference indicates the existence of an additional source of background in the total rate that is non-uniform and larger in Module 1. Evidence for a localized excess has been identified in prominent  $^{232}\text{Th}$   $\gamma$ -peaks in the natural spectra from two detectors in Module 1.

### Section 3.4: Evidence in Module 1 Natural Detector Spectra for a $^{232}\text{Th}$ Excess

The top detector (D1) in each string sits at an elevation that is just below the opening of the hollow copper crossarm where it interfaces with the cryostat. The two strings closest to this opening in both Module 1 (C1) and Module 2 (C2) are strings 2 (P2) and 7 (P7)<sup>2</sup> (see Figure 2.2).

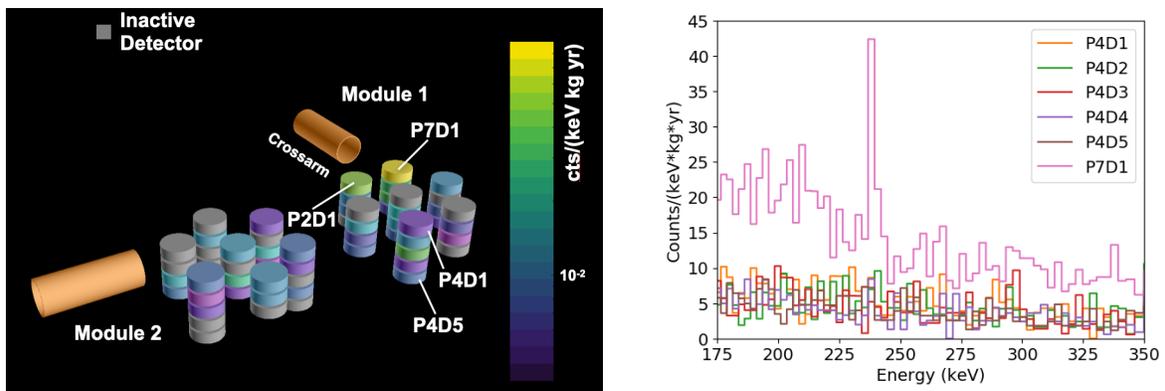


Figure 3.3: Direct evidence for a non-uniform  $^{232}\text{Th}$  excess in Module 1 natural detectors, based on rates of  $^{232}\text{Th}$ -chain  $\gamma$ -ray peaks. **Left:** Rate of the 2615-keV peak in individual detectors through datasets before the removal of enriched detectors. Detectors C1P2D1 and C1P7D1 have the highest rates. Detectors in gray were not operational during this period. **Right:** Rate of the 238-keV peak in operable Module-1 natural detectors through datasets before the removal of enriched detectors. Detector C1P7D1 has a much higher rate than the other detectors. Detector C1P2D1 is excluded due to limited statistics as it was biased down early in the dataset.

In datasets prior to DS9, natural detectors C1P2D1 and C1P7D1 observed an excess in the 2615-keV rate compared to other natural detectors in Module 1 (see Figure 3.3 left). This suggests that a localized source of  $^{232}\text{Th}$  existed in a component inside or near the crossarm opening of Module 1. Furthermore, detector C1P7D1 observed an excess in the 238-keV rate compared to other natural detectors (see Figure 3.3 right). The 238-keV  $\gamma$ -line is attenuated by more than 70% in 1 cm of copper, which suggests that minimal shielding exists in the line of sight between the source of the excess and C1P7D1.

These rates suggested that the excess either was located in the region of the crossarm opening or was integral to the crystals themselves. When enriched detectors were removed from the DEMONSTRATOR at

<sup>2</sup>When identifying these detectors, the individual IDs for the module, string, and elevation of the detector are used together. The elevation ID convention is as follows: D1 for the highest elevation in the string and DN for lowest elevation in the string, where N is the number of detectors in the string. As an example, the top detector (D1) in the Module 1 (C1) string located in the center of the cryostat (P1) is named C1P1D1.

the end of DS8, the natural detectors from Module 1 were installed into Module 2. C1P7D1 and C1P2D1 in particular were placed in new locations within Module 2. When observing the integrated rate of these two detectors in datasets when they operated in Module 1 versus Module 2, their overall rate fell to a comparable level with the other natural detectors in DS9 (see Figure 3.4). This measurement eliminates the possibility that the excess originates in these two detectors. The search for the source of the  $^{232}\text{Th}$  excess in the region of the crossarm opening will be discussed in section 4.4.

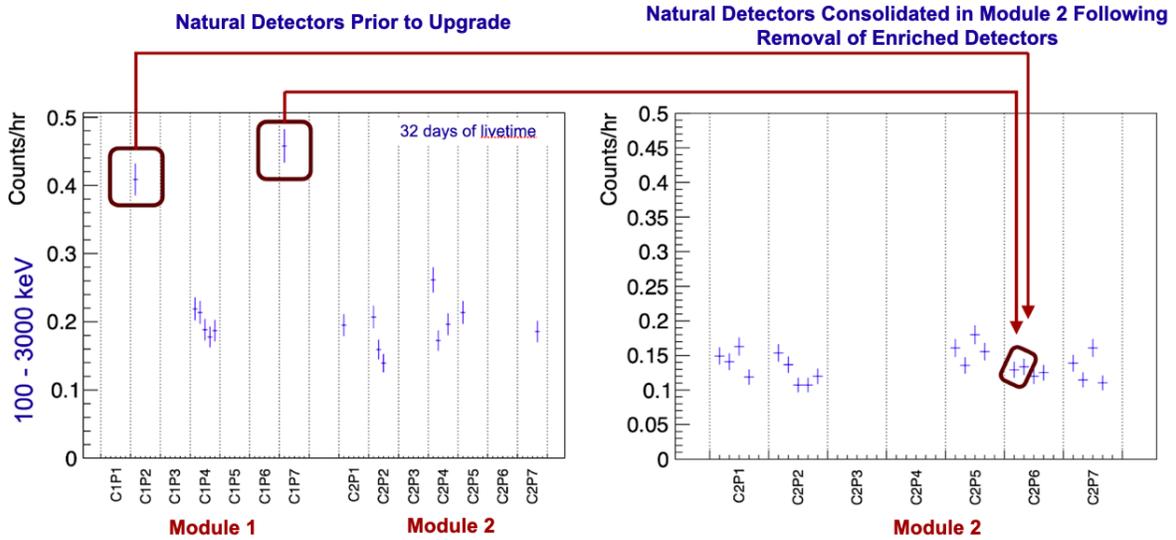


Figure 3.4: The two natural detectors which saw an excess of events between 100 and 3000 keV when positioned next to the Module 1 crossarm no longer had elevated rates after being moved to a new location in Module 2.

## CHAPTER 4: Simulations of the DEMONSTRATOR

The background model, whether fixed to radioassay measurements or constructed from Bayesian inferences, requires Monte-Carlo radiation transport simulations. Simulations of the DEMONSTRATOR are conducted using the Geant4 application `MaGe` [47][48], a program written explicitly for the MAJORANA DEMONSTRATOR and GERDA experiments. `MaGe` reproduces the geometry and materials of over 4000 components used in the DEMONSTRATOR. Conducting simulations of primary decays from these components provide detector efficiencies for every part of the apparatus. Provided with radioassay data or inferred activities as well as exposure values from relevant datasets, these detector efficiencies (counts/decay) can be converted into background indices (counts/keV·kg·yr).

### Section 4.1: Component Groupings

While it is feasible to conduct a separate simulation for every component and every potential decay chain observed by the DEMONSTRATOR, attempting to fit 4000 components  $\times$  several decay chains to data would be challenging given the large number of parameters (activities) to float and limited statistics. A more reasonable approach is to group components of the same material together, requiring that their array-wide detector efficiencies are relatively degenerate. This requirement has been extended further than in previous works, where electroformed copper components were grouped together regardless of location, for example.

A full table of component groupings, their corresponding components, and simulated decay chains used in this dissertation are shown in Appendix B. In general the component groupings are split by module in order to capture variations in the model between the two arrays, though groupings with components that span both modules (like the copper and lead shielding) are not split.

In general component groupings are simulated with an entire decay chain, though long decay chains like  $^{232}\text{Th}$  and  $^{238}\text{U}$  are split up into segments, where each segment represents a series of decays that are assumed to be in secular equilibrium. A summary of each decay chain and their corresponding segments can be found in Appendix C.

By default, component groups are simulated with 37.5 million primaries of a particular decay-chain segment. Only a subset of these primaries are used for component groupings with high detector efficiency (i.e. the LMFES) to reduce disk space usage. Conversely, component groups with low detector efficiency

(i.e. the bellows and lead shield), are simulated with about 3 billion primaries in order to obtain enough statistics in the resulting efficiency spectrum.

For most components, the primaries are generated within the bulk of the simulated component volume and decay according to an isotropic primary decay generator (`RDMiso`). This generator will allow a primary and its progeny to decay to a predetermined end-point (usually a stable isotope or an isotope that ends a primordial element decay segment) within the time window of one day. There are a few exceptions. For the Rn-purge shielding layer and detector unit PTFE (`DUPTFE`) parts, primaries are populated on the surface of the component volume in addition to the bulk. This is done using the `MaGE` generator `GSS` (Generic Surface Sampler), which ensures a random and uniform distribution of primaries on surfaces. For the lead shield volume, a generator that more accurately simulates the bremsstrahlung effect of betas passing through the lead is used rather than the standard isotropic primary decay generator. Finally, the bellow volumes, which represent thin-walled stainless steel components that isolate the cryostat and internal components vibrationally from the rest of the cryogenic and vacuum systems, have a small detector efficiency because they sit on the opposite end of the crossarms from the cryostat. Therefore, simulations of the bellows use a photon generator (`g4gun`) which emits  $^{232}\text{Th}$  decay chain gammas, energy weighted by intensity, in a cone from a random point in the bulk volume towards the detectors. The number of decays is corrected at the simulation processing stage in order to match the number of decays for an isotropic primary decay generator.

## Section 4.2: Processing with GAT and MPP and PDF Generation

The raw simulation file will contain un-normalized hit data, much like raw digitized data that is sent to NERSC (see section 2.2.4). In order to analyze the simulated data like a normal dataset, it must be processed with the same software that processes the digitized `DEMONSTRATOR` data. Originally this was done with the Germanium Analysis Toolkit (`GAT`). `GAT` sourced a database of detector information organized by dataset in order to produce simulation files for each configuration. These files, like their data counterparts, include parameters from offline data analysis, like corrections from energy non-linearity analysis and information from pulse shape discrimination.

A processing application designed for `LEGEND`, the `MaGe-Post-Processor` (`MPP`), was made backwards compatible with `DEMONSTRATOR` data and has become the preferred processor for simulation data. It runs with the same features as `GAT`, but also with additional features that produce systematic parameters important for the background model analysis. This includes measurements of the detector dead layer (the bulk region of the germanium crystal near the surface that cannot be depleted, see section 2.2.3), corrections to the DCR cut (see section 2.5.2), and parameters that vary the peak-shape of particular gamma decay

peaks. MPP also organizes detector data it sources explicitly by configuration rather than dataset (i.e. DS1-2 and DS3-6 instead of DS1, DS2, DS3, etc.). MPP calculates a total runtime for each configuration that incorporates channel selection and data cleaning, so that datasets with detectors that were only operational for part of the time are accounted for appropriately.

The processed simulation data from MPP is sourced into a python script that histograms the hit data and normalizes the hit data by the number of primaries for each component group. Two sets of histograms are produced up to 10 MeV - one set with a variable binning scheme to be used in the frequentist statistical model, and the other set with a 1keV fixed binning scheme to be used in the Bayesian statistical model. The latter set form the probability distribution functions (PDFs) that source the model discussed in section 5.3. These histograms are further subdivided by module, detector type, and multiplicity cut type.

### Section 4.3: Assay-Based Background Model

The PDFs produced from the processed simulation data exist in units of counts per decay (per keV with 1 keV binning). When combined with the measured activity of a component group from the radioassay program and the measured masses of the detectors, these histograms can be expressed in units of the background index:

$$\text{Model}_{\text{dtg},\text{dc}} = \sum_{\text{hw}} \text{PDF} \times \frac{\text{activity}_{\text{hw},\text{dc}}}{m_{\text{dtg}}} = \sum_{\text{hw}} \frac{\text{cts.}}{\text{decay}} \times \frac{\text{decays}}{\text{sec} \cdot \text{kg}} \times \frac{\text{sec}}{\text{yr}} = \sum_{\text{hw}} \frac{\text{cts.}}{\text{kg} \cdot \text{yr}} \quad (4.1)$$

where **dtg** is the detector grouping (“M1Enr”, “M2Nat”, etc.), **dc** is the decay chain, and **hw** is the hardware component grouping. These background indices can be summed together over the detector groupings and decay chains to produce the total assay-based background model (see Figure 3.2):

$$\text{Model} = \sum_{\text{dc}} \left( \sum_{\text{dtg}} \text{Model}_{\text{dtg},\text{dc}} \times \frac{m_{\text{dtg}}}{M_{\text{tot}}} \right) \quad (4.2)$$

For a more detailed description of the assay-based background model, see [46].

### Section 4.4: Analysis of <sup>232</sup>Th Excess through Simulations

In section 3.4 the data from natural detectors C1P2D1 and C1P7D1 suggested the presence of a non-uniform thorium radiation excess in a component near the opening of the Module 1 crossarm. In order to examine the likelihood of these components exhibiting this excess in full or in part, simulations of these components were conducted and compared to data from the DEMONSTRATOR as well as updated radioassay measurements.

#### 4.4.1: Module 1 Bulk Component Simulation Studies

Several sources in the region of the Module 1 (M1) crossarm and cryostat hoop junction (CCJ) were identified as potential sources for the thorium excess. These sources include:

- The vespel-made thermal buffer plate near the CCJ (`VespelSpider`)
- The copper adapter plates above strings 2 and 7 (directly above C1P2D1 and C1P7D1, `StringAdapterPlates`)
- The vespel pins that support the coldplate's position in the cryostat (`VespelPins`)
- The high voltage and signal cables that sit atop the coldplate (`CPHVCables` and `CPSigCables`)
- The high voltage and signal cables that line the length of the crossarm (`CrossarmCables`)
- The cold plate interface copper block at the end of the thermosyphon (`CPInterface`)
- The copper bolts that mount the cold plate interface block to the coldplate (`TSBolts`)
- The copper bolts and nuts that mount the cryostat body to a copper flange, along the lower rail (`RailBolts`)
- Four welds that join copper components in the region:
  - The circular weld at the CCJ, referred to here as `CW1_Annulus`
  - The circular weld joining the two crossarm tube pieces between the Front Spider and the CCJ, referred to here as `CW2_Annulus`
  - The rectangular weld joining the thermosyphon tube adapter to the cold plate interface block, referred to here as `TSW3_AnnularBox`
  - The circular weld joining the thermosyphon tube to the thermosyphon adapter, referred to as `TSW4_Annulus`

The approximate locations of these components are shown in Figure 4.1. All of these components were simulated as part of the standard background modeling simulation campaign, described in sections 4.1 and 4.2. Simulations processed by MPP in the DS3-6 configuration using all segments of the  $^{232}\text{Th}$  decay chain were then used for the  $^{232}\text{Th}$  excess studies. The welds were simulated separately on the `hcdData` computing cluster at Oak Ridge National Laboratory (ORNL). The volume of each weld was approximated as an annulus or hollow box at their appropriate locations within `MaGe`. 2.5 million  $^{232}\text{Th}$  decay chain primaries were populated in the bulk of these volumes with the `RDMiso` generator.

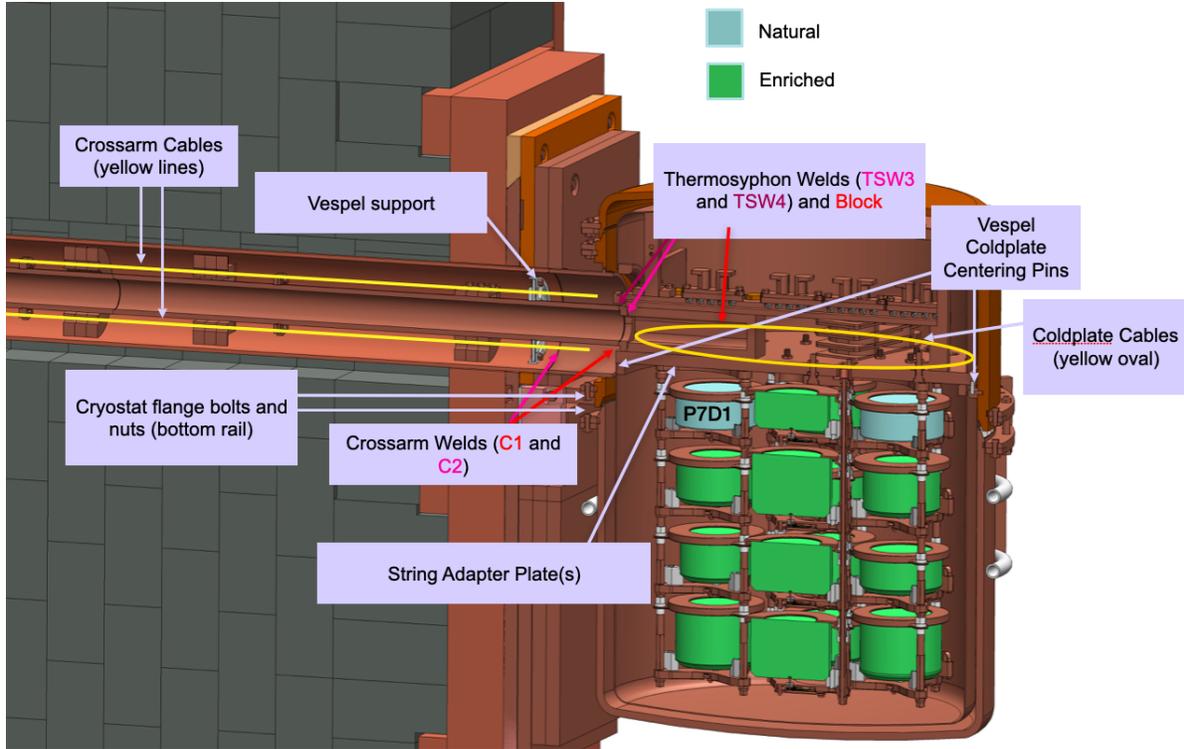


Figure 4.1: The location of components identified as potential source for the  $^{232}\text{Th}$  excess in Module 1. One of the two natural detectors that observed a high rate from the  $^{232}\text{Th}$  chain, C1P7D1, is labeled.

The weld simulations were compared to DS3-6 data (excluding DS5a due to excess noise) with a peak scaling method. The 238 keV gamma peak of one weld simulation was allowed to scale up to match the amplitude of the 238 keV gamma peak of detector C1P7D1 in data. The required activity in mBq to produce the 238 keV peak in data for each weld through this peak scaling method is shown in table 4.1, with the caveat that this activity would have to explain all of the thorium excess. The welds, if true sources of the  $^{232}\text{Th}$  excess, would need to have a bulk activity on the order of mBq.

Collaborators determined that the configuration of the DEMONSTRATOR in DS9, which was ongoing at the time of this study, could be used as an assay instrument with enough sensitivity to check the activities of the welds. An in-situ assay took place throughout the month of February 2022 in which the welded parts from Module 1 were cut out from their original components and placed beside the Module 2 cryostat, within the inner copper shield. Data were taken over 18 days, producing upper limits on the weld activities shown in table 4.1. These upper limits are consistently below the required activities from the peak scaling study, and consequently the welds were ruled out as primary sources of the  $^{232}\text{Th}$  excess.

A required activity was extracted from the other candidate components using the same peak scaling method, with a few minor differences. This analysis was not performed on the string adapter plates, which had since been ruled out as a candidate as their presence did not make a noticeable contribution to the

Component	Required Activity [mBq]	Assayed Activity or U.L. [mBq]	Assay Location	Notes
CW1_Annulus	0.89	<0.52	SURF (USA)	in-situ assay in Module 2
CW2_Annulus	5.82	<1.3	SURF (USA)	in-situ assay in Module 2
TSW3_AnnularBox	2.47	<0.72	SURF (USA)	in-situ assay in Module 2, TS welds assayed on same piece
TSW4_Annulus	5.89	<0.72	SURF (USA)	in-situ assay in Module 2, TS welds assayed on same piece
Vespe1Spider	7.14	<0.09	SNOLAB (Canada)	Re-assay TBD at HADES
Vespe1Pins	1.18	0.11	HADES (Belgium)	
RailBolts	1.25	<0.32	HADES (Belgium)	12 bolts assayed, 4 simulated, activity scaled to 4 bolts
TSBolts	1.40	<1.80	HADES (Belgium)	4 bolts assayed, 6 simulated, activity scaled to 6 bolts
CPInterface	1.78	<4.0	LSC (Spain)	Re-assay at LNGS (Italy), see section 4.4.4
CPHVCables	3.10	<2.66/bundle	LSC (Spain)	2 bundles is <5.32 mBq
CPSigCables	4.88	<3.2/bundle	LSC (Spain)	2 bundles is <6.4 mBq
CrossarmCables	40.33	<2.66/b. or <3.2/b.	LSC (Spain)	4 bundles is <11.72 mBq

Table 4.1: Table of estimated required activities for candidate components to explain the  $^{232}\text{Th}$  excess using the 238 keV peak scaling method as well as their assayed activities.

total rate in DS9. Simulations of the remaining components were sourced from the DS3-6 MPP production simulations, using a primary cut to isolate events from the relevant volumes. In **MaGe**, primaries are populated in every specified volume with a probability that is scaled by the fractional mass of that volume against all specified volumes. Therefore, the number of primaries of each simulation was scaled to match the fractional mass of each component in their component group. For example, simulations of the **Vespe1Pins** come from a set of simulations of the **M1ThermosyphonAndShieldVespe1** group, which contain 37.5 million primaries in total. The **Vespe1Pins** only make up 0.5% of the mass in that group. Therefore the total number of primaries for that set of simulations is scaled down from 37.5 million to 187,500 when normalizing the simulated spectra to units of counts per decay. These normalized spectra are scaled to the C1P7D1 238 keV peak in DS3-6 data to match the MPP configuration. The required activities are shown in table 4.1. Meanwhile, these components were shipped from SURF to various gamma counting stations in Canada and

Europe with the intention of re-assaying them down to the sensitivity of mBq. The results of this assay campaign are also listed in table 4.1. The results tentatively rule out components like the “Vespel Spider” and vespel pins as primary contributors, while not quite reaching the sensitivity needed to rule out components like the cables, thermosyphon bolts, and thermosyphon block. However, given the caveat that the peak scaling study assumes that all of the  $^{232}\text{Th}$  excess comes from one component in particular, there remains a possibility that any of these components could be contributors to the excess.

Simulations of these components were also used to investigate the impact of a Module 1  $^{232}\text{Th}$  excess on the Module 2 rate. First, the 238 keV and 2615 keV gamma peak efficiencies of these components in Module 1 and Module 2 detectors<sup>1</sup> were compared with one another (see Table 4.2, note that some components were combined for this study).

Component	M1\M2 efficiency (238 keV)	M1\M2 efficiency (2615 keV)
VespelSpider (Enr)	85.4	9.4
VespelSpider (Nat)	10.8	7.3
Rail Bolts, TSBolts, and CPInterface (Enr)	21.2	39.5
Rail Bolts, TSBolts, and CPInterface (Nat)	64.1	155.2
CPHVCables and CPSigCables (Enr)	116.4	17.5
CPHVCables and CPSigCables (Nat)	107.3	19.8
CrossarmCables (Enr)	17.5	11.4
CrossarmCables (Nat)	18.6	9.6
DS6 DATA (Enr)	$1.1 \pm 1.3$	$1.9 \pm 1.8$
DS6 DATA (Nat)	$1.7 \pm 1.1$	$4.3 \pm 3.9$

Table 4.2: Table of 238 keV and 2615 keV peak comparisons between Module 1 and Module 2. The second and third columns list the magnitude of a peak in Module 1 over the magnitude of a peak in Module 2 for a given set of enriched or natural detectors. The peaks are background subtracted. The uncertainties on the data ratios were calculated through the standard Poisson error  $\sqrt{N}$  that was propagated through the sideband subtraction and ratio.

When compared to DS6 data, most of the simulated component groups produce much larger ratios between the two modules, suggesting they would have a minimal impact on Module 2 if they were the source of the  $^{232}\text{Th}$  excess. The one exception is the simulation of the **VespelSpider** as a  $^{232}\text{Th}$  source, which produces a ratio between the modules in the 2615 keV peak and natural detectors that sits within the uncertainty window of the ratio from DS6 data. This makes sense despite the disfavoring of this component group in Table 4.1, as the location of the **VespelSpider** has a line of sight to Module 2 that encounters less

<sup>1</sup>All detectors enabled in DS6.

attenuation from shielding than the other component groups listed in Table 4.2. An an additional study, a DS6 BI for Module 2 was calculated for every component group simulation, assuming that component group was the dominant contributor to the DS6 BI for Module 1 (see table 4.3). Under those calculations, most component groups only contributed between 13-23% to the enriched DS6 BI of Module 2, further disfavoring a large contribution to Module 2 from a Module 1  $^{232}\text{Th}$  excess. The `VespelSpider` simulation is another exception to this conclusion in the natural detectors, producing a theoretical 53.27% of the counts seen in the DS6 Module 2 BI. However, this exception exists as a result of the geometrical reasons stated above, and the `VespelSpider` is disfavored by the assay results shown in Table 4.1. It is important to stress that a “worst-case scenario” framework was adopted for this study, where all contributions to the Module 1 BI were from one hypothetical  $^{232}\text{Th}$  source. This scenario is very unlikely as it ignores other decay chain sources and excludes the possibility that multiple sources are responsible for the  $^{232}\text{Th}$  excess. This framework was chosen merely to illustrate the low likelihood of a Module 1  $^{232}\text{Th}$  excess significantly contributing to the Module 2 rate.

Component	BI in M2 (c/FWHM-t-yr)	% of DS6 M2 rate in BEW
<code>VespelSpider</code> (Enr)	1.99	23.3
<code>VespelSpider</code> (Nat)	6.43	53.27
Rail Bolts, TSBolts, and CPInterface (Enr)	1.19	14.01
Rail Bolts, TSBolts, and CPInterface (Nat)	1.80	14.93
CPHVCables and CPSigCables (Enr)	1.16	13.59
CPHVCables and CPSigCables (Nat)	1.97	16.33
CrossarmCables (Enr)	2.02	23.71
CrossarmCables (Nat)	4.53	37.51

Table 4.3: Table of theoretical background indices from components in Module 2, assuming all of the counts in the DS6 Module 1 BEW can be attributed to the component(s) in the first column.

#### 4.4.2: Point Source Studies: Peak Ratios

Another strategy to constrain the region of the excess is to simulate a number of point sources around the CCJ region and observe how well the 238 keV to 2615 keV peak ratio compares to data. One can also observe the relative rates in different detectors from these sources to see how well they match with data (see Figure 3.3). There is also a possibility, given the localized nature of the excess, that the excess itself is a point source. This could exist in the form of a particulate settling, or perhaps accidentally incorporated into, a component near the Module 1 crossarm opening. Such a contaminant could exhibit a high activity

even when made of a material with a low activity density due to its point-source like nature.

An initial series of simulations were carried out in a number of point locations based on their approximate detector efficiency and amount of attenuation from shielding materials:

- A point on the detector surfaces of C1P2D1 and C1P7D1, specifically such that these two points face each other (`P2D1Surface` and `P7D1Surface`)
- Center spot on the bottom surface of the cold-plate interface block (`CPInterfaceBottom`)
- The bottom surface of the end of the thermosyphon tube near the junction with the thermosyphon adapter (`ThermoTubeBottom`)
- The center bottom surfaces of the two vespel coldplate centering pins nearest the Module 1 crossarm opening (`VespelPin01` and `VespelPin03`)
- The center cryostat-facing surfaces of the four cryostat bottom rail bolts closest to the Module 1 crossarm opening (`RailBot101`, `RailBot102`, `RailBot111`, and `RailBot112`)
- The center cryostat-facing surfaces of the two cryostat top rail bolts closest to the Module 1 crossarm opening (`TopRailBolt01` and `TopRailBolt12`)
- The center of the top surface of the clamp plate, directly above `CPInterfaceBottom` (`ClampPlateTop`)
- The two inner copper shield bolts closest to the crossarm opening at the elevation of the cryostat top lid (`InnerCuShieldBolt01` and `InnerCuShieldBolt12`)
- A point each on the surface of the upper and lower part of the nearest crossarm baffle plate to the cryostat (`BafflePlateTop` and `BafflePlateBottom`)

These point source locations are shown more clearly in Figure 4.2

Simulations of these points were carried out using `MaGe`. A virtual spherical volume measuring a micrometer in diameter was placed at each specified location in the geometry. Up to a million primaries were populated at each virtual volume. These simulations were processed by MPP in the DS3-6 configuration. A peak ratio study for two prominent  $^{232}\text{Th}$  gamma peaks (at 238 keV and 2615 keV) was conducted using these simulations and DS3-6 open and blind data. The counts of each peak in C1P7D1 (one of the two “hot” detectors) were determined through sideband subtraction, covering bands of twice the peak width. The results of the study are shown in table 4.4.

These results formed two distinct groups - sources with too little shielding relative to C1P7D1 (ratio > 2) and sources with too much shielding (ratio < 2). This led to a search for a point source between the

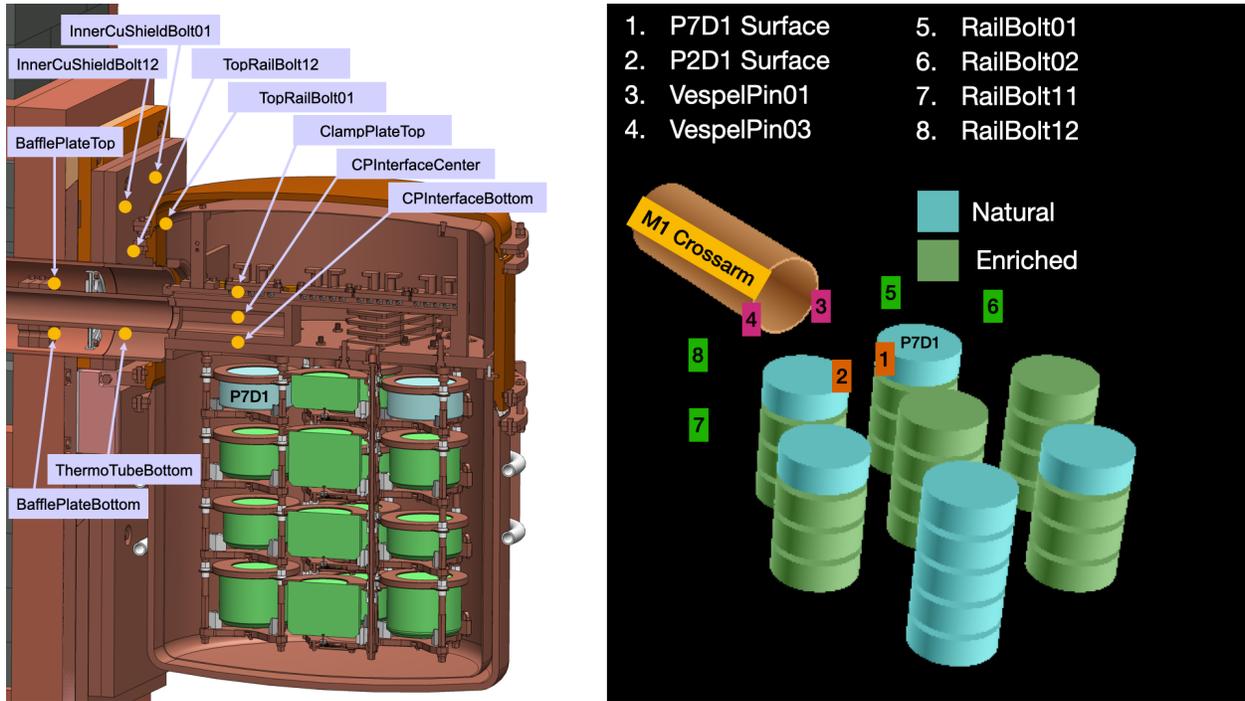


Figure 4.2: Locations of the point sources simulated to evaluate the possibility of a  $^{232}\text{Th}$  excess from a point source. Left is a cutaway graphic of the module 1 cryostat (point sources are yellow circles). Right is a full 7-string graphic of Module 1 (point sources are the numbered rectangles).

Component	Ratio in C1P7D1	Component	Ratio in C1P7D1
P2D1Surface	$36.0 \pm 1.7$	BafflePlateTop	$1.1 \pm 0.2$
P7D1Surface	$13.7 \pm 1.4$	BafflePlateBottom	$1.5 \pm 0.3$
ThermoTubeBottom	$10.5 \pm 1.9$	ClampPlateTop	$1.4 \pm 0.2$
VespelPin01	$11.9 \pm 1.1$	TopRailBolt01	$0.2 \pm 0.2$
VespelPin03	$6.6 \pm 1.1$	TopRailBolt12	$0.2 \pm 0.2$
RailBolt01	$9.1 \pm 1.0$	InnerCuShieldBolt01	$1.0 \pm 0.7$
RailBolt02	$9.2 \pm 1.0$	InnerCuShieldBolt12	$0.3 \pm 0.7$
RailBolt11	$1.2 \pm 0.5$	CPInterfaceBottom	$8.2 \pm 0.7$
RailBolt12	$7.6 \pm 1.3$	CPInterfaceCenter	$1.9 \pm 0.2$
DS3-6 DATA		$2.0 \pm 0.45$	

Table 4.4: Table of ratios of counts in the 238 keV gamma peak (from  $^{212}\text{Pb}$  in the  $^{232}\text{Th}$  chain) over the 2615 keV peak (from  $^{208}\text{Tl}$ ). The sideband subtracted peak counts are taken from the hot detector that was primarily online throughout the DS3-6 configuration, C1P7D1. The error on the data value was calculated through the standard Poisson error  $\sqrt{N}$  that was propagated through the sideband subtraction and ratio.

two groups with the right amount of copper thickness in the line of sight between the point and C1P7D1. Though there were three point sources within the uncertainty window of the ratio given by DS3-6 data, the center of the cavity within the thermosyphon block (referred to as `CPInterfaceCenter`) looked to be an ideal candidate, with its vertical position directly between `CPInterfaceBottom` and `ClampPlateTop` (see Figure 4.2), which produced peak ratios of 8.2 and 1.4 respectively. The resulting ratio from this point source, shown in table 4.4, matched the ratio in data with a difference of only 0.1. This result highlighted the interior of the thermosyphon block as a leading candidate for the source excess. However, relative rates between detectors in these simulations required review in order to eliminate sources with similar amounts of gamma attenuation to C1P7D1.

#### 4.4.3: Point Source Studies: Relative Rates

Simulations of the point source in the middle of the thermosyphon block cavity produce the closest 238 keV to 2615 keV peak ratio to data in detector C1P7D1. However, the location itself does not have a strong physical motivation as a point source, as it is in a region of nitrogen flow where it evaporates from liquid to gas. A more likely point source would exist on the bottom surface on the cavity. A simulation was run with a point source placed at the center of this bottom surface, directly underneath the `CPInterfaceCenter` source. Both simulations returned a ratio in C1P7D1 of 1.9, and their rates un-normalized by exposure in relevant detectors in the DS1-2 configuration <sup>2</sup> are shown in Figure 4.3. Though there is a slight variation

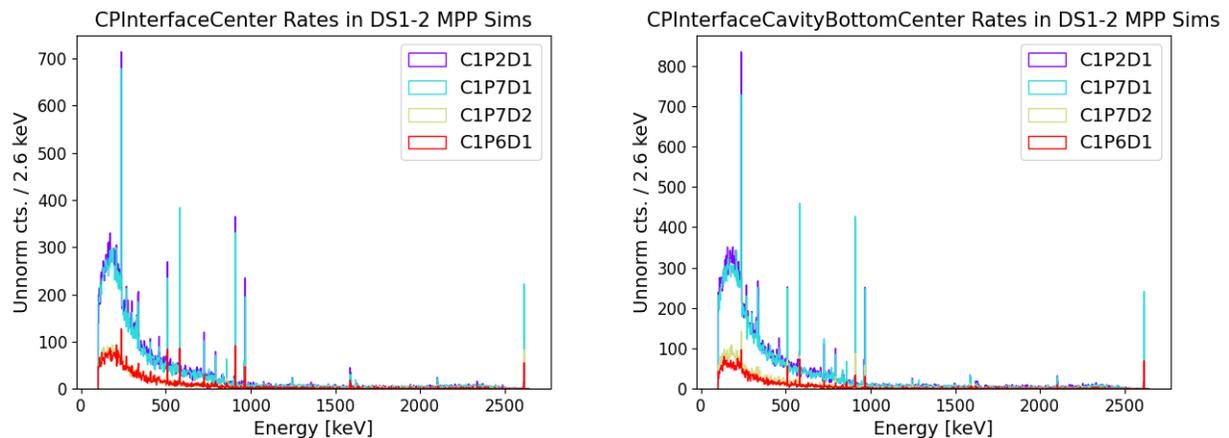


Figure 4.3: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the point source `CPInterfaceCenter`. **Right:** Rates from the point source `CPInterfaceCavityBottomCenter`

<sup>2</sup>The DS1-2 configuration is used when studying relative rates from the point sources because detector C1P2D1 is most active during this period.

Component	Ratio in det. C1P7D1	Component	Ratio in det. C1P7D1
CPInterfaceCenter	$1.9 \pm 0.2$	CPInterfaceBottom	$8.2 \pm 0.7$
CPICBCenter	$1.9 \pm 0.2$	ClampPlateTop	$1.4 \pm 0.2$
CPICBFront	$1.2 \pm 0.1$	TSBoltsBulk	$3.3 \pm 0.3$
CPICBBack	$1.3 \pm 0.2$	TSBoltsSurf	$4.7 \pm 0.4$
BafflePlateBottom	$1.7 \pm 0.4$	CPICBSurface	$1.9 \pm 0.2$
CPICBTM2	$1.9 \pm 0.2$	DS3-6 DATA	$2.0 \pm 0.45$

Table 4.5: Table of ratios of counts in the 238 keV gamma peak (from  $^{212}\text{Pb}$  in the  $^{232}\text{Th}$  chain) over the 2615 keV peak (from  $^{208}\text{Tl}$ ). The sideband subtracted peak counts are taken from the hot detector that was primarily online throughout the DS3-6 configuration, C1P7D1. The error on the data value was calculated through the standard Poisson error  $\sqrt{N}$  that was propagated through the sideband subtraction and ratio. CPICB is an abbreviation for CPInterfaceCavityBottom.

in counts reaching the two detectors with the next highest rates after the hot detectors (C1P7D2 and C1P6D1), the overall shape and ratios of the detector spectra are similar. For the rest of the relative rates studies, CPInterfaceBottomSurfaceCenter is used as a baseline, and the plot of relative rates from this simulated source shown in figure 4.3 will be used as a reference plot on the right of all remaining relative rate comparison figures. Simulations were carried out on point sources along the three axes that extend from the baseline source - moving forward and backward relative to the length of the crossarm, moving left or right perpendicular to the crossarm, and moving up and down which only requires revisiting the relative rates of CPInterfaceBottom and ClampPlateTop. The 238 keV to 2615 keV peak ratios for these new simulations (and some revisited ones) are all summarized in table 4.5. The locations of the cavity sources are shown in Figure 4.4.

Along the crossarm axis, simulations were generated at point sources at the very front (end) of the thermosyphon block cavity as well as the very back where it interfaces with the TSW3.AnnularBox weld. At the front, the ratio reveals too much shielding because of the angle at which radiation must travel through the bottom thickness of the thermosyphon block and the coldplate. The rate in C1P6D1 also increases relative to the rates of the two hot detectors, as this point lies closer to that detector (see Figure 4.5). In fact, any point further in this direction relative to CPInterfaceBottomSurfaceCenter will raise the relative rates of P6 and P4 string detectors relative to the hot detectors, and are therefore disfavored as excess locations. The point source at the back of the cavity similarly produces a ratio that suggests too much shielding. However, at this location the rate in C1P7D2 increases relative to the rate of C1P7D1 (see Figure 4.6). This is because the presence of the germanium material of C1P7D1 attenuates the rate of C1P7D2 to a relative rate that matches data. Therefore, in the absence of additional germanium shielding, point source positions further

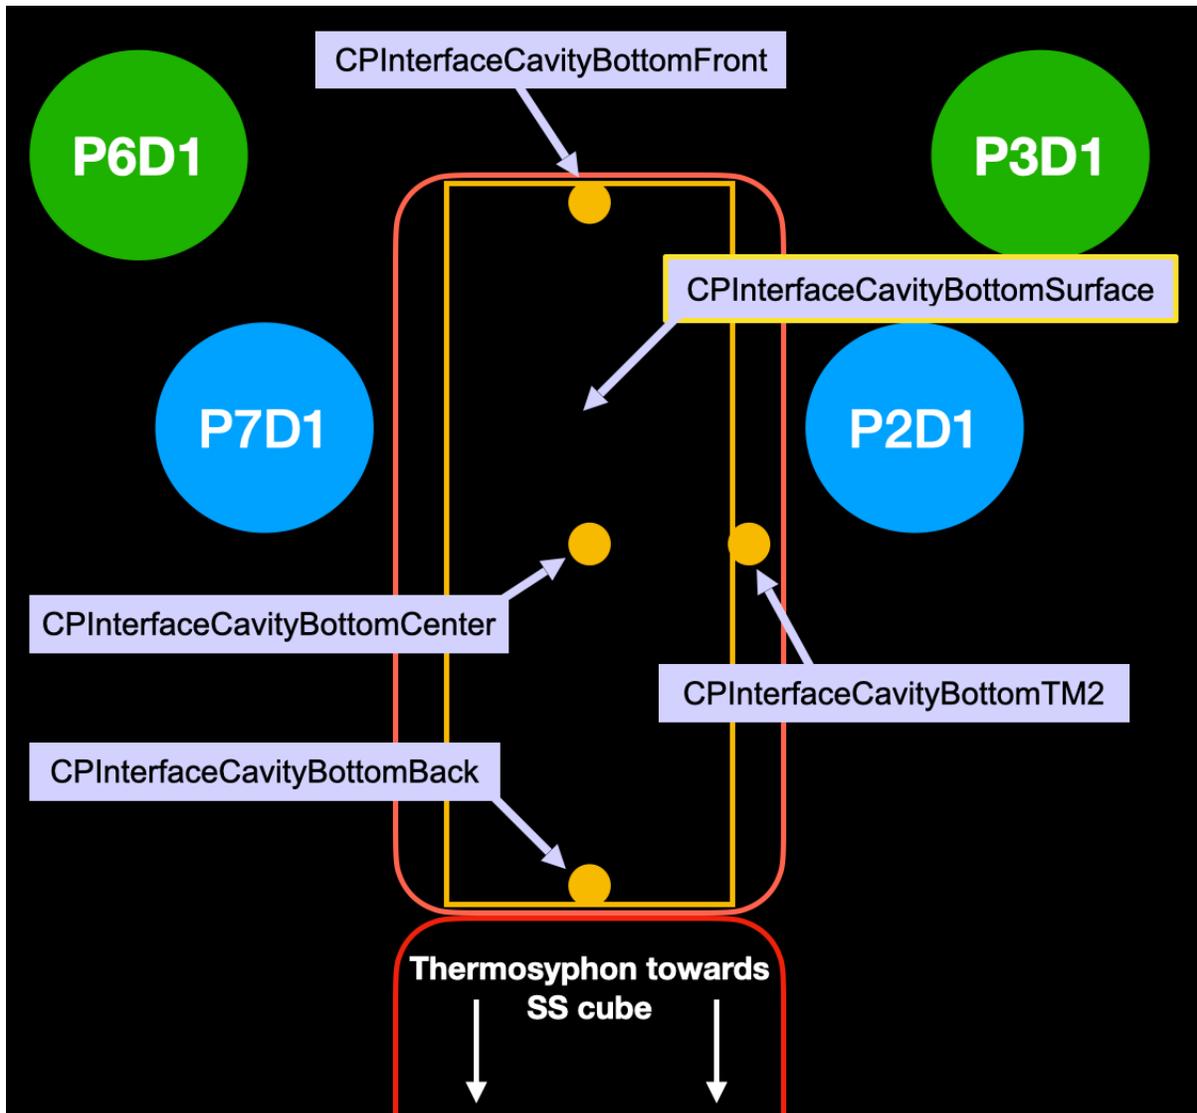


Figure 4.4: A top-down view of the thermosyphon block showing locations of the sources simulated on the bottom surface of the cavity to study the relative rates from the  $^{232}\text{Th}$  chain in different detectors. The yellow box outlines the surface area where primaries were generated for the `CPIInterfaceCavityBottomSurface` source. The detector positions are not to scale, and in reality C1P7D1 and C1P2D1 are partially beneath the thermosyphon block.

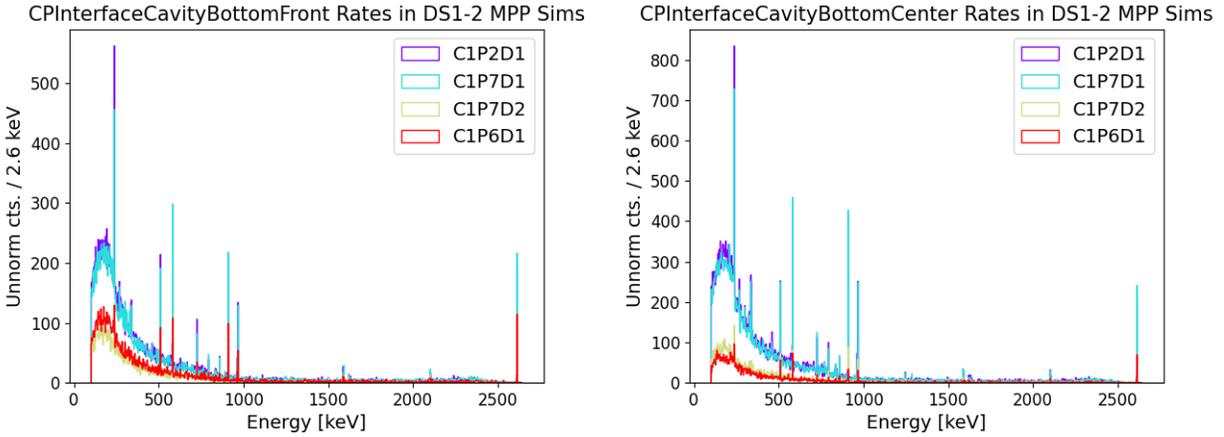


Figure 4.5: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the point source `CPInterfaceCavityBottomFront`. **Right:** Rates from the point source `CPInterfaceCavityBottomCenter`

back into the crossarm raise the relative rate in the lower P7 and P2 detectors, which disfavors these points as the excess source. Returning to Figure 4.1, there are points further down the crossarm and away from

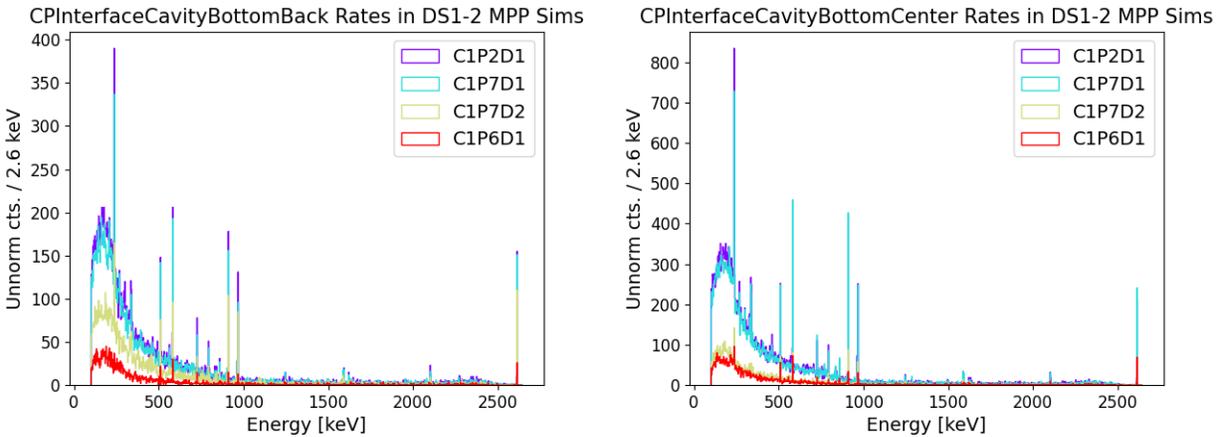


Figure 4.6: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the point source `CPInterfaceCavityBottomBack`. **Right:** Rates from the point source `CPInterfaceCavityBottomCenter`

the cryostat where attenuation to C1P7D2 increases again due to the corner material that sits below the CCJ. To study this, the relative rates from point source `BafflePlateBottom` were plotted and compared to baseline (see Figure 4.7). Though there is some attenuation to detector C1P7D2, there is not enough relative to C1P7D1 to match the ratios of those spectra seen from the thermosyphon block cavity. Moving any further back from the baffle plate would attenuate the rates to both detectors at a similar rate, and thus

disfavors all points along the crossarm that sit behind the baseline point source.

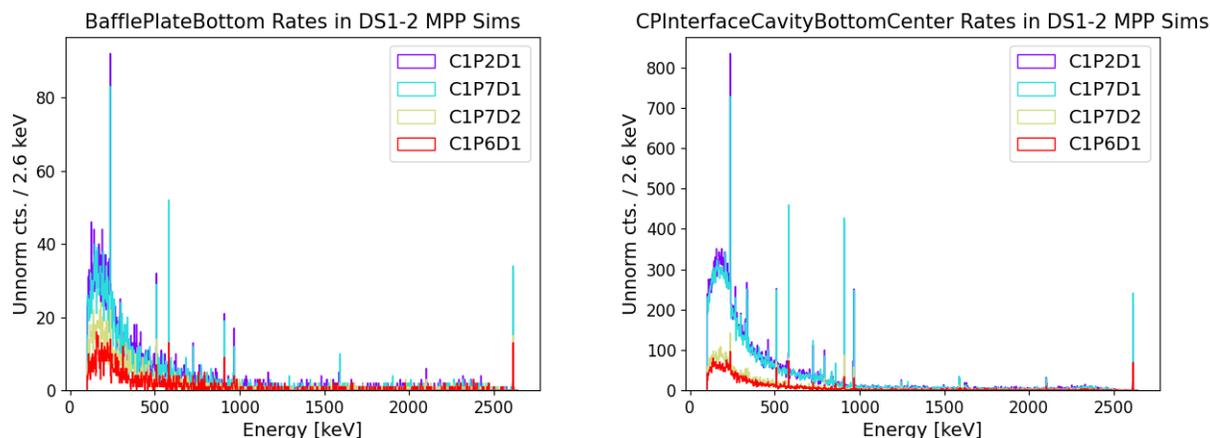


Figure 4.7: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the point source `BafflePlateBottom`. **Right:** Rates from the point source `CPIInterfaceCavityBottomCenter`

Considering the axis perpendicular to the crossarm, a simulation was conducted with a point source to the side of the baseline source, towards the position of Module 2 (`CPIInterfaceCavityBottomTM2`). Though the ratio is within the uncertainty window of the ratio in data, the relative rate of C1P2D1 increases significantly compared to C1P7D1 (see Figure 4.8). The relative rates between the two detectors in data are only different by about 25%, so this disfavors any points that lie outside a relatively tight range on either side of the crossarm axis. When comparing the relative rates from `CPIInterfaceBottom` and `C1ampPlateTop`

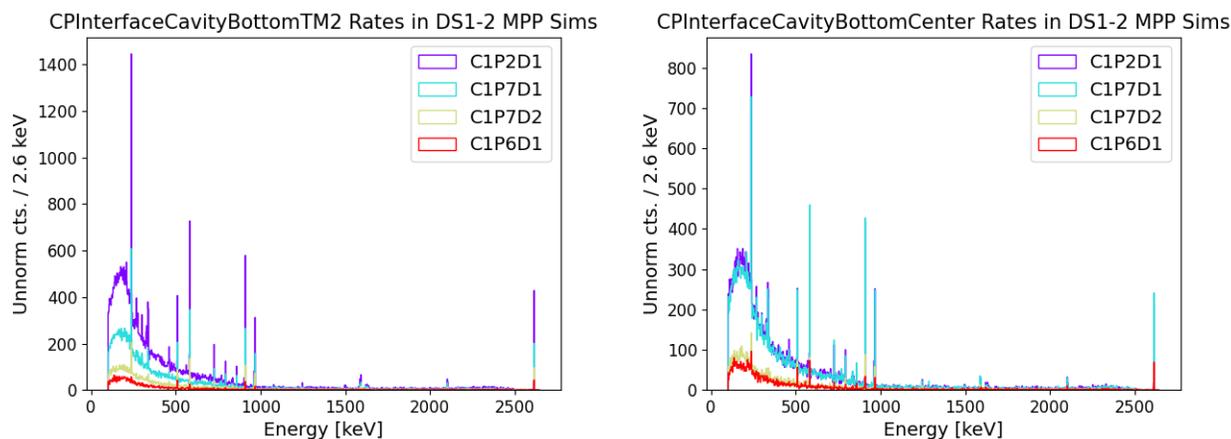


Figure 4.8: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the point source `CPIInterfaceCavityBottomTM2`. **Right:** Rates from the point source `CPIInterfaceCavityBottomCenter`

to those from the baseline point source, the ratio of the 238 keV peak to the continuum below 238 keV from

those sources does not match what is seen from `CPInterfaceBottomSourceCenter`, disfavoring points above and below the cavity (see Figures 4.9 and 4.10).

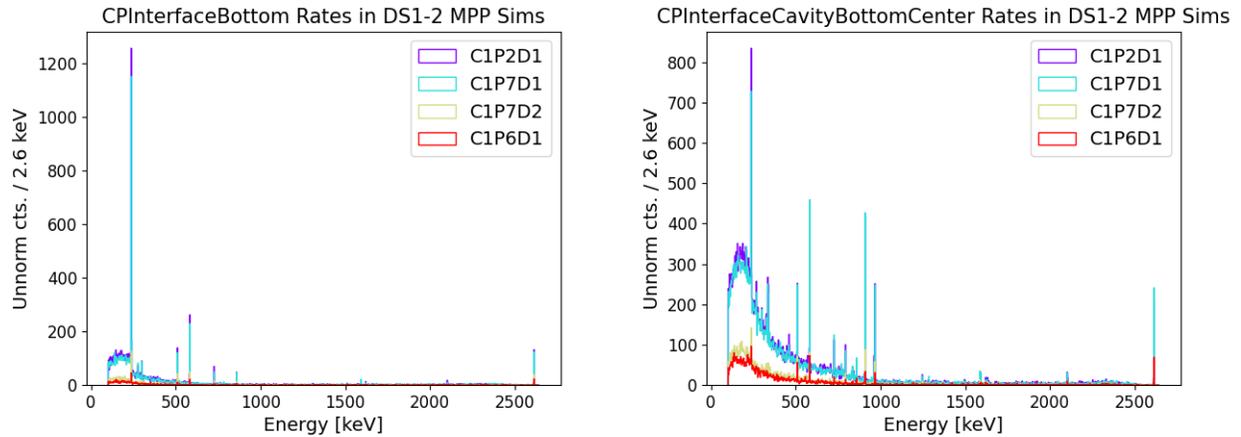


Figure 4.9: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the point source `CPInterfaceBottom`. **Right:** Rates from the point source `CPInterfaceCavityBottomCenter`

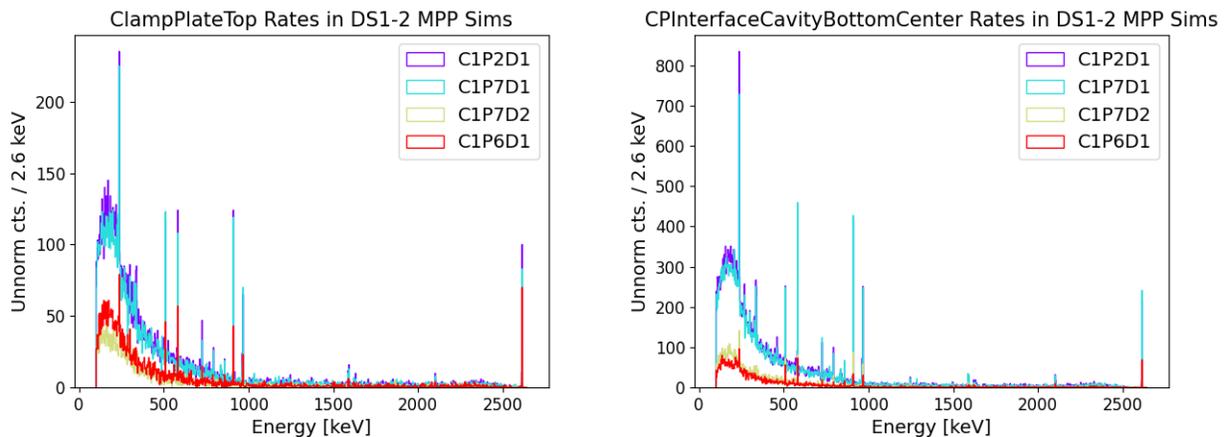


Figure 4.10: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the point source `ClampPlateTop`. **Right:** Rates from the point source `CPInterfaceCavityBottomCenter`

Theoretically it is possible to have multiple sources that symmetrically surround the thermosyphon block cavity and reproduce the ratios and rates from the cavity. The candidates for this are the `TSBolts` (see Figure 4.11). Bulk simulations of these bolts were re-run with the same number of primaries as the previous set of point source simulations. An additional surface-only simulation was also conducted with the `GSS` generator. The resulting ratios suggest, however, that these bolts have too little shielding between themselves and the hot detectors. When comparing their relative rates to the `CPInterfaceCavityBottomCenter` point source, the bolts are further disfavored as the source of the excess based on the ratio of the 238keV peak to the

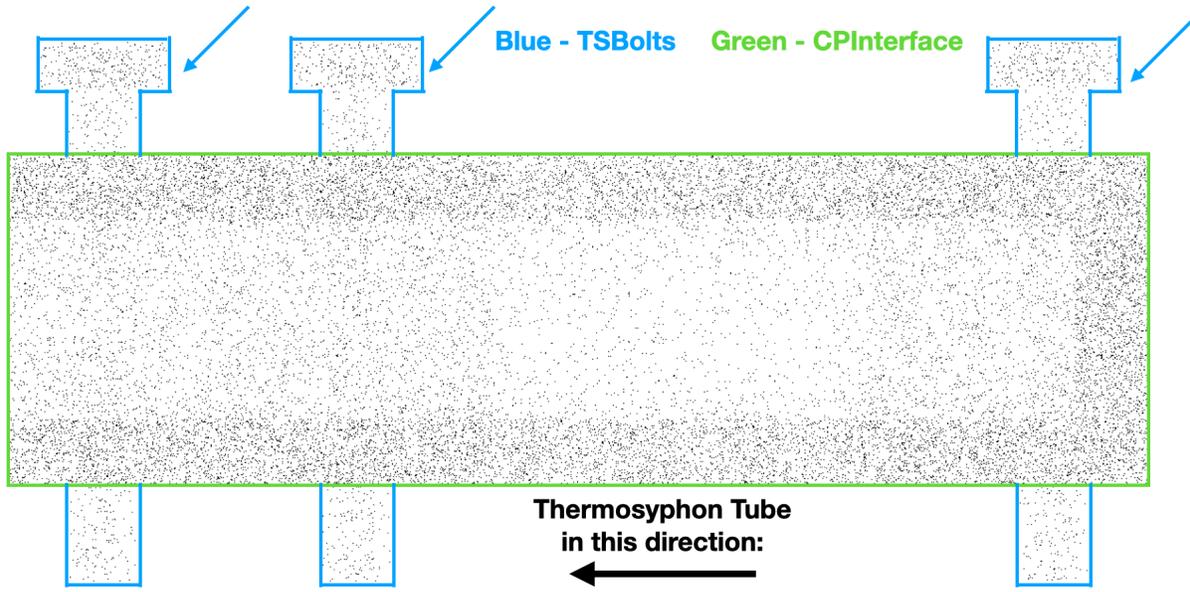


Figure 4.11: A side view of simulated primaries generated inside the copper bulk of the thermosyphon block (CPIInterface) as well as the thermosyphon bolts (TSBoltsBulk) that hold the block in place. The components are outlined in green and blue respectively.

continuum below 238 keV (see Figures 4.12 and 4.13).

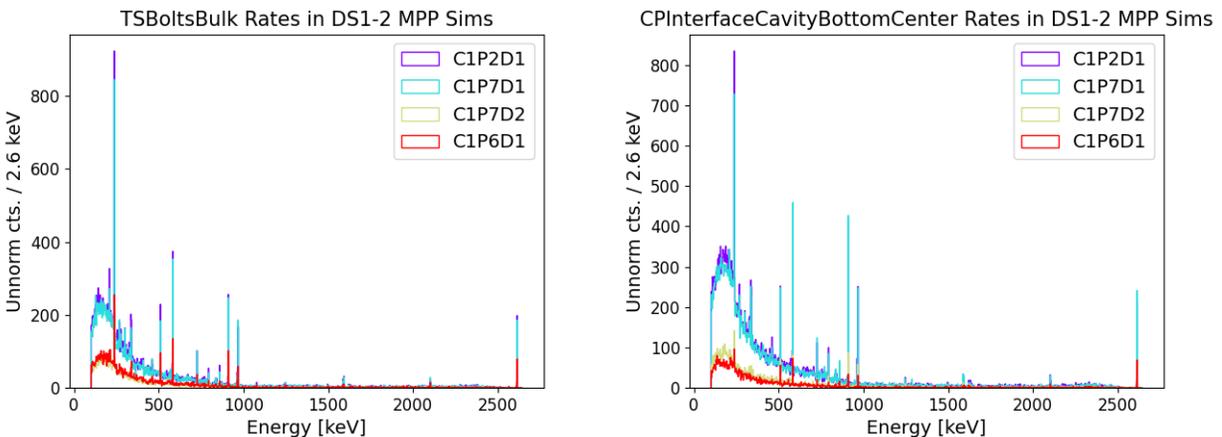


Figure 4.12: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the grouping TSBoltsBulk. **Right:** Rates from the point source CPIInterfaceCavityBottomCenter

At this stage, the cavity of the thermosyphon block is the most plausible location for the  $^{232}\text{Th}$  excess. However, even though the baseline point source is more physically motivated than CPIInterfaceCenter, a point source in general is highly specific and far less likely than a contaminant spread over an area, like the entire bottom surface of the cavity. A simulation was carried out that populated primaries across this surface. The resulting ratio matched that of the baseline point source and the relative rates are similar

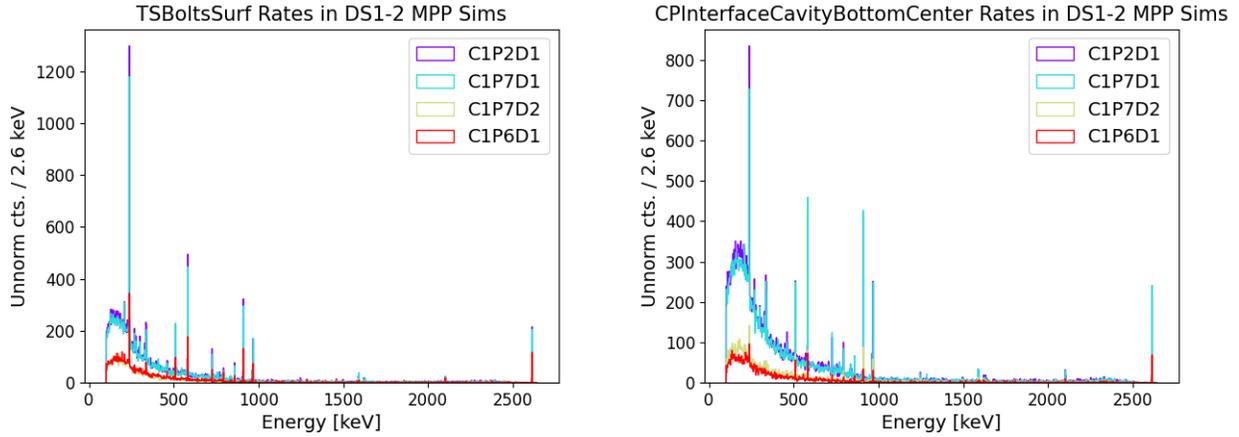


Figure 4.13: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the grouping `TSBoItsSurf`. **Right:** Rates from the point source `CPInterfaceCavityBottomCenter`

between the surface and point source (see Figure 4.14). This makes sense, as the surface simulation takes the effects from the front cavity source and back cavity source and cancels them out to produce results that match the center source. As the most plausible source, the bottom surface of the thermosyphon block cavity

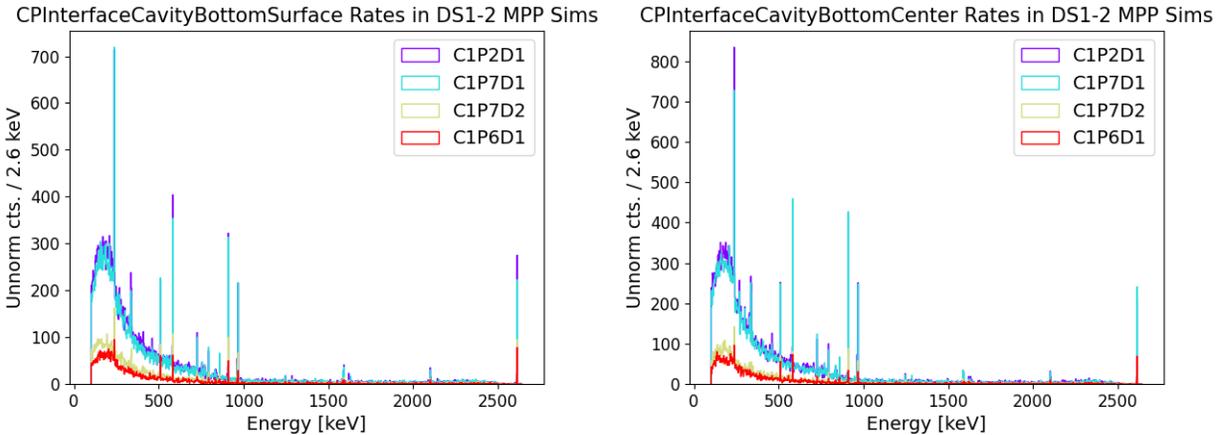


Figure 4.14: Comparison of relative and un-normalized detector rates in the DS1-2 configuration using MPP simulations. Note that the range on the y-axis is different between the two figures. **Left:** Rates from the surface source `CPInterfaceCavityBottomSurface`. **Right:** Rates from the point source `CPInterfaceCavityBottomCenter`

has been selected as our candidate source of the  $^{232}\text{Th}$  excess and is now included in the background model as its own component group. Simulations of the bottom surface were fed into the production pipeline and PDFs were generated from these simulations through the same method as discussed in section 4.2. This was done for the cavity in Module 1 and Module 2 in order to prove the hypothesis that the Module 2 cavity does not contribute any excess  $^{232}\text{Th}$  counts.

#### 4.4.4: Re-assay of the Thermosyphon Block

Table 4.1 lists the upper limit on the assay of the thermosyphon block (`CPInterface`) as 4.0 mBq. However, due to its plausibility as the source of the  $^{232}\text{Th}$  excess, the part was sent to LNGS in Italy to be re-assayed. After 39 days of counting, the upper limit on the activity of the bottom surface of the part was constrained to 0.57 mBq [49]. Given the estimated required activity of 1.78 mBq (in the bulk volume), the results of this assay indicate that the postulated contaminant is no longer on the bottom surface of the thermosyphon block cavity. However, this does not exclude the possibility of its presence during the runtime of the experiment, especially considering the evidence presented in the preceding subsections. A possible explanation is that a contaminant from the  $^{232}\text{Th}$  decay chain was introduced to the closed loop of nitrogen within the thermosyphon during installation, when the loop was not yet sealed. This contaminant was then deposited onto the bottom surface of the thermosyphon block cavity as the liquid nitrogen evaporated into its gaseous form. The contaminant then escaped the closed loop once the thermosyphon was opened after operations in Module 1 were concluded.

## CHAPTER 5: Building a Bayesian Statistical Model

To make inferences from a set of data ( $D$ ), the data is modeled with a particular framework ( $M$ ). This model has a set of parameters ( $\theta$ ) that vary with their own distributions. The viability of these model parameters can be obtained via a multi-dimensional probability distribution of the model parameters given the data and model type. This distribution,  $P(\theta|D, M)$  is called the “**posterior**” and is calculated with Bayes’ theorem [50]:

$$P(\theta|D, M) = \frac{P(D|\theta, M)P(\theta|M)}{P(D|M)} \quad (5.1)$$

where  $P(D|\theta, M)$  is the “**likelihood**” of a set of data given a set of model parameters and the model framework,  $P(\theta|M)$  is the likelihood of the model parameters having a particular range of values given “**prior**” knowledge of the model type, and  $P(D|M)$  is the probability of generating the data given a particular model type by integrating over the model’s parameters (this variable is called the “**evidence**”). The ideal model is one with a set of parameters that maximizes the probability in the posterior. At this maximum point, the expectation value of each parameter ( $E[\theta]$ ) is taken to return the desired information from the inference.

Bayesian inference methods on datasets are powerful statistical tools, but they rely on an ability to calculate the posterior. This is usually too difficult to do analytically except for the most basic of likelihood and prior distributions. To get around this, the posterior can be sampled and the sample distribution can be used as a proxy. If the posterior distribution is not too complicated and the model only has one or two parameters (dimensions), a grid evaluation could be used for sampling. However, the grid needs to cover the region of highest density (probability) in the posterior and this location is less obvious the more complex the distribution is. This is especially true if the posterior is multimodal or highly peaked. As the number of parameters and dimensions increase, a grid evaluation becomes unusable. The number of grid points required to sample the posterior for one parameter and maintain the same uncertainty exponentially increases with the number of dimensions. Making the grid calculations in this case becomes tedious and pointless, since most of the grid points will be in regions of low density.

At high complexity and dimensionality of the posterior, Monte Carlo methods are more effective for sampling. To find the region(s) of highest density, where the expectation value  $E[\theta]$  may be taken, Monte

Carlo can be used to approximate an integration over the posterior:

$$E[\theta] = \frac{\int \theta P(\theta|D, M) d\theta}{\int P(\theta|D, M) d\theta} \approx \frac{1}{N} \sum_{i=1}^N \theta_i \quad (5.2)$$

where  $N$  is the number of random Monte Carlo samples taken from the posterior. The uncertainty in the Monte Carlo integration approximation is proportional to  $N^{-1/2}$  and has no dependence on the number of dimensions in the posterior. Therefore the sample size does not have to be unreasonably large to handle a model with several parameters.

### Section 5.1: Markov Chain Monte Carlo

Random sampling of a distribution is the core tenet of a Monte Carlo method, but these methods aren't all created equal when solving a Bayesian inference problem. For example, a Monte Carlo method using completely random sampling of the posterior tends to get stuck in low density regions in large parameter spaces, much like a grid evaluation. Markov Chain Monte Carlo (MCMC) can get around this by reducing the randomness and setting up a walk over parameter space that "prefers" regions of high density [50]. This random walk uses the Markov chain method, which evaluates the probability of moving from one parameter state  $\theta_t$  to another state  $\theta_{t+1}$  using a transition kernel ( $Q$ ) specific to the MCMC method being used and then either accepts or declines the move based on that probability. This kernel, which only depends on the current state and does not remember previous states, drives the random walk to move towards the high density regions. A requirement of the Markov chain method is that over enough samples, the distribution of samples will reach a stationary state that is representative of the posterior. This requirement can be fulfilled if the chain is reversible, such that the transition kernel specifies the probability of transitioning from state  $\theta_a$  to  $\theta_b$  is just as probable as transitioning from  $\theta_b$  to  $\theta_a$ . One of the simplest MCMC methods, the Metropolis algorithm [50], has this reversible chain. To begin, a sample  $\theta_{t+1}$  is drawn from the transition kernel  $Q(\theta_{t+1}|\theta_t)$ . Then, the Metropolis ratio is calculated:

$$\rho = \frac{q(\theta_{t+1})}{q(\theta_t)} \quad (5.3)$$

where  $q(\theta)$  is the probability density at the specified parameter value. If  $\rho \geq 1$ , the sample value  $\theta_{t+1}$  is accepted and the algorithm repeats itself with  $\theta_{t+1}$  becoming the new value for  $\theta_t$ . If  $\rho < 1$ , then a random number,  $r$ , is drawn from the uniform distribution  $U(0, 1)$ . If  $\rho \geq r$ , the sample value is accepted, and if not, it is rejected. In the case of a rejection, a new sample is taken from  $Q(\theta_{t+1}|\theta_t)$  and the algorithm continues.

The Metropolis algorithm can sample a posterior much more efficiently than grid sampling, especially in posteriors with large parameter ranges. However, it also suffers at high dimensionality due to a concept called “concentration of measure” [51]. With every additional dimension, the hypervolume of the parameter space becomes very large, leading to a situation where most of the probability mass lies far away from the mode of the posterior distribution. Since the transition kernel of the Metropolis algorithm only takes into account the local space of the initial sampling rather than the full posterior, it will aimlessly become stuck in large regions of low density. This is illustrated in Figure 5.1 where distributions of varying dimensionality are sampled using the Metropolis algorithm.

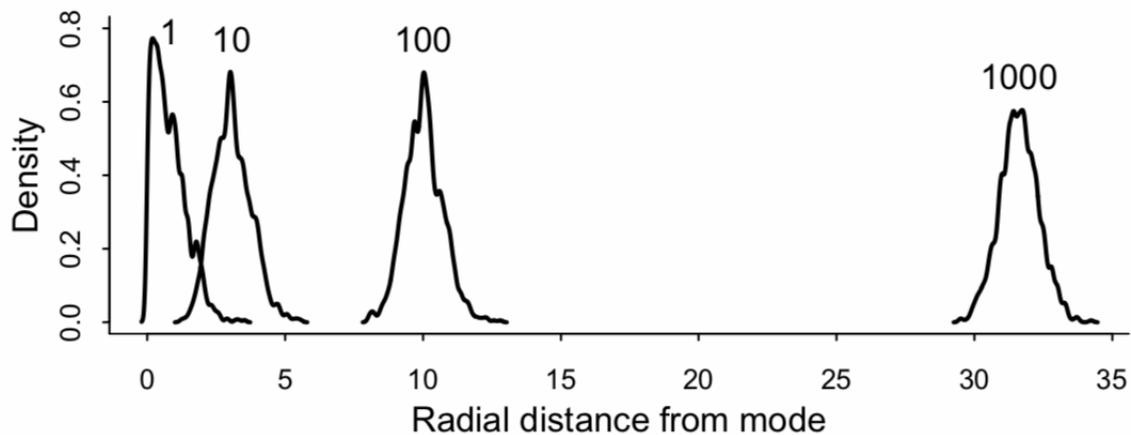


Figure 5.1: Sampling results from different distributions using the Metropolis algorithm. Density refers to the density of samples, and each peak contains 1000 samples of a different distribution. The dimensionality of the distribution is shown above the peak. This demonstration was done by Richard McElreath for his book “Statistical rethinking : a Bayesian course with examples in R and Stan”. [51]

## Section 5.2: Hamiltonian Monte Carlo, NUTS, and PyMC3

To bypass the issue of “concentration of measure”, we require an MCMC algorithm that makes smarter sampling choices and can account for the global shape of the posterior. Hamiltonian Monte Carlo (HMC) is a good candidate [51]. As the name suggests, HMC utilizes Hamiltonian mechanics to govern the transitions between samples. A simulation takes place in which the posterior is converted into a multi-dimensional “surface” with lower elevations correlated to high probability densities. A frictionless particle is placed on this surface at coordinates that correspond to some combination of model parameters ( $\theta_t$ ). The particle is given some random initial momentum vector and is allowed to travel along the surface for some distance, with the slope of that surface informed by the posterior’s gradient. At the end of the allowed distance, the simulation freezes, and the new coordinates of the particle are taken as the sample ( $\theta_{t+1}$ ). The particle is

then given another random momentum vector and travels along the surface, repeating the process. Since the particle will physically travel towards lower regions of the surface (areas of high probability density), the sampling is done effectively. The amount of distance between samples is controlled by two variables, “leapfrog steps” and “step size”. The leapfrog steps control how many lengths a particle can travel before a sample is taken, and the step size controls the length of the individual leapfrog steps, which in turn effects how quickly the particle can change direction. An example of HMC in practice, sampling from a 2D Gaussian posterior, is shown in figure 5.2.

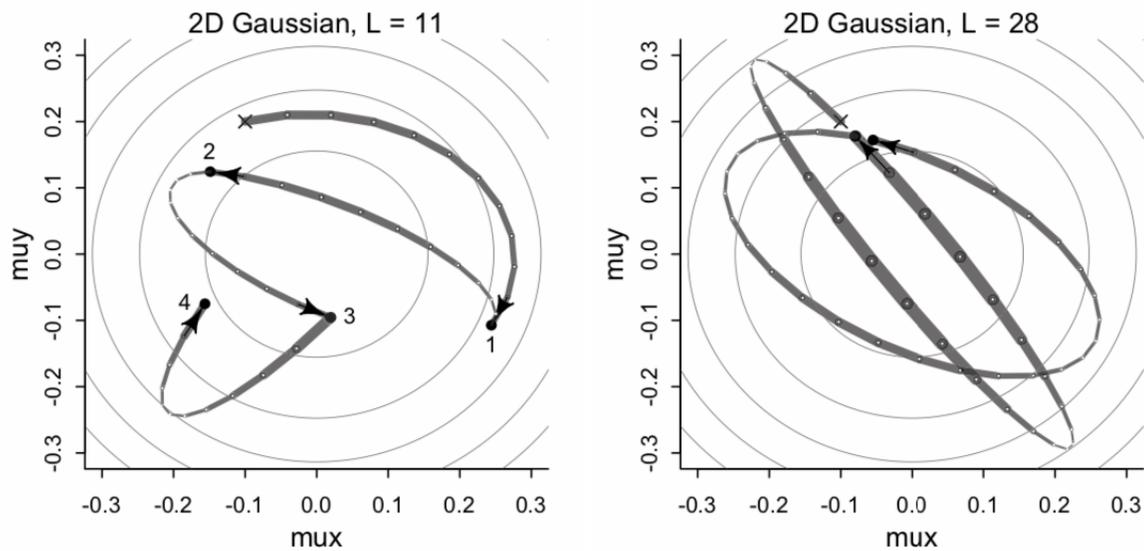


Figure 5.2: Sampling results after four samples (X’s) from a 2D Gaussian posterior using HMC and two different leapfrog step values ( $L=11$  and  $L=28$ ) with the same step sizes. **Left:** With a good  $L$  value, HMC explores the posterior efficiently. **Right:** With a poor  $L$  value, the particle can return to the area near the origin at each sampling point, leading the sampler to getting stuck. This can be avoided by using NUTS. This demonstration was done by Richard McElreath for his book “Statistical rethinking : a Bayesian course with examples in R and Stan”. [51]

There is a flaw in HMC where the combination of leapfrog steps and step size can cause a particle to leave its origin, rebound off a surface with low probability density, and return to the origin when a sample is taken. This potentially traps the particle in a region away from the mode. An HMC-based algorithm called the “No U-Turn Sampler” (NUTS) can get around this problem by not allowing the trajectory of the particle to turn around. If the trajectory double backs on itself, the simulation immediately freezes, and the sample is taken.

A python package, PyMC3 [52], utilizes NUTS to facilitate an effective sampling of the posterior for the purpose of making useful inferences from data. PyMC3 leverages another package, Theano, to handle the vector information from HMC in C and compiles it into machine code - boosting CPU performance. PyMC3

automatically tunes the number of leapfrog steps and the size of those steps based on input from NUTS, allowing the user to sample the posterior without resorting to a trial-and-error implementation of HMC.

### Section 5.3: A Statistical Model of the DEMONSTRATOR

In order to leverage PyMC3 to provide insights into the background data presented in section 3.3, we need to translate a model of MAJORANA DEMONSTRATOR backgrounds into the language of Bayesian probabilistic programming. Producing a posterior from which PyMC3 will sample requires definitions of a set of likelihoods and priors in accordance with equation 5.1<sup>1</sup>.

The DEMONSTRATOR detects discrete radiation events that happen at a particular rate. This situation is best modeled by a Poisson distribution:

$$P(r|\lambda) = \frac{e^{-\lambda}\lambda^r}{r!} \quad (5.4)$$

where  $r$  is the number of events observed in a given time interval and  $\lambda$  is the number of events expected in that same time interval. To formulate an energy spectrum these variables are measured for each energy bin ( $j$ ). The expected rate in an energy bin is equivalent to the sum of the activities of the radioactive sources ( $\theta_i$ ) multiplied by the efficiencies of those sources per bin ( $\epsilon_{ij}$ ):

$$\lambda_j = \sum_i \theta_i \epsilon_{ij} \quad (5.5)$$

Since the data we observe relies on this Poisson-like process, we choose the Poisson distribution as our likelihood. With  $N$  energy bins, our likelihood becomes:

$$P(D|\theta, M) = \prod_{j=1}^N \frac{e^{-\sum_i \theta_i \epsilon_{ij}} (\sum_i \theta_i \epsilon_{ij})^{r_j}}{r_j!} \quad (5.6)$$

The Monte-Carlo radiation transport simulations discussed in chapter 4 provide the efficiencies  $\epsilon_{ij}$ . The activities,  $\theta_i$ , are the parameters we are trying to infer.

Since the prior involves our previous knowledge about the model parameters, which in this case are the source activities, it makes sense to source the priors from the assay-based background model. Each prior distribution is treated as a truncated Gaussian with the mean equal to the assayed activity of the source and the standard deviation equal to the error bounds on the assay measurement. The prior distributions are truncated at zero to prevent the inclusion of un-physical assay values.

<sup>1</sup>Defining the evidence is unnecessary here, since only the shape of the posterior is needed to perform the sampling. The evidence will become important later when calculating the Bayes Factor, to be discussed in chapter 6.

With these likelihood and prior distributions, PyMC3 samples the posterior with three separate NUTS chains, each taking 10000 samples after an initial “burn-in” phase of 500 samples to tune the leapfrog steps and sizes. The model parameters ( $\theta_i$ ) are the activities for every tuple of hardware sources and decay chains (for example, [M1LMFEs,  $^{232}\text{Th}$ ]). The model itself is broken into a number of submodels, each with its own set of tuple activities that cover a particular dataset configuration (DS1-2, DS3-6, etc.), a particular module (Module 1 or Module 2), a particular detector type (enriched or natural), and a particular multiplicity ( $m=1$  or  $m>1$ ). The posterior is informed by all of these submodels simultaneously. In an “unpooled” probabilistic programming framework, the tuple activities in each submodel would be allowed to float independently of their other submodel equivalents. However, there exists a physical constraint on the activities such that they cannot vary from submodel to submodel (i.e. the source is the same in every submodel). Therefore, this work relies on a “pooled” probabilistic programming framework in which the tuple activities of each submodel are fixed to their equivalents in other submodels <sup>2</sup>.

---

<sup>2</sup>The cosmogenic activity tuples vary from submodel to submodel, but by a known decay constant that is multiplied to the DS0 activity value prior to fitting. See section 6.2 for more details.

## CHAPTER 6: Implementation of Inference Methods

### Section 6.1: Structure of Spectra and Component Groupings

In order to supply the statistical model with as much information as possible for the most successful sampling of the posterior, the model itself is divided into several submodels, each corresponding to a data spectrum of a particular type. The parameters that define the submodel are detector configuration, detector enrichment, multiplicity, and module. There are 17 submodels in total, each with its own likelihood constructed from PDFs of that particular set of parameters (see section 4.3). Each submodel also draws from a particular set of priors that match the associated parameters, though in practice the priors only change for cosmogenic decay chains where the central assay value depends on the dataset (see section 6.2).

The submodels are fit simultaneously to the corresponding spectra (data spectra or simulated data spectra) as part of a pooled statistical model, as one would expect the activity density of component groupings to not differ between submodels of different parameters. Additionally, using a pooled statistical model reduces the overall parameter space, which is already quite large. The component groupings to be used by the statistical model are the same as those introduced in the latest set of simulation and PDF generation campaigns (see sections 4.1 and 4.2).

### Section 6.2: Handling of Cosmogenic Decay Chains

The cosmogenic decay chains  $^{68}\text{Ge}$ ,  $^{60}\text{Co}$ ,  $^{57}\text{Co}$ , and  $^{54}\text{Mn}$  have short enough half-lives that their activity in components changes over the runtime of the experiment. This is accounted for in the fitting algorithm by multiplying a numerical factor to the DS0 activity value for a particular tuple involving one of the aforementioned decay chains. This factor changes value depending on the detector configuration of the submodel, which itself is an indicator of the amount of time that has passed since DS0.

The DS0 specific activity for the component, decay-chain tuple is calculated with the following equation if the component is saturated before being transported underground:

$$\text{S.A.} = P e^{-\lambda \cdot t_{\text{und}}} \quad (6.1)$$

where  $P$  is the cosmogenic production rate in atoms/kg/day,  $\lambda$  is the decay rate of the cosmogenic isotope,

and  $t_{\text{und}}$  is the amount of time in days the component was underground prior to the beginning of DS0. If the component is not saturated prior to being transported underground, then the following equation is used instead:

$$\text{S.A.} = \lambda P t_{\text{surf}} e^{-\lambda \cdot t_{\text{und}}} \quad (6.2)$$

where  $t_{\text{surf}}$  is the amount of time the component was exposed on the surface before being transported underground.

With the DS0 activity represented by  $A_0$ , we can calculate the activity of the component, decay-chain tuple with any detector configuration using the following exponential function, in which  $\lambda$  has been replaced with its reciprocal  $\tau$ , the time constant of the relevant decay chain in days:

$$A = A_0 \lambda e^{-t_{\text{eff}}/\tau} \quad | \quad t_{\text{eff}} = t_{\text{offset}} + t_{\text{mean}} \quad (6.3)$$

where  $t_{\text{eff}}$  is the amount of time in days that has elapsed between the beginning of DS0 and the beginning of the detector configuration ( $t_{\text{offset}}$ ) in addition to an average time in days that represents decays in the specified detector configuration ( $t_{\text{mean}}$ ). Finding  $t_{\text{mean}}$  requires finding the expectation value of the amount of time elapsed in the detector configuration, which is defined as:

$$t_{\text{mean}} = E[t] = \frac{\int_0^b t e^{-t/\tau} dt}{\int_0^b e^{-t/\tau} dt} \quad (6.4)$$

where  $b$  is the run time of the detector configuration in days. Through integration by parts and algebra, equation 6.4 reduces to:

$$t_{\text{mean}} = \frac{\tau(1 - (b/\tau + 1)e^{-b/\tau})}{1 - e^{-b/\tau}} \quad (6.5)$$

### Section 6.3: Handling of $^{40}\text{K}$ and $^{60}\text{Co}$ Priors

As discussed in section 5.3, the priors that support this statistical model of the DEMONSTRATOR are truncated Gaussian distributions with means informed by the assayed activities of components. However, some of the assayed activities, especially those from  $^{40}\text{K}$  groups and the LMFE-based  $^{60}\text{Co}$  groups, are large upper limits due to insufficient sensitivity. This creates a problem, as using these upper limits as the prior means will lead to an overfit of the data, as the upper limits are orders of magnitude larger than what is suggested by the strength of the  $^{40}\text{K}$  peaks in data. Therefore, it is important to find a solution for these

prior activity estimates that is far below the value of the assayed upper limits, yet not so low as to lead to an underfit of the data.

In initial tests of the Bayesian fitting algorithm, it was discovered that fits of the DEMONSTRATOR statistical model to simulated datasets were somewhat unreliable due to the limited statistics in the dataset and an over-dependence of the posterior result on the prior values. These specific tests are not discussed in this work because they used an older version of the PDF's described in section 4.2, but an updated set of these tests are described later in section 6.4. In any case, one idea put forth was to let one  $^{232}\text{Th}$  component group at a time float with an uninformative (flat) prior, and let all others float with their truncated gaussian priors. This lets the fitting algorithm focus on placing posterior counts into that particular  $^{232}\text{Th}$  group, which takes on a temporary role as the “Th excess” for the fit, and allows the algorithm to scale the  $^{40}\text{K}$  groups (and two  $^{60}\text{Co}$  groups) to match the environment.

Five fits were run against DS0-7 open and blind data in this scheme, where each fit allowed one particular  $^{232}\text{Th}$  component group to float with a flat prior. The component groups span the locations near and far to the detectors and are listed in the first row of tables 6.1 and 6.2. All  $^{40}\text{K}$  component groups with absolute assay measurements and uncertainties were fixed at those values, allowing the fitting algorithm to allocate the remaining  $^{40}\text{K}$  counts across those groups with high upper limits. The results of these five fits are shown in tables 6.1 and 6.2.

To evaluate which set of  $^{40}\text{K}$  and  $^{60}\text{Co}$  activities was the best to choose for the model, the performance of each fit between 800 and 1440 keV was analyzed. This is the region where the  $^{40}\text{K}$  spectrum is most prominent up until just below the  $^{40}\text{K}$  peak at 1460 keV. When calculating the root mean square of the pulls between the total model and data of each 10 keV bin inside the aforementioned energy window, the fit where the M1LMFES were allowed to float with a flat prior produced the lowest overall root mean square value. Therefore, the set of activity values in the “Th\_M1LMFES” column from table 6.1 was chosen as the set of prior means to use for future fitting. The associated prior uncertainty of each value was chosen as the calculated standard deviation of all five fitted activities for that particular component group. All of the activity values in the “Th\_M1LMFES” column are either comparable to or far below their corresponding upper limit values from the radioassay program, as shown in the last column of tables 6.1 and 6.2. Therefore, this revised set of prior means is far less likely to overfit the  $^{40}\text{K}$  and  $^{60}\text{Co}$  gamma peaks in data.

#### **Section 6.4: Simulated Dataset Studies**

In order to verify the efficacy of the statistical model when it is fit to data, it is important to start with a controlled set of data with known quantities. These data are constructed from the same set of simulations

Groups	Th_M1LMFEs	Th_M1Connectors	Th_M1CCPCables	Upper Limit
K_M1Connectors	4.42E-03	1.68E-03	2.02E-03	1.03E-03
K_M1CCPCables	1.05E-02	8.26E-03	7.40E-03	3.13E-03
K_OuterCuShield	8.82E-07	1.14E-02	7.48E-03	3.81E+00
K_M1CCF	2.24E-03	1.81E-03	1.31E-03	1.39E-03
K_M1CCN	7.65E-09	9.72E-09	3.38E-09	4.59E-03
K_M1DUStringCu	5.54E-10	2.12E-10	1.13E-09	4.68E-04
K_M1LMFEs	4.61E-10	4.70E-06	3.35E-06	1.09E-05
K_M1StringCables	2.41E-09	5.63E-09	1.51E-11	2.43E-04
K_M2Connectors	2.19E-03	2.35E-03	2.28E-03	1.03E-03
K_M2CCPCables	5.14E-03	5.84E-03	5.76E-03	3.13E-03
K_M2CCF	1.47E-03	1.44E-03	1.36E-03	1.39E-03
K_M2CCN	1.14E-08	4.18E-08	1.96E-09	4.59E-03
K_M2DUStringCu	1.44E-09	4.34E-09	5.04E-09	4.65E-04
K_M2LMFEs	7.11E-06	5.75E-06	5.87E-06	1.09E-05
K_M2Seals	4.90E-06	4.57E-06	4.95E-06	4.47E-06
K_M1StringCables	4.99E-10	8.48E-09	1.66E-09	2.43E-04
K_InnerCuShield	3.84E-08	4.96E-08	1.70E-09	6.83E-02
Co_M1LMFEs	1.07E-10	2.95E-10	6.22E-11	1.68E-04
Co_M2LMFEs	6.61E-11	9.04E-11	5.64E-10	1.68E-04

Table 6.1: Posterior activities of  $^{40}\text{K}$  and  $^{60}\text{Co}$  component groups in Bq that have upper limits in their assay measurements, following fits where a particular  $^{232}\text{Th}$  group was allowed to float with a flat prior and other groups floated with truncated gaussian priors.  $^{40}\text{K}$  groups with absolute assay measurements were fixed for these fits. The last column is the original upper limit on the assayed activity value.

Groups	Th_InnerCuShield	MTh_M1CCF	Upper Limit
K_M1Connectors	4.15E-03	3.23E-03	1.03E-03
K_M1CCPCables	1.06E-02	5.69E-03	3.13E-03
K_OuterCuShield	1.08E-06	5.86E-03	3.81E+00
K_M1CCF	2.30E-03	1.31E-03	1.39E-03
K_M1CCN	5.13E-09	3.38E-09	4.59E-03
K_M1DUStringCu	6.30E-11	3.64E-10	4.68E-04
K_M1LMFEs	2.56E-10	3.35E-06	1.09E-05
K_M1StringCables	6.75E-10	3.14E-09	2.43E-04
K_M2Connectors	1.68E-03	2.43E-03	1.03E-03
K_M2CCPCables	3.72E-03	6.17E-03	3.13E-03
K_M2CCF	1.42E-03	1.36E-03	1.39E-03
K_M2CCN	1.00E-08	1.96E-09	4.59E-03
K_M2DUStringCu	3.26E-09	1.19E-09	4.65E-04
K_M2LMFEs	8.66E-06	5.87E-06	1.09E-05
K_M2Seals	4.56E-06	4.95E-06	4.47E-06
K_M1StringCables	1.35E-04	4.43E-09	2.43E-04
K_InnerCuShield	3.07E-08	1.77E-08	6.83E-02
Co_M1LMFEs	4.79E-11	4.24E-12	1.68E-04
Co_M2LMFEs	3.49E-11	3.98E-11	1.68E-04

Table 6.2: Posterior activities of  $^{40}\text{K}$  and  $^{60}\text{Co}$  component groups in Bq that have upper limits in their assay measurements, following fits where a particular  $^{232}\text{Th}$  group was allowed to float with a flat prior and other groups floated with truncated gaussian priors.  $^{40}\text{K}$  groups with absolute assay measurements were fixed for these fits. The last column is the original upper limit on the assayed activity value.

that inform the Assay-Based Background Model, as described in sections 4.1 through 4.3. By fitting to these simulated datasets, the behavior of the fitting algorithm is not only evaluated at the same level of statistics as the DEMONSTRATOR dataset, it may optionally be evaluated at higher levels of statistics. This provides insight into how such a fitting algorithm would perform in a high-background dataset.

#### 6.4.1: Overview and Approach

Construction of the simulated datasets is a relatively simple matter of sampling the Assay-Based Background Model. As a reminder, the Assay-Based Background Model spectrum is informed by the simulated PDFs of each decay chain and component group tuple and a scaling factor that corresponds to the assayed activity of each tuple. A random variable of the `rv_discrete` class from the Python package `SciPy` [53] is created for each tuple. The `rv_discrete.rvs()` function is then called which samples the distributions associated with each tuple with a number of samples that corresponds to the number of integrated counts in the Assay-Based Background Model for that tuple multiplied by a special factor. When creating simulated datasets with a similar amount of integrated counts as the DEMONSTRATOR datasets, this factor is simply 1. For simulated datasets that represent “high-statistics datasets”, a factor much greater than 1 is used, typically 1000 in this work.

The samples are organized into their own `numpy` arrays that represent each tuple and combine to form the simulated dataset. The special factor used to scale the number of counts is also used to scale the activity values for each tuple prior before fitting to match the “high-statistics dataset” environment.

After the sampling of the posterior is complete, the fit is analyzed in two ways. The first is to look at plots of the energy spectrum from 100 to 2620 keV comparing the fitted model and the dataset (see figure 6.9 for an example). A set of pulls for each 10 keV bin are calculated in a plot below the spectral plot, to give an overall sense of how well the fit performed. The second is to look at the number of integrated counts that are fit into each component group and decay chain. This information is displayed in bar graph form, with a gray bar representing the number of integrated counts in a component group as given by the prior activity, and a colored bar representing the number of integrated counts in the component group as given by the posterior activity. The color of the bar also represents the relevant decay chain, with green representing  $^{232}\text{Th}$ , pink representing  $^{40}\text{K}$ , and so on. If the fit involves a simulated dataset based on activities that do not completely match the prior activities, a blue bar placed between the gray and colored bars will represent the number of integrated counts in the simulated dataset for that component group. The remaining sections and subsections will represent fitting results in these two formats.

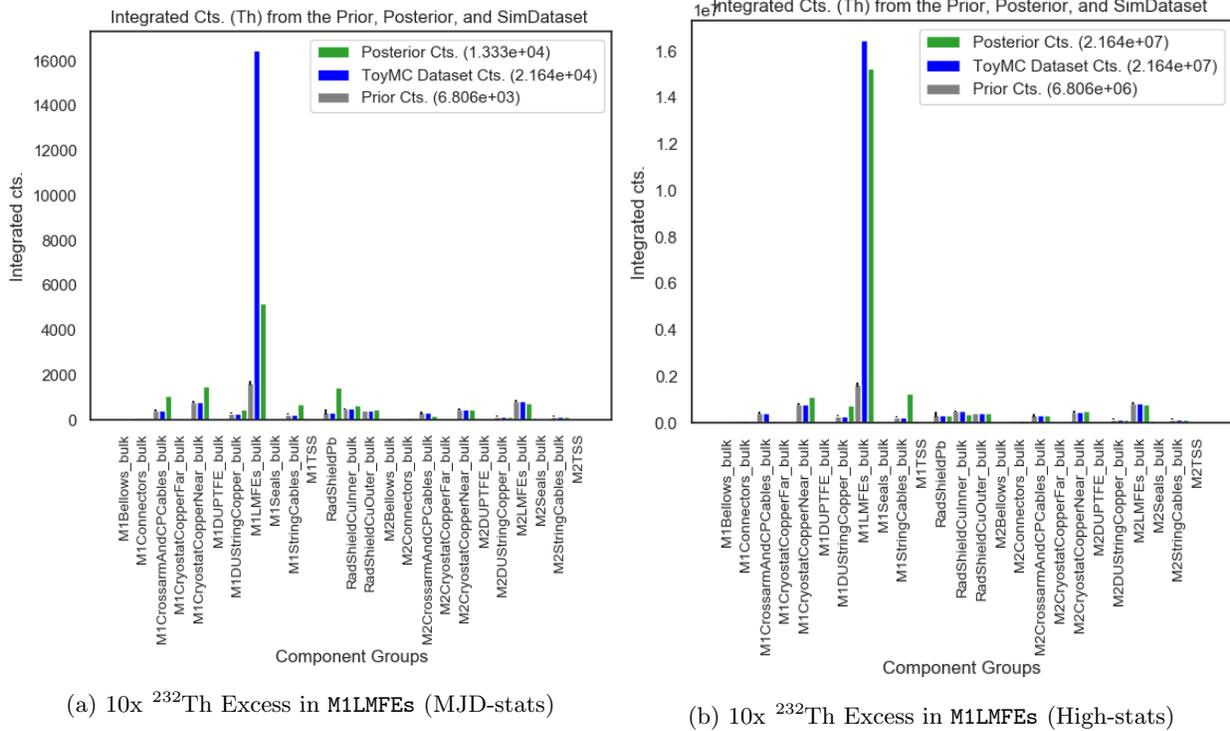
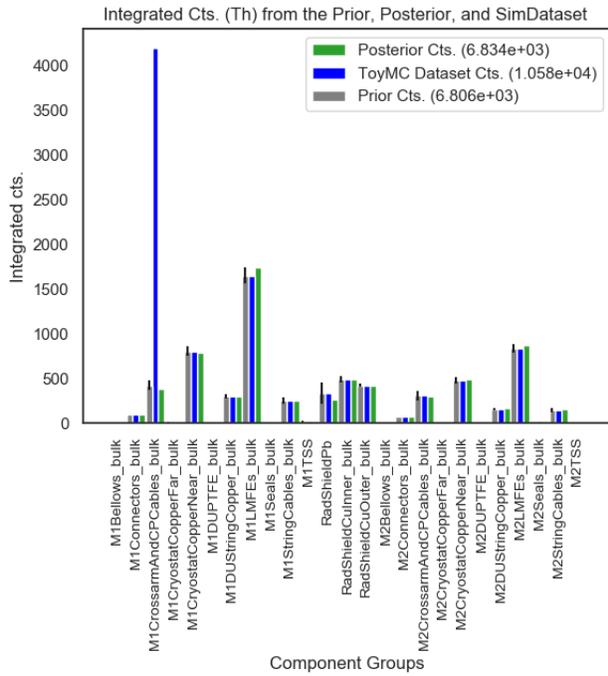


Figure 6.1: Fits to simulated datasets with a  $10x$   $^{232}\text{Th}$  excess in the M1LMFEs group.

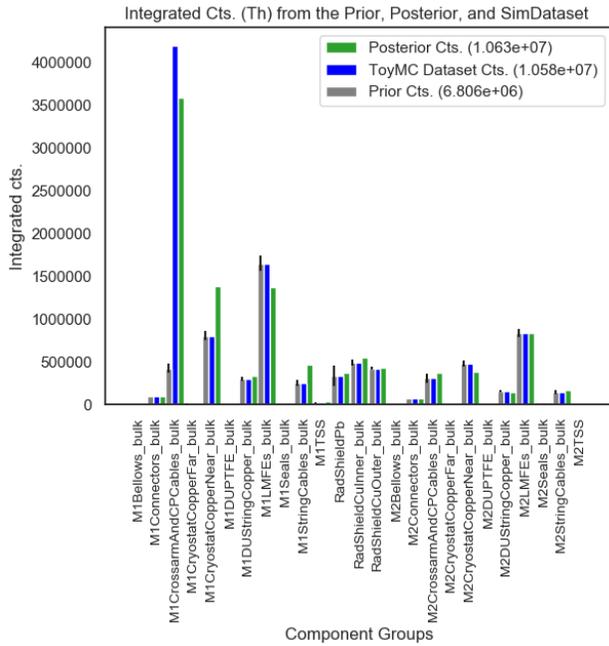
#### 6.4.2: Fits with a $^{232}\text{Th}$ excess in Various M1 Component Groupings

Two simulated datasets were constructed by sampling the assay-based model, with one at “MJD-stats” and the other at “high-stats” in the manner described by the previous section. A factor of  $10$   $^{232}\text{Th}$  excess is inserted into one M1  $^{232}\text{Th}$  component group to see if the fitting algorithm can make up for the excess number of counts in the posterior that are in the simulated dataset. All groups are floated in these fits with truncated Gaussian priors where the mean of the prior is the assayed activity value and the width of the prior is based on the uncertainty of the assayed activity value. Particular  $^{40}\text{K}$  and  $^{60}\text{Co}$  groups have prior activity values that are estimates from previous Bayesian fits rather than the assay program, see section 6.3.

When placing the excess in the M1LMFEs group, the fit to the “MJD-stats” dataset produces the  $^{232}\text{Th}$  integrated counts result in figure 6.1a. The fitting algorithm has a difficult time placing counts in the posterior to make up for the  $^{232}\text{Th}$  excess in the simulated dataset. There is too much of a dependence on either the prior mean or prior uncertainty. When fitting to the “high-stats” dataset, the fit produces the result in figure 6.1b. The fitting algorithm is able to see the excess in M1LMFEs and place the corresponding posterior counts in that group. This same behavior for both the “MJD-stats” and “high-stats” datasets is seen for other groups, like the M1CrossarmAndCPCables group and the PbShield group. See figures 6.2a through 6.3b.

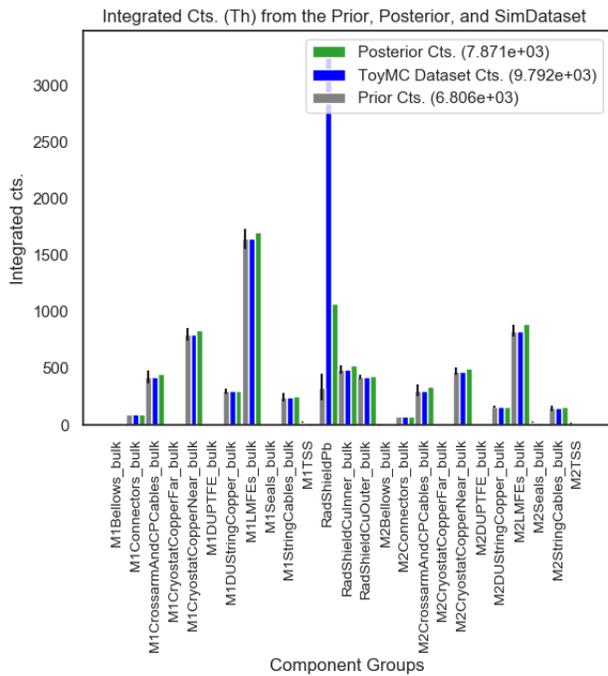


(a)  $10x$   $^{232}\text{Th}$  Excess in **M1CrossarmAndCPCables** (MJD-stats)

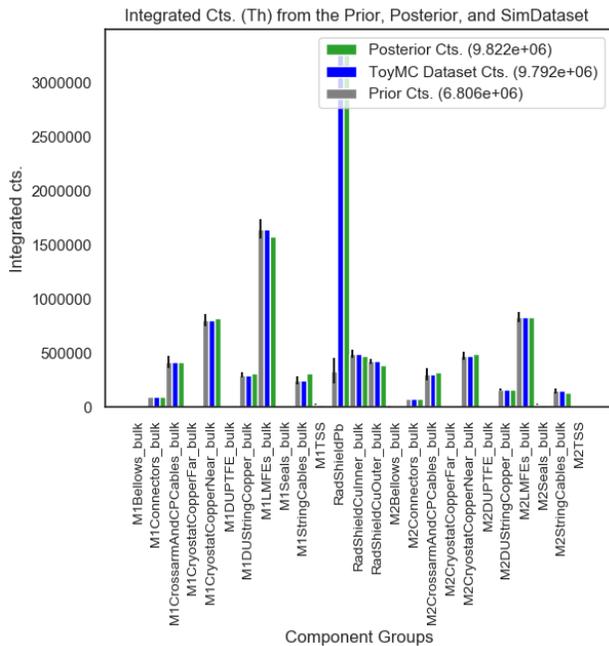


(b)  $10x$   $^{232}\text{Th}$  Excess in **M1CrossarmAndCPCables** (High-stats)

Figure 6.2: Fits to simulated datasets with a  $10x$   $^{232}\text{Th}$  excess in the **M1CrossarmAndCPCables** group.

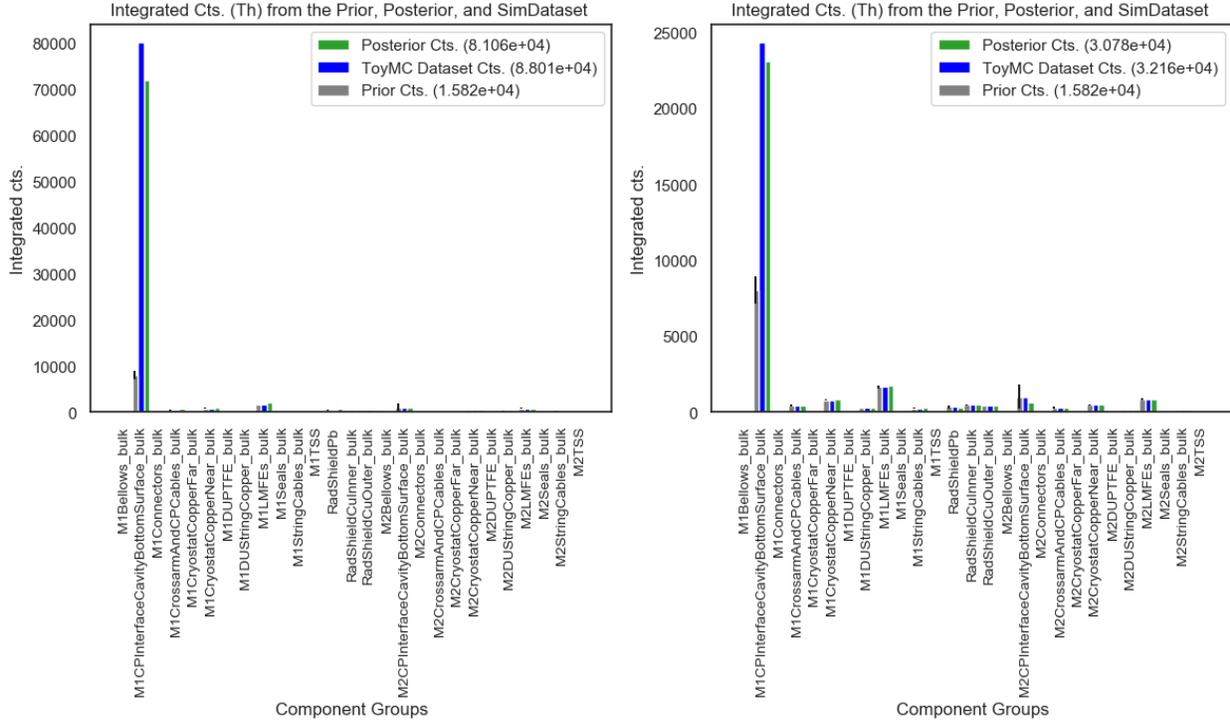


(a)  $10x$   $^{232}\text{Th}$  Excess in **PbShield** (MJD-stats)



(b)  $10x$   $^{232}\text{Th}$  Excess in **PbShield** (High-stats)

Figure 6.3: Fits to simulated datasets with a  $10x$   $^{232}\text{Th}$  excess in the **RadShieldPb** group.



(a)  $10x$   $^{232}\text{Th}$  Excess in **M1Cavity** (MJD-stats)

(b)  $3x$   $^{232}\text{Th}$  Excess in **M1Cavity** (MJD-stats)

Figure 6.4: Fits to a “MJD-stats” simulated dataset with some  $^{232}\text{Th}$  excess in the **M1Cavity** group.

#### 6.4.3: Fits with a $^{232}\text{Th}$ Excess in the **M1Cavity** Group

The same type of fit as described in section 6.4.2 is run, but now including the **M1CPInterfaceCavityBottomSurface** group (referred to as **M1Cavity** for the remainder of the chapter) in our list of component groups. A factor of  $10$   $^{232}\text{Th}$  excess is placed in this group relative to its assay prior (an upper limit of  $0.57$  mBq), and the fit to the “MJD-stats” dataset produces the result in figure 6.4a. The number of counts in the posterior is a lot closer to what is in the simulated dataset than when putting the  $^{232}\text{Th}$  excess in other groups. This is also true if we change the  $^{232}\text{Th}$  excess factor to around  $3x$ , to match the difference between the assay upper limit ( $0.57$  mBq) and the best estimate for the true cavity activity given the peak scaling studies discussed in section 4.4 (about  $1.78$  mBq, see figure 6.4b).

One noticeable difference between the fit with a  $^{232}\text{Th}$  excess in the **M1Cavity** group versus fits with the  $^{232}\text{Th}$  excess in other groups, is that the number of  $^{232}\text{Th}$  integrated counts in the **M1Cavity** prior is already a factor of four or greater than that of other groups, and therefore could potentially be influencing the posterior. Another fit to the “MJD-stats” dataset is done with the number of counts in the **M1Cavity** prior reduced artificially by a factor of four. The result of this fit is shown in figure 6.5, where the performance of the fit remains unchanged.

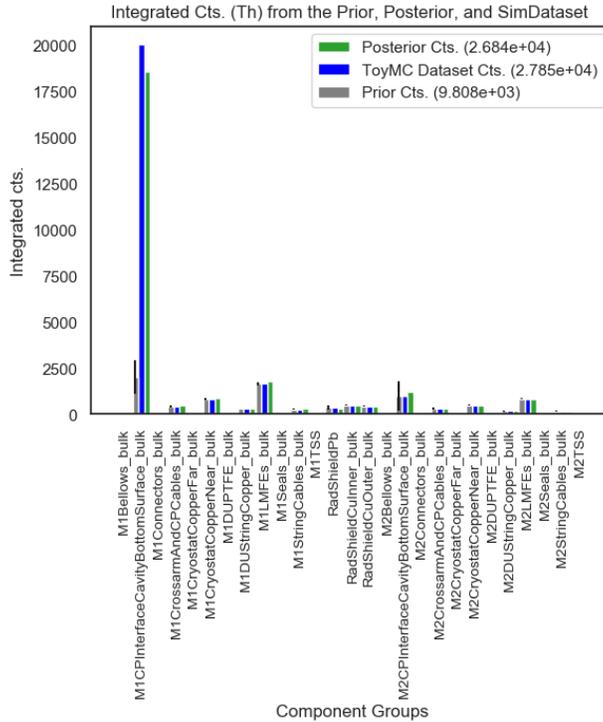


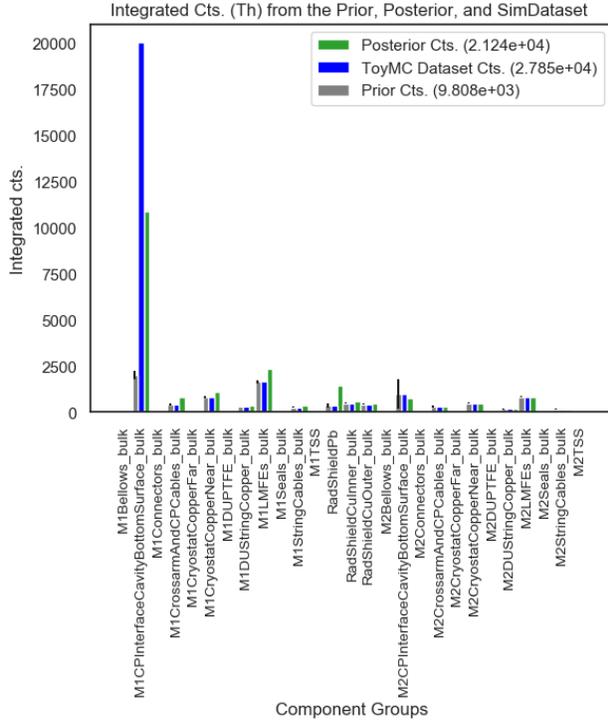
Figure 6.5: Fit to a “MJD-stats” simulated dataset with a  $3 \times 232\text{Th}$  excess in the M1Cavity group with a 0.25x prior.

Why do these fits perform much better than the other models? It is possible that the M1Cavity group is distinct enough from other groups that the fitting algorithm is able to identify the excess a lot more easily, but it is also possible the posterior receives more counts in the fit because the uncertainty on the M1Cavity group is larger than all other groups. To resolve this degeneracy, additional fits to simulated datasets were done with changes made to the prior uncertainties.

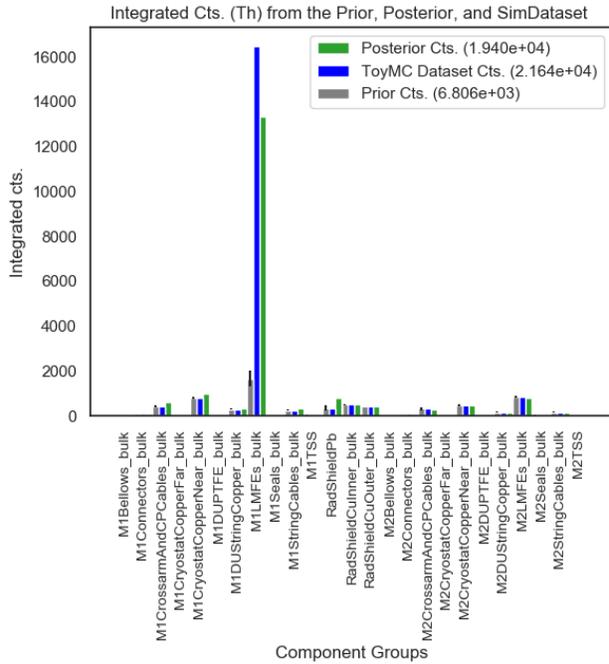
#### 6.4.4: Fits with a Change to the Prior Uncertainties

The same fit of the reduced prior on the M1Cavity group to the “MJD-stats” simulated dataset with around a factor of  $3 \times 232\text{Th}$  excess in the dataset is run, but this time the uncertainty on the M1Cavity group prior is tightened by a factor of five. When this is done, the fit produces the result shown in figure 6.6a. The fitting algorithm no longer has the ability to fit a number of counts in the posterior that are needed to match the counts in the simulated dataset. This tightening of the M1Cavity group’s prior uncertainty illustrates that the posterior is in fact dependent on the prior uncertainty value of the M1Cavity relative to other prior uncertainty values.

To evaluate the reciprocal situation, the same fit from section 6.4.2 is run with a factor of  $10 \times 232\text{Th}$  excess in the M1LMFES group against the “MJD-stats” simulated dataset, but this time the uncertainty on



(a) Fit to a “MJD-stats” simulated dataset with a  $3x$   $^{232}\text{Th}$  excess in the **M1Cavity** group with a 0.25x prior and a 0.2x prior uncertainty.



(b) Fit to a “MJD-stats” simulated dataset with a  $10x$   $^{232}\text{Th}$  excess in the **M1LMFes** group with a 5x prior uncertainty.

Figure 6.6: Comparison of fits to a “MJD-stats” simulated dataset when prior uncertainties are changed.

the **M1LMFes** group prior is loosened by a factor of 5. When this is done, the fit produces the result in figure 6.6b. The fitting algorithm is able to fit many more counts in the posterior than before and is less restricted by the prior. In fact, the **M1LMFes** fit to the simulated dataset with a loosened prior uncertainty fits more counts into the proper posterior than the **M1Cavity** fit to the simulated dataset with a tightened uncertainty.

These fits highlight the Bayesian fitting algorithm’s dependence on the prior uncertainty values, which makes sense as the shape of the priors are truncated Gaussian distributions with widths defined by these prior uncertainties. If one loosens or tightens these uncertainties, the priors can change dramatically and have a large effect on the fit when the dataset is small, or at “MJD-stats”.

#### 6.4.5: Fits with the **M1Cavity** Region in the Simulated Dataset, But Not in the Component Groupings List

Where does the fit place counts when the **M1Cavity** group is in the simulated dataset, but not in the list of component groups? The result of this fit is shown in figure 6.7a. The fit distributes posterior counts into particular groups, specifically **RadShieldPb**, **M1LMFes**, **M1CroystatCopperNear**, and **M1CrossarmAndCPCables**. At first this does not make sense, since it is fitting counts that would otherwise go into the **M1Cavity** group

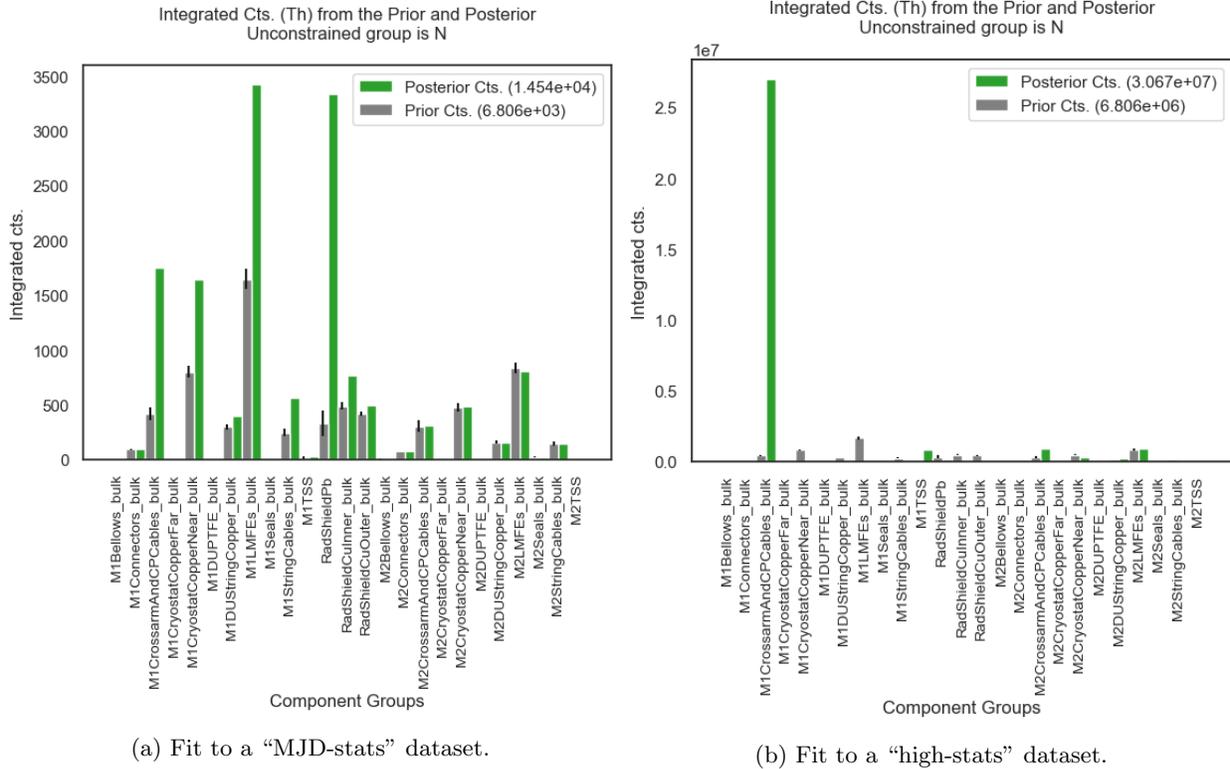


Figure 6.7: Fit with the `M1Cavity` group in the dataset but not in the list of components. Ignore the second line of the plot title. The number of counts in the simulated dataset for each component is equal to the prior counts, so the blue bars are not shown here.

into “near” groups like the `M1LMFes` and “far” groups like `RadShieldPb`. However, it is clear that these four groups have the largest uncertainties on their priors of all the `M1` groups and shielding groups involved in the fit. Again, the fitting algorithm when fitting to a “MJD-stats”-like dataset is heavily influenced by the prior uncertainties. When fitting to a “high-stats” dataset, almost all of the counts that would normally go into the `M1Cavity` group go instead into the `M1CrossarmAndCPCables` group. This makes sense, as that is one of the closest groups to the location of the `M1Cavity` region (see figure 6.7b).

#### 6.4.6: Fits with Flat Priors

If dependence on the prior uncertainties is an issue, should those uncertainties be ignored by the fit? In other words, what if truncated Gaussian priors are abandoned in favor of truncated flat priors, where the prior uncertainty has no influence. The only information from the prior in this scenario is the starting point of the chain, which will be at the prior activity value.

The initial set of simulated dataset fits from sections 6.4.2 and 6.4.3 were repeated against the “MJD-stats” simulated dataset, injecting some factor of a  $^{232}\text{Th}$  excess into particular component groups. The

results of these fits are shown in figures 6.8a through 6.8d. For fits with the excess in components like the `M1Cavity` and `M1LMFEs`, the posterior can make up for most of the counts that are in the simulated dataset. Alternatively, fits with the excess in the `M1CrossarmAndCPCables` and `PbShield` groups perform rather poorly.

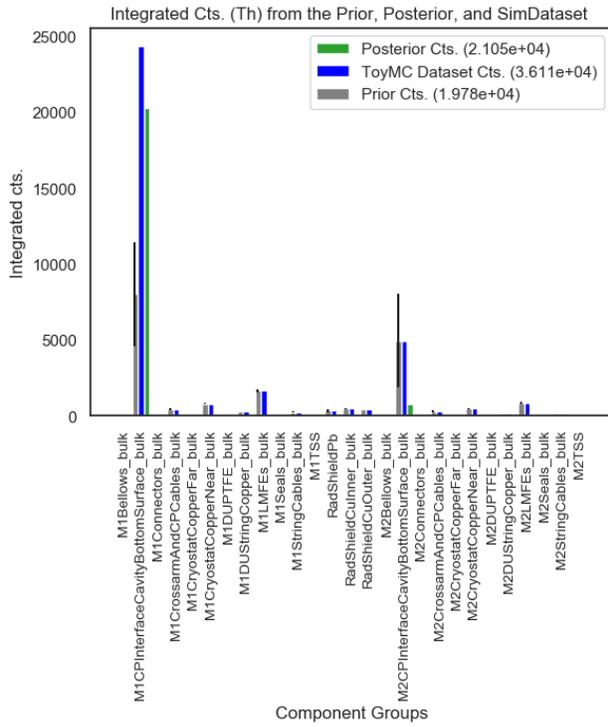
However, in all four fits the pertinent result is in the energy spectra (see figures 6.9a through 6.9d), where the data is severely underfit by the model, especially when compared to fits with Gaussian priors. This is further supported by integrated count plots of other decay chains (see figures 6.10a through 6.10d), where aside from the distinct  $2\nu\beta\beta$  spectrum and the  $^{60}\text{Co}$  spectrum in natural detectors, the fit struggles to place counts in the posteriors of component groups. Without any tight bound on the prior and limited information from the dataset, it seems the fit defaults to not placing activity in the component groups. Thus, attempting to fit to the `DEMONSTRATOR` dataset with flat priors seems unworkable with the Bayesian fitting algorithm.

#### 6.4.7: Simulated Dataset Study Conclusions

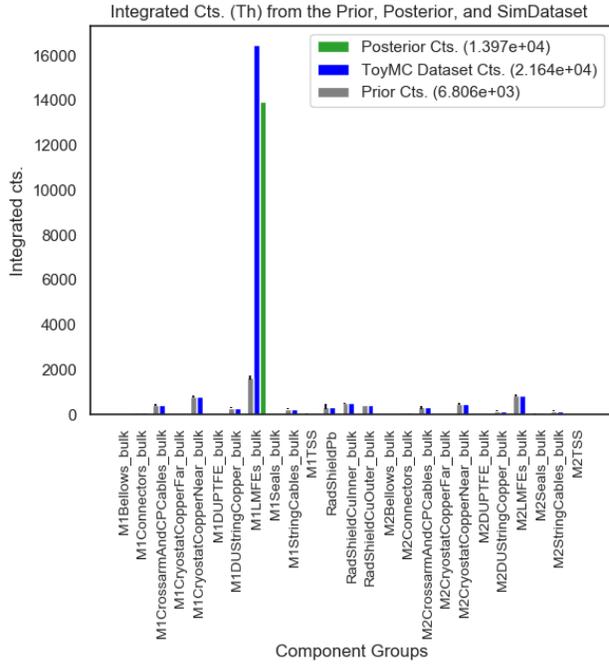
Simulated dataset studies with “MJD-stats” datasets show a strong dependence of the posterior on the relative size of the prior uncertainties. A tighter prior uncertainty means fewer counts are fit into a group while a looser prior uncertainty means more counts are fit into the group. If the dataset were about a thousand times larger than the amount of data collected by the `DEMONSTRATOR`, fits to “high-stats” simulated datasets show a large reduction of this dependence on prior uncertainties.

This result gets at some of the pitfalls of using Bayesian inference methods. When there is an unexpected element in a dataset (like the  $^{232}\text{Th}$  excess) that is not accounted for in the prior, and there are not enough statistics in the dataset to highlight particular features in this multi-parameter space, the application of these inference methods begin to break down. Complicating matters is the fact that some of the priors used in this study are somewhat unreliable, either because they had relatively high limits since the radioassay was not sensitive enough to determine their true activity values (many of the  $^{40}\text{K}$  groups and a couple  $^{60}\text{Co}$  groups), or because their true value changed after the radioassay program concluded. For a theoretically longer running `MAJORANA DEMONSTRATOR` experiment, these issues would resolve with higher statistics.

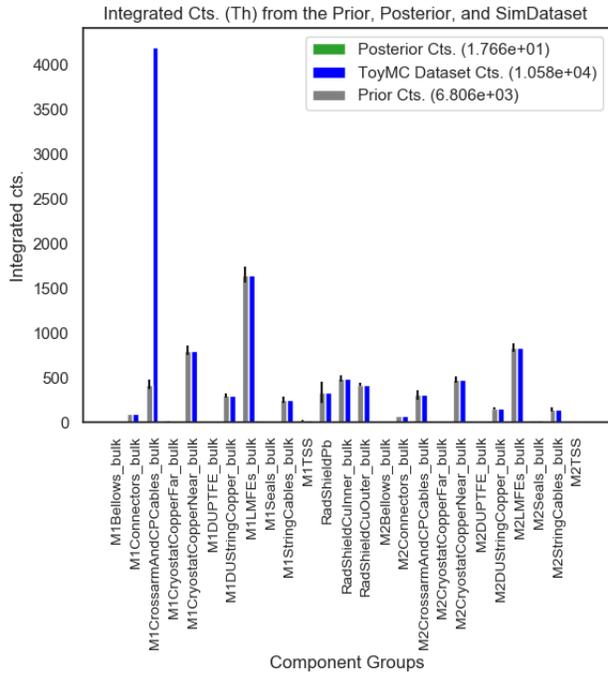
Simulated Dataset studies where the truncated Gaussian priors are replaced with uninformative flat priors perform worse for every component group and decay chain except for those with unique features in the spectrum, like the  $2\nu\beta\beta$  spectrum and  $^{60}\text{Co}$  spectrum in the natural detectors. With such a high-dimensional posterior to sample from, it is critical to have as much information as possible in order to resolve degeneracies between similarly shaped PDFs. An appropriate range of activity values is necessary information for this implementation of inference methods to work.



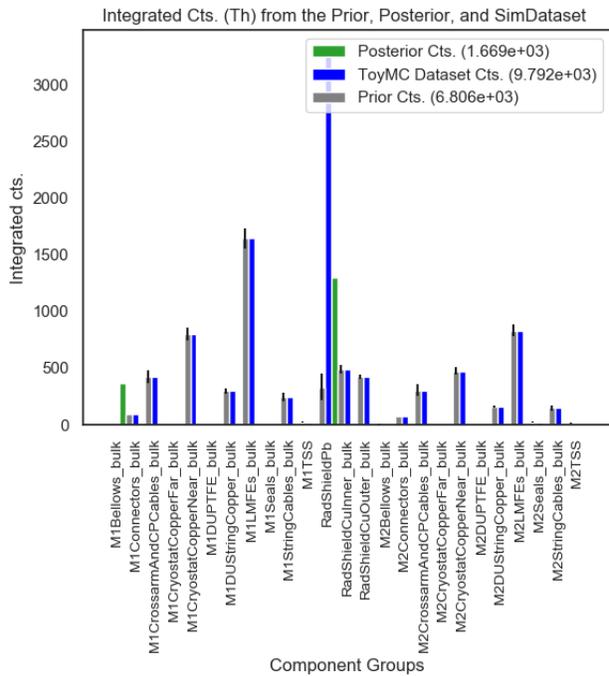
(a)  $3x$   $^{232}\text{Th}$  Excess in the M1Cavity group with flat priors.



(b)  $10x$   $^{232}\text{Th}$  Excess in M1LMFEs with flat priors.



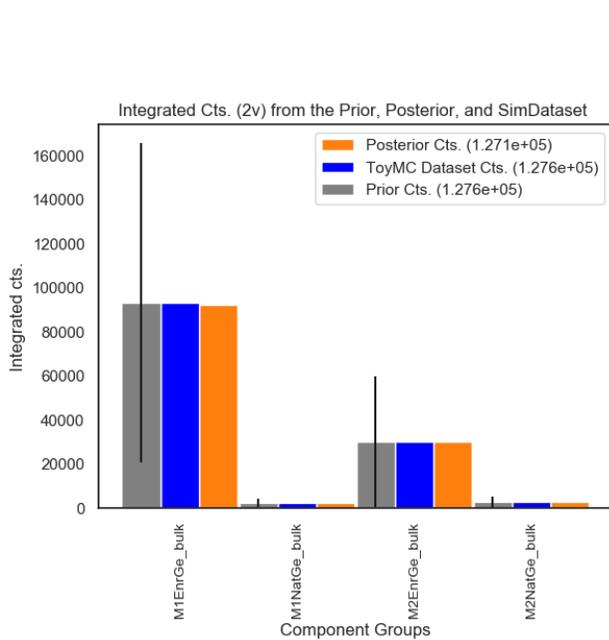
(c)  $10x$   $^{232}\text{Th}$  Excess in M1CrossarmAndCPCables with flat priors.



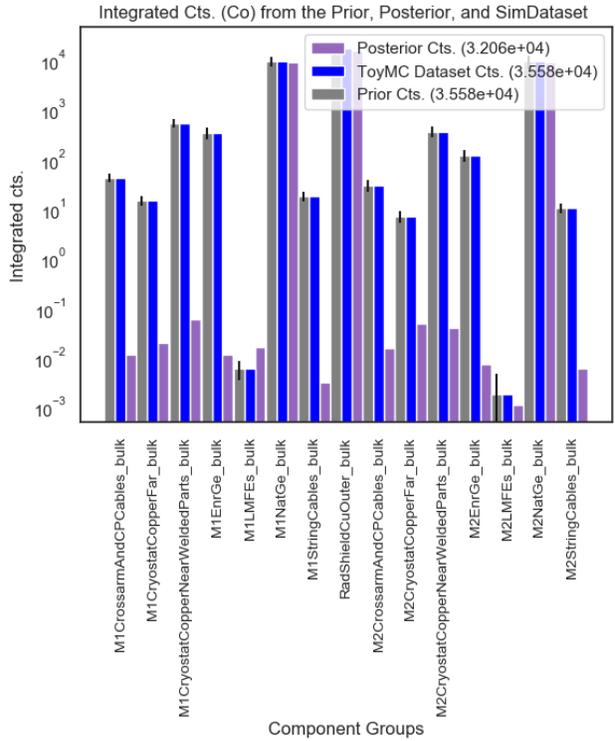
(d)  $10x$   $^{232}\text{Th}$  Excess in PbShield with flat priors.

Figure 6.8: Fits with flat priors to a “MJD-stats” simulated dataset.

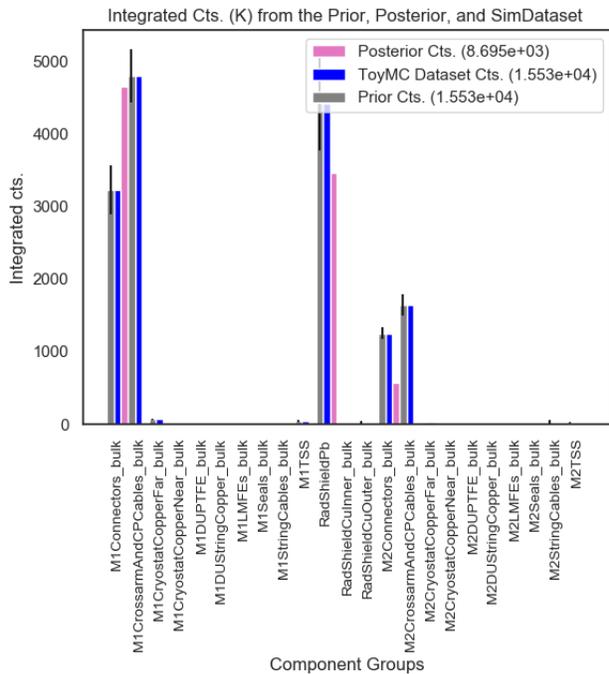




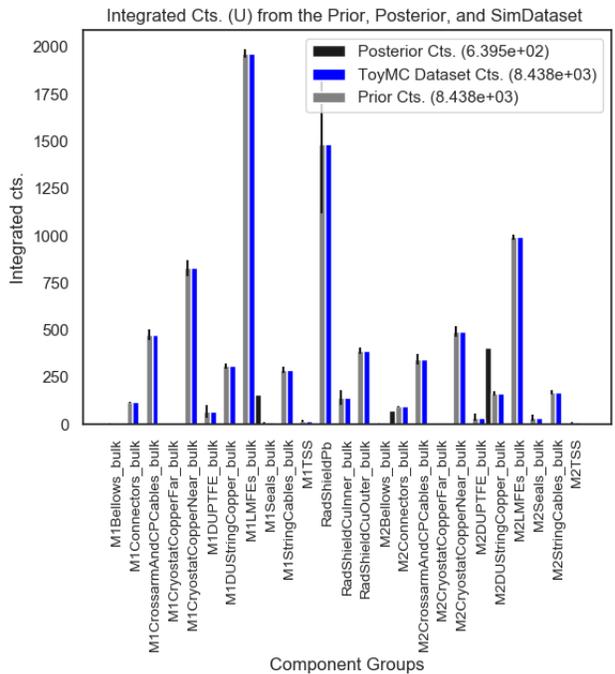
(a)  $3x$   $^{232}\text{Th}$  Excess in the M1Cavity group with flat priors, 2v integrated counts.



(b)  $3x$   $^{232}\text{Th}$  Excess in the M1Cavity group with flat priors,  $^{60}\text{Co}$  integrated counts.



(c)  $3x$   $^{232}\text{Th}$  Excess in the M1Cavity group with flat priors,  $^{40}\text{K}$  integrated counts.



(d)  $3x$   $^{232}\text{Th}$  Excess in the M1Cavity group with flat priors, U integrated counts.

Figure 6.10: Integrated count plots from various decay chains. These are from fits using flat priors to a “MJD-stats” simulated dataset, specifically with a  $3x$   $^{232}\text{Th}$  excess in the M1Cavity group.

## Section 6.5: Determining the Most Likely Model with Bayes Factors

Based on the work discussed in chapter 4, there is evidence that the source of the  $^{232}\text{Th}$  excess seen in the data is the region around the cavity of the M1 thermosyphon. A set of PDFs were created from simulations of the bottom surface of this component and used in the fits to simulated datasets discussed in the previous section. Since there is a degree of belief that the M1Cavity component group will describe the  $^{232}\text{Th}$  excess in the data, one may hypothesize that a model that includes the M1Cavity component group will describe the data with higher probability than any other model.

This hypothesis can be tested by evaluating a metric known as the Bayes Factor. The Bayes Factor is calculated by taking the ratio of the evidences (see chapter 5) of two different models:

$$BF = \frac{P(D|M_1)}{P(D|M_2)} \quad (6.6)$$

A Bayes Factor of this form assumes a priori that one model is not inherently favored over the other (see [50]). A Bayes Factor greater than 10 usually means that there is confidence the model in the numerator describes the data with higher probability than that model in the denominator. The following section describes how such a Bayes Factor may be calculated and what method is necessary for calculating factors that may test the aforementioned hypothesis.

### 6.5.1: Methodology

Following the rules of probability, Bayes theorem (equation 5.1) can be rewritten as an expression for the marginal likelihood:

$$P(D|M) = \int P(D|\theta, M)P(\theta|M)d\theta \quad (6.7)$$

where  $P(D|\theta, M)$  represents the Poisson-distributed likelihood of our model and  $P(\theta|M)$  represents the truncated Gaussian priors of our model. In essence, the marginal likelihood is the integration of the likelihood over the parameter space of the prior. In practice, computing the integral analytically is usually impractical, so instead Monte-Carlo techniques are used by drawing samples  $\{\theta_l\}$  from the prior [50]:

$$P(D|M) = \frac{1}{N_s} \sum_{l=1}^{l=N_s} P(D|\theta_l, M) \quad (6.8)$$

where  $P(D|\theta_l, M)$  is the likelihood evaluated at a sampled activity from the prior distribution, and  $N_s$  is the number of samples drawn from the prior. Normally this would be the right approach to take in

order to calculate the marginal likelihoods. However, as evidenced by the underestimate of the Assay-Based Background Model when compared to data, the component priors span a parameter space of activities that in some cases differ from data by a large margin. Therefore, the value of the likelihood at sampled prior values as evaluated in equation 6.8 will be relatively low and uninformative. A better parameter space to draw from would be that of the posterior, which will more closely resemble the range of activities that are seen in the data.

The concept of “importance sampling” can be used here to facilitate this, as outlined by Owens [54]. The expectation value of a function  $f(x)$  is:

$$E[f(x)] = \int f(x)p(x)dx \quad (6.9)$$

where  $p(x)$  is a PDF that defines the space that the integrand  $f(x)$  will be integrated over. Multiplying in the ratio of some other positive PDF over itself results in:

$$\int \frac{f(x)p(x)}{q(x)}q(x)dx = E\left[\left(\frac{f(x)p(x)}{q(x)}\right)\right] \quad (6.10)$$

where now the expectation value of the function is being evaluated over the parameter space defined by  $q(x)$ , with an adjustment factor of  $p(x)/q(x)$  called the “likelihood ratio”. The format of equation 6.7 matches that of equation 6.9, if we let  $f(x)$  represent the likelihood and  $p(x)$  represent the prior. Continuing this correspondence through equation 6.10, a  $q(x)$  is constructed to have the same shape as the prior, a truncated Gaussian, but with a mean equal to the posterior activity value rather than that of the prior. Making these substitutions and rewriting the structure of equation 6.10 in terms of a numerical Monte Carlo calculation that draws samples  $\{\theta_l^q\}$ , the evaluation of the marginal likelihood becomes:

$$P(D|M) = \frac{1}{N_s} \sum_{l=1}^{l=N_s} P(D|\theta_l^q, M) \frac{P(\theta^q|M)}{q(\theta^q)} \quad (6.11)$$

Remember from equation 5.6 that the total likelihood of one submodel is the product of the likelihoods for each individual 10 keV energy bin. In order to simplify the marginal likelihood calculation, the logarithm of the marginal likelihood is taken, leading to the following evaluation of the total log marginal likelihood:

$$\log(ML)_{\text{total}} = \sum_{\text{submodels}} \left( \frac{1}{N_s} \sum_{l=1}^{l=N_s} \left( \sum_{j=1}^{j=N_{\text{bins}}} \log(P_j(D|\theta_l^q, M)) \right) + \log(P(\theta^q|M)) - \log(q(\theta^q)) \right) \quad (6.12)$$

To evaluate the hypothesis that the model containing the M1Cavity component group will describe the

data with higher probability than any other model, several single-component fits are run where one <sup>232</sup>Th component group is allowed to float and all other component groups are fixed to their prior activity values. Other than the fit that floats the `M1Cavity` component group, the `M1Cavity` component group is not included in the list of components. Following every fit, a method is called from the `scipy.stats` [53] class that randomly samples from a truncated Gaussian distribution with a mean equal to the posterior activity for the floating component group. The total number of samples for this work was 10,000 per submodel. Using `scipy.stats`, the log probabilities were evaluated for the individual likelihoods, the floating component prior, and our chosen importance sampling distribution using each sampled value  $\theta^q$ . These log probabilities were then summed together as defined by equation 6.12.

The total marginal likelihood of the fit that floats the `M1Cavity` group is divided one at a time by the total marginal likelihoods of all other fits, producing a set of Bayes Factors. Since these marginal likelihoods are in logarithmic form, the log marginal likelihoods from fits that do not contain the `M1Cavity` group are subtracted from the log marginal likelihood of the fit that floats the `M1Cavity` group. If the resulting (log) Bayes Factors are greater than  $\ln(10)$ , then the hypothesis is accepted.

#### 6.5.2: Results with Fits to DS0-7 Open and Blind Data

Single-component fits were done to DS0-7 open and blind data in order to evaluate which model describes the data with the highest probability. Each model consists of one M1 thorium component group that floats while every other group is fixed. If `M1Cavity` is not floating in the fit, then the cavity groups are not included in the fit. A summary of the results is shown in table 6.3. Note that the marginal likelihood and Bayes Factors are in logarithmic form. This may seem a bit alarming at first as the marginal likelihoods are quite small and the Bayes Factors are quite large, but this is merely the result of evaluating the log marginal likelihood of each energy bin (reasonable values in the single digits), and then summing over all 252 bins and all 17 submodels (split by dataset, module, det. type, and multiplicity) to arrive at a total value for the fit.

The single-component fit of the `M1Cavity` group has a much higher log marginal likelihood than all other groups, and thus produces very large Bayes Factors when compared to all other models, supporting the hypothesis that the model that includes the `M1Cavity` group describes the data with higher probability than all other models.

#### 6.5.3: Results with Fits to Module 2 DS3-6 Open and Blind Data

There has not been any conclusive evidence from data or simulation studies that points to a particular component group in Module 2 as the source for the (smaller) excess in that module. It is possible that

Groups	M1Cavity	M1LMFEs	M1Connectors	M1CCF	M1DUPTFE	M1Seals
Log ML	-11881	-35592	-20302	-20284	-20278	-20312
Log BF	n/a	23710	8421	8403	8396	8431
Groups	M1TSS	OuterCu	M1Bellows	M1CCPCables	M1CCN	M1DUStringCu
Log ML	-20742	-23870	-20294	-39578	-36258	-24271
Log BF	8861	11989	8412	27696	24377	12390
Groups	M1StringCables	PbShield				
Log ML	-29203	-39822				
Log BF	17322	27940				

Table 6.3: Log marginal likelihoods and log Bayes Factors of several component groups from single component fits where one  $^{232}\text{Th}$  group is allowed to float and all other remain fixed. All DS0-7 open and blind data are used in these fits. All log Bayes Factors are compared to the M1Cavity group.

running single component fits of Module 2  $^{232}\text{Th}$  component groups and calculating the Bayes Factors for those models will point to a particular group being favored over all other groups in explaining the Module 2 excess. The same set of fits were run as described in the previous section, but this time only against Module 2 data with Module 2  $^{232}\text{Th}$  groups floating one at a time. Again, fits where the M2Cavity group was not floating did not include the M2Cavity group in the fit. The results are shown in table 6.4.

The results here are surprising, as the group with the highest marginal likelihood (and Bayes Factor when compared to all other models) is the M2Cavity. This doesn't make much sense, as freely floating Bayesian fits to data (to be described in section 6.6) show that the M2Cavity should fit to an activity near zero, and there is no evidence in the data or simulations that the cavity in Module 2 should be favored as a  $^{232}\text{Th}$  excess source. Remember, however, that the uncertainty on the prior for the M2Cavity is the same as that for its Module 1 counterpart, and relatively large compared to the other groups. If prior uncertainties played a large role in influencing the results of fits to simulated datasets, it would make sense to play a large role in single-component fits to data. To confirm, it is necessary to do another set of single-component fits to Module 1 data only, and then tighten the uncertainty on the M1Cavity group to see if the log marginal likelihood decreases.

#### 6.5.4: Results with Fits to Module 1 DS0-7 Open and Blind Data

The results of the single-component fits to Module 1 data are shown in table 6.5. As expected, the model with the M1Cavity group has the largest log marginal likelihood. Is this still the case if the prior uncertainty is tightened by a factor of five? The results of such a single component fit are shown in table 6.6. The

Groups	M2Cavity	M2LMFEs	M2Connectors	M2CCF	M2DUPTFE	M2Seals
Log ML	-3084	-4335	-4116	-4115	-4115	-4119
Log BF	n/a	1251	1032	1031	1031	1035
Groups	M2TSS	OuterCu	M2Bellows	M2CCPCables	M2CCN	M2DUStringCu
Log ML	-4122	-4140	-4116	-4451	-4355	-4164
Log BF	1038	1056	1032	1367	1271	1080
Groups	M2StringCables	PbShield				
Log ML	-4244	-4380				
Log BF	406	502				

Table 6.4: Log marginal likelihoods and log Bayes Factors of several component groups from single component fits where one  $^{232}\text{Th}$  group is allowed to float and all other remain fixed. Only M2 DS3-6 data is used in these fits. All log Bayes Factors are compared to the M2Cavity group. Note that these Bayes Factors are smaller in general than those in table 6.3. Only 4 submodels are used when fitting to Module 2 data, so the log marginal likelihood of each single-component fit is artificially reduced due to there being fewer submodels than the original set of 17, which leads to smaller log Bayes Factors.

log marginal likelihood drops by noticeable margin, and the Bayes Factors are reduced, but ultimately both metrics still favor the model containing the M1Cavity as the best descriptor of the data.

#### 6.5.5: Results when Changing Other Prior Uncertainty Values

One configuration of priors to consider is the “worst case scenario”, where the uncertainty on the M1Cavity prior is tightened by a factor of five and the uncertainty on another group, like the M1LMFEs, is loosened by a factor of five. The results of this scenario are shown in table 6.7. While the log marginal likelihood of the M1LMFEs group does increase by a noticeable margin, it is still far lower in value than the log marginal likelihood of the M1Cavity with a tightened uncertainty. To reach the level in which the model with the M1Cavity group would be disfavored relative to the other models, the uncertainties on these groups would have to change to unreasonable values. Therefore, even with this dependence on the prior uncertainty values, the single-component fits still favor the model with the M1Cavity group when explaining the data.

One additional observation is that loosening the prior uncertainties on particular M1  $^{232}\text{Th}$  groups for their single-component fits (like M1Connectors and M1CryostatCopperFar) actually have no effect on their log marginal likelihoods. This is because the activity fit by these groups is about one sigma or less from their prior activity values. Groups where the fitted activity is many sigma away from their prior activity values, like M1LMFEs and M1Cavity, are the groups where loosening (or tightening) the prior uncertainty has a noticeable effect on the log marginal likelihood.

Groups	M1Cavity	M1LMFes	M1Connectors	M1CCF	M1DUPTFE	M1Seals
Log ML	-8254	-28269	-16738	-16725	-16720	-16746
Log BF	n/a	20015	8485	8471	8466	8492
Groups	M1TSS	OuterCu	M1Bellows	M1CCPCables	M1CCN	M1DUStringCu
Log ML	-17063	-18775	-16732	-30917	-28506	-19846
Log BF	8809	10521	8478	22663	20252	11592
Groups	M1StringCables	PbShield				
Log ML	-23579	-29952				
Log BF	15325	21698				

Table 6.5: Log marginal likelihoods and log Bayes Factors of several component groups from single component fits where one  $^{232}\text{Th}$  group is allowed to float and all other remain fixed. Only M1 DS0-7 data is used in these fits. All log Bayes Factors are compared to the M1Cavity group.

Groups	M1Cav. (.2x unc.)	M1LMFes	M1Connectors	M1CCF	M1DUPTFE	M1Seals
Log ML	-10218	-28269	-16738	-16725	-16720	-16746
Log BF	n/a	18051	6521	6507	6502	6528
Groups	M1TSS	OuterCu	M1Bellows	M1CCPCables	M1CCN	M1DUStringCu
Log ML	-17063	-18775	-16732	-30917	-28506	-19846
Log BF	6846	8557	6514	20699	18288	9628
Groups	M1StringCables	PbShield				
Log ML	-23579	-29952				
Log BF	13362	19734				

Table 6.6: Log marginal likelihoods and log Bayes Factors of several component groups from single component fits where one  $^{232}\text{Th}$  group is allowed to float and all other remain fixed. Only M1 DS0-7 data is used in these fits. All log Bayes Factors are compared to the M1Cavity group.

Groups	M1Cav. (.2x unc)	M1LMFes (5x unc)
Log ML	-10218	-15356
Log BF	n/a	5138

Table 6.7: Log marginal likelihoods and log Bayes Factors of the M1Cavity and M1LMFes groups in a “worst-case scenario” configuration from single component fits where one  $^{232}\text{Th}$  group is allowed to float and all other remain fixed. Only M1 DS0-7 data is used in these fits. The M1LMFes log Bayes Factor is compared to the M1Cavity group.

## Section 6.6: Final Fit to DS0-7 Open and Blind Data and Estimate of the $2\nu\beta\beta$ Rate

A fit to DS0-7 open and blind data was carried out, floating all component groups with truncated Gaussian priors. While simulated dataset studies have shown that such fits are highly dependent on the prior uncertainties of the various component groups, the same studies have shown that the posterior of the  $2\nu\beta\beta$  spectrum remains stable due to its prominence in the model. Therefore, it should be possible to make an estimate of the  $2\nu\beta\beta$  rate in the DEMONSTRATOR using the results of this fit.

The  $^{232}\text{Th}$  integrated count results of this fit are shown in figure 6.11. The priors for the cavity groups were set at their upper limits from their most recent assay result, 0.57 mBq. As this assay result is an upper limit, their uncertainty is set at the same value. Therefore, it is not surprising that this fit to data places more posterior counts in these two groups than other groups, which will lead to an overfitting of the data as the M2Cavity group should have a much lower activity. That being said, the fitting algorithm recognizes from the data that more posterior counts should go into the M1Cavity group, which improves the fit overall relative to the comparison between the Assay-Based Background Model and data. The posterior activity of the M1Cavity group is 2.26 mBq and the posterior activity of the M2Cavity group is 0.62 mBq.

The energy spectrum of the posterior model and the data are shown in figure 6.12. As expected, based on the number of counts fitted into the posterior of the M2Cavity group, the model overfits the data in the  $^{208}\text{Tl}$  peak at 2615 keV. This also occurs in other  $^{232}\text{Th}$  decay chain peaks, such as the 2109 keV single-escape peak (see section 2.5.1) as well as the 919 keV  $^{228}\text{Ac}$  peak. Meanwhile, the 2204 keV peak of  $^{214}\text{Bi}$  from the  $^{238}\text{U}$  decay chain is slightly underfit. The 10 keV bin below the 1460.8 keV  $^{40}\text{K}$  peak is also underfit, but this is a residual effect of using static 10 keV wide bins. Overall, the fit is an improvement over the comparison between the Assay-Based Background Model and data (see figure 3.2).

The posterior activity for  $2\nu\beta\beta$  decay is 1.23 mBq for the M1EnrGe group and 0.94 mBq for the M2EnrGe group. This translates to specific activities of  $7.31 \times 10^{-5}$  and  $7.29 \times 10^{-5}$  Bq/kg respectively. Such measurements correspond well with the frequentist measurement of the  $2\nu\beta\beta$  decay specific activity,  $7.326 \times 10^{-5}$  across all enriched detectors [55].

The specific activities can be converted into a  $2\nu\beta\beta$  decay half life using the following expression:

$$T_{1/2}^{2\nu} = \frac{\ln(2)N_A}{m_A} \frac{\eta}{w_{2\nu}} \quad (6.13)$$

Where  $N_A$  is Avogadro's number,  $m_A$  is the molar mass of  $^{76}\text{Ge}$ ,  $\eta$  is the enrichment fraction of enriched detectors, and  $w_{2\nu}$  is the specific activity. This expression yields a  $2\nu\beta\beta$  decay half life of  $2.1 \times 10^{21}$  years.

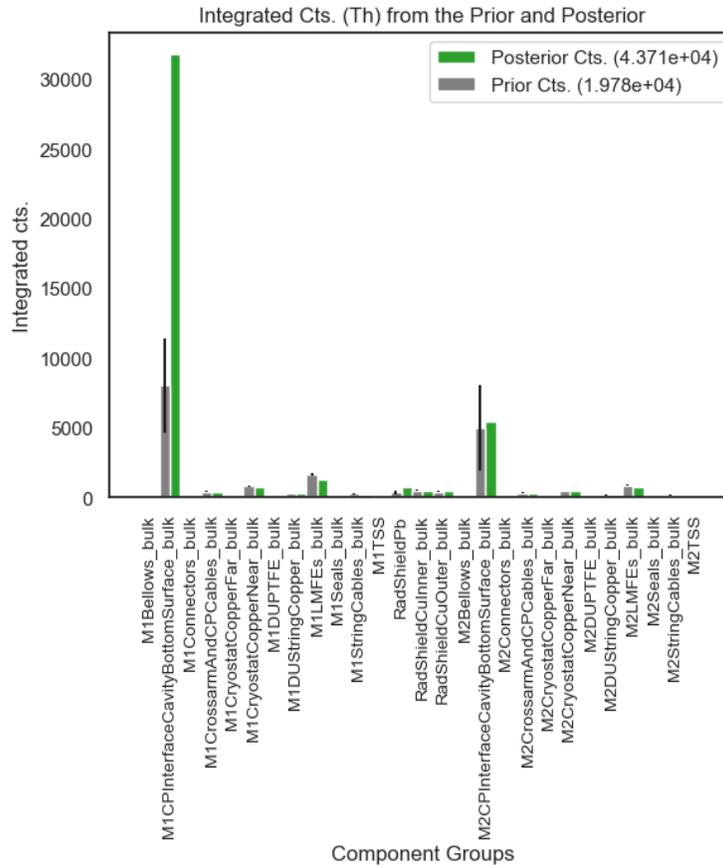


Figure 6.11: Fit to DS0-7 open and blind data, enriched and natural detectors. These are the integrated  $^{232}\text{Th}$  counts in all component groups, with the gray bars representing the number of counts as given by the prior, and the green bars representing the number of counts as given by the posterior.

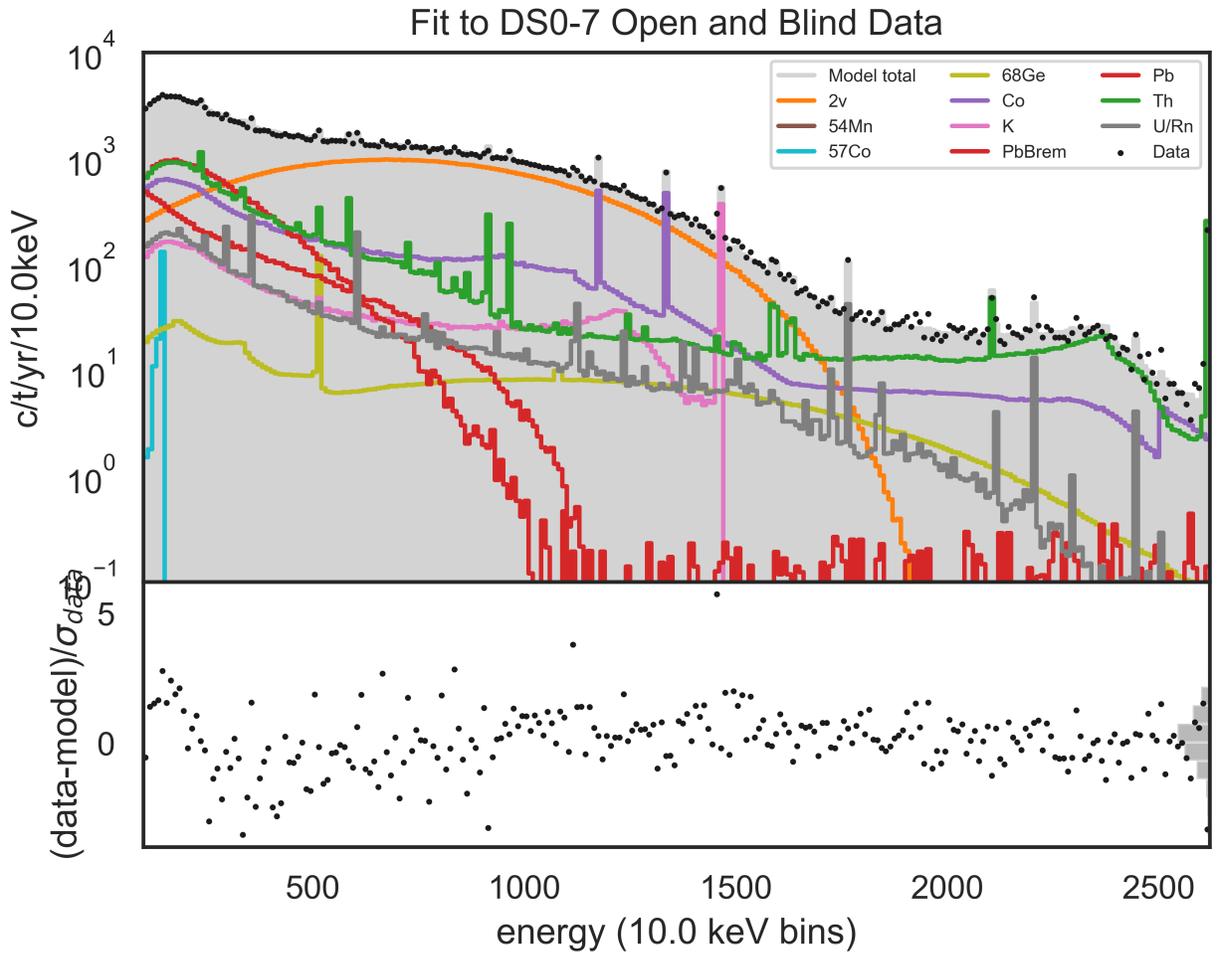


Figure 6.12: Fit to DS0-7 open and blind data, enriched and natural detectors.

## CHAPTER 7: Conclusions

A radioassay program was undertaken [41] prior to construction of the DEMONSTRATOR to measure the radioactivity of bulk components, specifically in the primordial decay chains of  $^{232}\text{Th}$ ,  $^{238}\text{U}$ , and  $^{40}\text{K}$ . These measurements culminated in the creation of the Assay-Based Background Model (section 4.3). Comparison of the Assay-Based Background Model to data showed a factor of five difference between the predicted rate and final rate of DS0-7 open and blind enriched backgrounds. Evidence in the energy spectrum from the data showed an excess in the 2615 keV  $^{208}\text{Tl}$  gamma peak as well as the 238 keV  $^{212}\text{Pb}$  gamma peak, both from the  $^{232}\text{Th}$  decay chain. This was initial evidence for a non-uniform  $^{232}\text{Th}$  contaminant in the DEMONSTRATOR.

Several supplementary simulation and data studies were carried out to determine the origin of this excess, as described in chapter 4. A combination of peak ratio and peak scaling studies were able to eliminate locations near the detector array and far from the detector array. A smattering of point source simulations centered around two high rate detectors, C1P7D1 and C1P2D1, were able to localize the likely source of the  $^{232}\text{Th}$  excess to the bottom surface of the Module 1 thermosyphon block cavity, referred to in this work as `M1ThermosyphonBlockCavityBottomSurface`, or shortened to `M1Cavity`. Peak ratios of the 238 and 2615 keV peaks in the high rate detectors as well as relative rate plots centered around the 238 keV peak eliminated all other possible sources. A parallel assay campaign was undertaken in an attempt to observe the contaminant in-situ, but a contaminant was not found. This included an assay of the Module 1 thermosyphon block cavity, which returned an upper limit of  $^{232}\text{Th}$  equal to 0.57 mBq, below the required activity of 1.78 mBq to explain the excess. Nevertheless, based on evidence from data and the aforementioned point source simulations, it is likely the contaminant was located in this component during operation of the experiment.

A Bayesian model of the DEMONSTRATOR was improved from previous work by regrouping component groups according to location and material in order to limit degeneracies between their corresponding PDFs. A PDF of the `M1Cavity` source (as well as its M2 equivalent) was added to the list of component groupings to be used in the model. A probabilistic programming framework was followed where prior information came from the results of the radioassay program, and a likelihood function was built for each individual energy bin based on PDFs from the simulations discussed in chapter 4. Equation 5.1 yielded a posterior from which a fitted model of DEMONSTRATOR backgrounds could be sampled. The Python package `PyMC3` facilitated this Hamiltonian Monte Carlo-like sampling with the No U-Turn sampler.

The fitting results were discussed in chapter 6. A number of fits to simulated datasets were carried out to assess the ability of the Bayesian fitting algorithm to place counts in the correct component group when an artificial  $^{232}\text{Th}$  excess is placed in that component group. It was observed that the Bayesian fitting algorithm was able to place the correct number of counts in component groups when the dataset was sufficiently large - approximately 1000 times larger than the size of the DS0-7 open and blind dataset. However, the Bayesian fitting algorithm struggled to fit the correct number of counts when the dataset had roughly the same number of counts as the DS0-7 open and blind dataset. Specifically, the algorithm tended to fit additional counts into groups with larger uncertainties on their prior. Thus, the Bayesian probabilistic programming framework is limited in its power to infer the activities of various components on the DEMONSTRATOR due to the low-background nature of the experiment.

By reducing dimensionality of the model to a single parameter, the fitting algorithm could be used to test the hypothesis that a model including the `M1Cavity` PDF will describe the data with a higher probability than a model without the PDF. Using the concept of a Bayes Factor as a goodness of fit metric and the methodology described in section 6.5.1, a series of single-component fits were carried out to test this hypothesis. Though it was discovered that Bayes Factors were influenced by the relative size of the `M1Cavity` prior uncertainty to other prior uncertainties, the value of the log marginal likelihood of the model with the floating `M1Cavity` group proved to be higher than all other models, given reasonable prior uncertainty values for all component groups. This observation further favors the presence of a  $^{232}\text{Th}$  contaminant on the bottom surface of the Module 1 thermosyphon block cavity during operation of the experiment.

A final fit to DS0-7 open and blind data was performed to obtain a measurement of the  $2\nu\beta\beta$  half-life. Though the results of this fit ultimately overestimated the amount of  $^{232}\text{Th}$  in component groups relative to energy spectrum from data, the prominence of the  $2\nu\beta\beta$  spectrum in the data meant there was little variation in the number of counts attributed to  $2\nu\beta\beta$  decay in this fit relative to all previous fits (to data or simulated datasets). The  $2\nu\beta\beta$  half-life as observed by the DEMONSTRATOR and measured by this work is  $2.1 \times 10^{21}$  years, which corresponds well with the frequentist measurement of  $2.085 \times 10^{21}$  years [55].

## Section 7.1: Future Studies

For continued use of the python-based software project that facilitates the Monte-Carlo sampling of the posterior to make model inferences for the DEMONSTRATOR, there are a number of updates that are required. The current project uses PyMC3 version 3.6 and `theano` version 1.0.3. Both of these packages are outdated, and require older versions of MacOS in order to run. Additionally, the `theano` package has been shown to be somewhat unreliable and prone to memory leak issues. Fortunately, the most recent version of PyMC

as of the release of this work, version 5.9.0, uses the `PyTensor` package which like `theano` can manipulate multi-dimensional arrays, but with greater support. The software project of this work would have to undergo a number of syntax updates to match the currently available set of functions for both `PyMC` and `PyTensor`, but the overall framework of the project should not have to undergo any major revisions. The official website of the `PyMC` package<sup>1</sup> has the documentation necessary to facilitate these updates.

Though the size of the DEMONSTRATOR’s low-background dataset ultimately limits the power of the Bayesian probabilistic programming framework from making accurate radioactivity inferences, a number of supplementary studies could be undertaken to learn more about the individual backgrounds observed by the experiment. In particular, a moderate excess of  $^{238}\text{U}$  has observed in the data of both modules, as shown in table 7.1.

Peak Energy	Decaying Isotope	% Assay Cts. v. Data (All)	% Assay Cts. v. Data (M1)	% Assay Cts. v. Data (M2)	% Assay Cts. v. Data (Enr)	% Assay Cts. v. Data (Nat)
351.9 keV	$^{214}\text{Pb}$	46	39	62	43	52
609.3 keV	$^{214}\text{Bi}$	45	42	50	46	43
1120.3 keV	$^{214}\text{Bi}$	36	35	36	32	47
1764.5 keV	$^{214}\text{Bi}$	45	41	53	52	35
2204.1 keV	$^{214}\text{Bi}$	35	34	38	33	43

Table 7.1: Table of prominent  $^{238}\text{U}$  chain peak counts in the assay-based model versus DS3-6 data. The detectors are divided into various groupings, including M1 only, M2 only, Enr only, and Nat only.

An ideal future study would carry out a number of single-component fits with  $^{238}\text{U}$  component groups and calculate the individual log marginal likelihoods for each model to observe if one group has a higher likelihood than other groups to explain this excess. Based on a study of relative rates in individual detectors around prominent  $^{238}\text{U}$  gamma peaks conducted by A. Reine [55], the hypothesis to test here is that the excess is uniform and not contributed by a single component.

The conclusions presented in this work have useful implications for the LEGEND-1000 experiment. The supplementary studies discussed in chapter 4 determined that the source of the  $^{232}\text{Th}$  excess could not be from near-detector components. Therefore, a lot of the same near-detector technology used on the DEMONSTRATOR and currently installed in LEGEND-200 could be reused for LEGEND-1000. Additionally, the fits to high-statistics simulated datasets discussed in section 6.4 illustrated that the Bayesian inference methods discussed in chapters 5 and 6 could work for high-statistics datasets. With a much larger exposure,

<sup>1</sup>Follow the guide at this hyperlinked url.

LEGEND-1000 may provide a dataset that will be more compatible with the methods explored in this work, and thus a Bayesian approach to constructing the background model for LEGEND-1000 maybe be a viable strategy.

## APPENDIX A: A CABLE AND CONNECTOR UPGRADE FOR THE MAJORANA DEMONSTRATOR

In November 2019, a cable and connector upgrade commenced on Module 2 of the MAJORANA DEMONSTRATOR. The collaboration implemented new materials and techniques to improve the reliability and robustness of cables and connectors in the DEMONSTRATOR. I discuss the motivation and implementation of this upgrade below.

### Section A.1: Motivation

Prior to the hardware upgrade, only 40 of 58 detectors were operational in the DEMONSTRATOR. Eight of these detectors had issues either with Vespel-made connectors that link signal cables together, or with damaged Low Mass Front End boards (LMFEs). Nine detectors remained unbiased due to problems with their corresponding high voltage (HV) cables [56]. These 17 detectors<sup>1</sup> had potential to be brought online through an upgrade, where all new high voltage and signal cables would be installed with connection and handling improvements.

### Section A.2: Overview of Cables

#### A.2.1: Signal Cables

For assembly reasons, signal cables installed in the DEMONSTRATOR consist of two cables that are connected together above a copper plate in the detector cryostat, referred to as the “coldplate” (see Figure A.1). Each detector requires four cables to read out signals. A cable company, Axon’, originally fabricated these cables in 2013 to feature a coaxial design with a total diameter of 0.4mm (see Table A.1).

The cable was acceptable to use for the upgrade, as the original design satisfied the radiopurity and performance requirements established by the collaboration. Axon’ delivered a new production batch of cable for the upgrade in 2019. These cables were pre-terminated with nano twist-pin connectors.

#### A.2.2: HV Cables

HV cables are run as one unbroken length, and each detector requires exactly one HV cable. Axon’ originally fabricated these cables in 2013 to feature a coaxial design with a total diameter of 1.2mm. They

---

<sup>1</sup>The 18<sup>th</sup> inoperative detector with serial ID P42909C remained unbiased due to problems with the crystal itself and was not made operational through the hardware upgrade. Initially it was ambiguous if this detector had problems with its crystal or cables and connectors, but the upgrade confirmed the crystal issue. The detector was removed in July 2020.

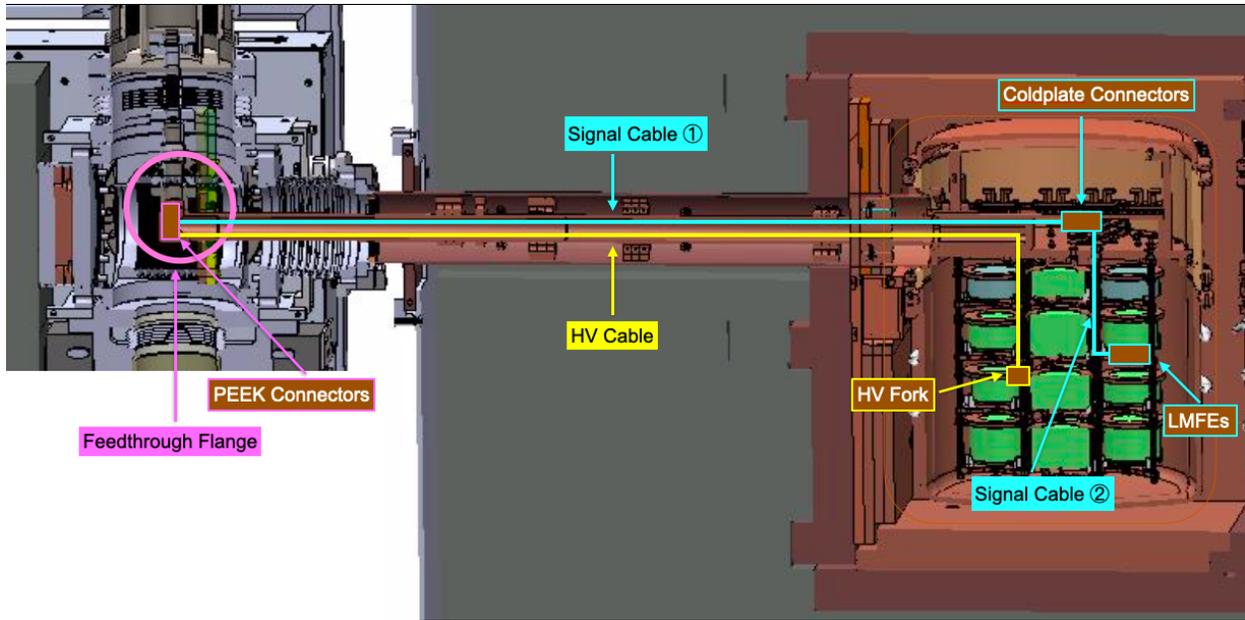


Figure A.1: Schematic of an MJD module, highlighting the locations of cables and connectors.

Axon' Signal Cable Properties	
Property	Value
Diameter	0.4 mm
Characteristic Impedance	50 $\Omega$
Capacitance per unit length	87 pF/m
Mass per unit length	0.4 g/m
Conductor Material	OFHC Cu
Conductor Gauge	AWG 40
Conductor Diameter	0.076 mm
Dielectric Material	FEP
Dielectric Diameter	0.254 mm
Shield Material	OFHC Cu
Shield Gauge	AWG 50
Number of Shield Strands	approximately 30
Jacket Material	FEP

Table A.1: Table of Axon' signal cable properties.

are rated up to 5kV DC (see Table A.2). The collaboration carried out R&D measurements on the original batch of cables to characterize them and determine their optimal configuration in the DEMONSTRATOR [57]. The study also analyzed Module 1 commissioning data (referred to as Dataset 0 or “DS0”) and found that the micro discharge rate of cables remained below the upper limits determined from characterization

Axon' HV Cable Properties	
Property	Value
Diameter	1.2 mm
Maximum Voltage Rating	5 kV DC
Mass per unit length	3 g/m
Conductor Material	OFHC Cu
Conductor Gauge	AWG 34
Conductor Diameter	0.152 mm
Dielectric Material	FEP
Dielectric Diameter	0.618 mm
Shield Material	OFHC Cu
Shield Gauge	AWG 50
Jacket Material	FEP

Table A.2: Table of Axon' HV cable properties.

measurements.

Prior to the upgrade, several detectors suffered from occasional high voltage breakdowns. This led to several interruptions in data taking, as breakdowns automatically trigger a ramp-down of voltage to all detectors in the same module as the detector that suffered the breakdown. The collaboration fixed this issue by disconnecting the high voltage cables of several problematic detectors from ground. A collaborator, David Radford, analyzed data taken after this disconnection to measure the impact of un-grounded detectors on the energy resolution of events at low energies. Radford found that detectors that were un-grounded in Module 1 had slightly worse energy resolution than those that remained grounded (FWHM = 317 eV un-grounded, FWHM = 257 eV grounded), though the collaboration agreed that this worsening of energy resolution at low energies was worth it for the reduction in breakdowns. The energy resolution at high energies, specifically measured at the 2615 keV  $^{208}\text{Tl}$  full escape peak, was negligibly effected.

The disconnection of particular HV cables from ground had the desired effect of reducing the number of breakdowns, but it did not address the source of the breakdowns themselves. To investigate the source of the breakdowns, collaborators placed HV cable samples between one of three sets of adjustable clamps - one of two flat surfaces, one of a flat surface and rounded surface, and one of a flat surface and sharp surface. Collaborators operated the cable for five minutes at its maximum voltage rating of 5 kV with a particular amount of pressure from the clamps, measured in units of clamping torque. The clamping torque increased with every trial until the cable exhibited a full breakdown. The test results are shown in Table A.3. While cables withstood significant clamping torque from the flat and rounded surfaces, a fraction of that torque

from the sharp surface was enough to induce a breakdown. These results suggested that breakdowns in the original set of HV cables were possibly a result of kinks or deformed sections of cable, which were damaged when pressed against sharp edges. It is likely that the sharp edges of the thermosyphon baffle plates kinked or pressured the cables during their installation along the thermosyphon.

HV Cable Clamping Test Results	
Clamp Type	Clamping Torque at which Breakdown was Triggered
Flat Surface & Flat Surface	Above 12 inch-lbs
Flat Surface & Round Surface	6-10 inch-lbs
Flat Surface & Sharp Surface	2 inch-lbs

Table A.3: Table of HV cable clamping test results.

Like the signal cables, the HV cable was acceptable for use in the upgrade, as it satisfied the radiopurity and HV performance requirements outlined by the collaboration<sup>2</sup>. The upgrade used the same production batch of cable from 2013. In order to avoid kinking cables prior to installation, the assembly team handled the cables as little as possible during the fabrication process and carefully wound them at a uniform radius when stored to keep tight loops from forming. The team bundled the cables together according to NASA specifications [58] prior to quality testing (see section A.4.1). At installation, the team secured these cables in the crossarm with newly designed baffle plates that feature rounded edges. These plates are discussed in section A.4.3.

### Section A.3: Overview of Connectors

#### A.3.1: Nano Twist-pin Connectors

For assembly reasons, signal cables installed in the DEMONSTRATOR consist of two cables that are connected together above the coldplate (see Figure A.1). One cable travels the length of the hollow copper crossarm, which serves as a channel through the lead shielding for cable passage. For this cable, one end terminates at a stainless steel flange in the high vacuum system (referred to as the “feedthrough flange”) and the other terminates above the coldplate. The second cable travels the length of the detector strings, with one end terminating above the coldplate, and the other terminating at the LMFE of a detector. There are four signal cables for every detector, as necessitated by the LMFEs that read out the charge collected at the point contact. The LMFEs with terminated cables are mounted to all detectors of the same string

<sup>2</sup>There was an investigation into the use of alternative HV cables, but ultimately collaborators chose the original design due to the time and cost restraints associated with these alternative cables.

individually before assembly of that string.

A custom Vespel connector, fabricated in-house, initially joined together the two ends of each cable that terminated above the coldplate. The collaboration optimized the design of this connector to minimize radioactivity and mass. Unfortunately, collaborators discovered during use of the connectors that this design was not robust enough to withstand repeated temperature cycling. The connector avoided using conventional spring components. These spring components, made from CuBe, would have resulted in unacceptably high  $^{232}\text{Th}$  and  $^{238}\text{U}$  activities. Without conventional spring components, these connectors suffered from intermittent and permanent connectivity issues, resulting in the inoperability of several detectors.

The collaboration’s project engineer, Matthew Busch, worked with several vendors to investigate alternative connector designs. The selected solution was a low-mass connector designed by Axon’ [59]. These “nano twist-pin” connectors (Figure A.2) feature gold alloy contacts encapsulated by an Ultem shell. Such connectors surpass the radiopurity of the previous Vespel connectors (see Table A.4) and passed performance tests at Axon’s production site in France as well as at the cable fabrication site at the University of North Carolina at Chapel Hill (UNC). Axon’ pre-terminated these connectors to signal cables before sending them to the collaboration.

Axon’ Nano Twist-Pin Connector Assay Results and Comparison		
Decay Chain	$^{238}\text{U}$ (mBq/kg)	$^{232}\text{Th}$ (mBq/kg)
Axon’ NTP Connector Full Body Assay	$0.692 \pm 0.088$	$0.257 \pm 0.032$
Vespel Connector Full Body Assay [41]	$1.16 \pm 0.017$	$0.365 \pm 0.006$

Table A.4: Table of Axon’ Nano Twist-Pin Connector assay results, compared to the previously used Vespel connectors. Assays performed at Pacific Northwest National Laboratory (PNNL).

### A.3.2: Signal Feedthrough Flange Connectors

The cable termination at the feedthrough flange was a custom-built 50 pin D-sub made from PEEK material to meet radiopurity requirements (see figure A.4a). These connectors did not suffer from continuity issues, but there was evidence of signal cable damage due to connection difficulties during installation. The connectors featured incorrect pin-spacing due to a machining error, which made the D-sub difficult to plug and unplug into the feedthrough flange.

Collaborators considered commercial 50 pin D-sub for the upgrade, but ultimately ruled them out due to radiopurity concerns. The upgrade used the original PEEK design with correctly machined pin-spacings to solve previous installation difficulties (see figures A.4b and A.8). Improved stripping of signal cables as well as new socket crimping techniques also mitigated the risk of damaging the signal cables during attachment

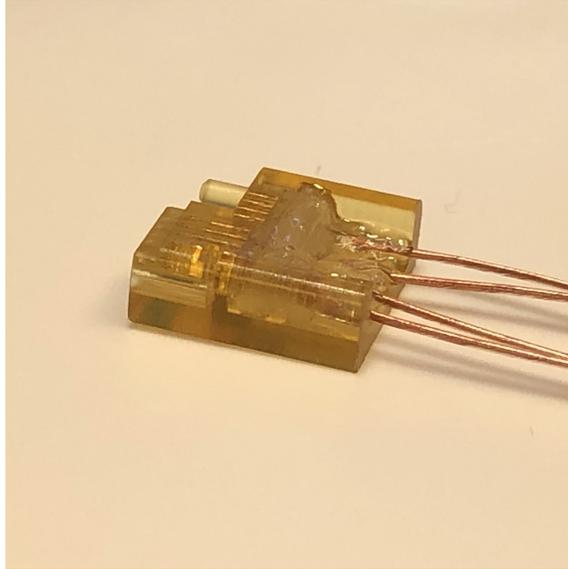


Figure A.2: One half of an Axon' Nano twist-pin connector. Axon' delivered these connectors pre-terminated to signal cables.

to the connector.

### A.3.3: HV Feedthrough Flange Connectors

HV cables run an unbroken length from the feedthrough flange to the detectors. Like the signal cables, the HV cables are terminated at the feedthrough flange via custom low-mass PEEK connectors. Originally, these connectors consisted of two sockets that collaborators terminated to the cable ground shield and central conductor with silver epoxy. A PEEK two-body shell housed the epoxied connection. Unfortunately, the sockets used were much smaller than what the design requested, and therefore the connectors could not effectively mate with the flange sockets and had a tendency to fall out of them when bumped into. The sockets were also not completely covered by the PEEK housing, allowing for potential bending of the sockets or catching the sockets on other cables during installation. Additionally, the small amount of inlet space and the central position of this space at the top of the connector, designed for cable passage, impacted strain relief (see Figure A.3a). These issues resulted in connectivity problems.

The upgrade used a revised design that extended the PEEK two-body shell to encapsulate both sockets and secure them in properly sized bore holes. This protected the sockets from bending and improved their alignment and holding force when inserted into the flange sockets. Collaborators also searched for a suitable replacement to the shorter sockets and found a product from Glenair [60] that provided a stronger and more consistent clamping force. The new design added a larger inlet space at the top of the PEEK case to address the issue of strain relief, and called for epoxying the cable in place during fabrication. Finally, the new design

added a notch to the top side of the case in order to easily identify which socket on the connector is crimped to the ground shield. See Figures A.3a, A.3b, A.4a, and A.4b for a comparison of both designs.

#### A.3.4: HV Fork Connectors

The HV cables are terminated at the detector by way of a copper “HV Fork”. This custom connector is clamped to a copper plate, called the “HV Ring”, that makes contact with the n+ surface of the detector (see Figures A.6a and A.6b). The original design terminated the HV cable at the HV Fork with a custom plug made of Vespel that pinned the exposed central conductor to a bored hole through the HV Fork’s surface (see Figure A.5a). During usage, collaborators found that this Vespel plug was not always secure in holding the exposed HV cable central conductor to the HV Fork, leading to connectivity issues.

The upgrade utilized a new HV Fork design that featured a crimped connection, as shown in Figure A.5b. The crimped connection of the central conductor directly to the HV Fork simplified assembly and resulted in a robust connection.

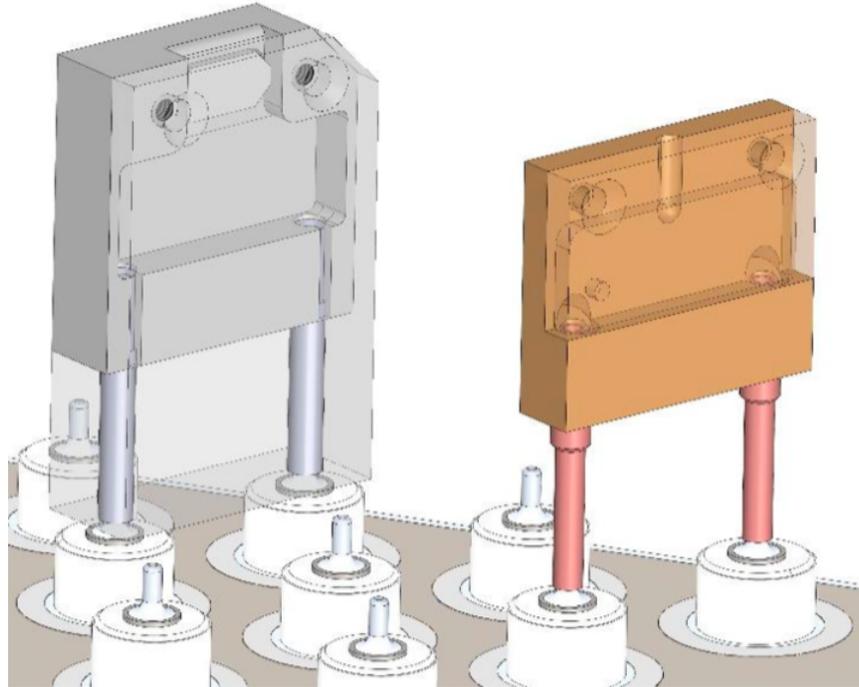
### Section A.4: Assembly and Fabrication at UNC

The assembly team at UNC assembled and fabricated cables and connectors for the upgrade in a class 1000 soft-walled cleanroom. Axon’ produced and shipped the signal cables and nano twist-pin connectors to UNC. Collaborators at SURF sent the 2013 production batch of HV cable from their storage space on-site to UNC. Professional machinists at UNC machined all of the new PEEK connectors. Collaborators purchased the sockets used for the feedthrough connectors from a commercial company, Glenair [60]. And finally, a professional machinist at SURF machined the HV Fork connectors out of electroformed copper and shipped them to UNC.

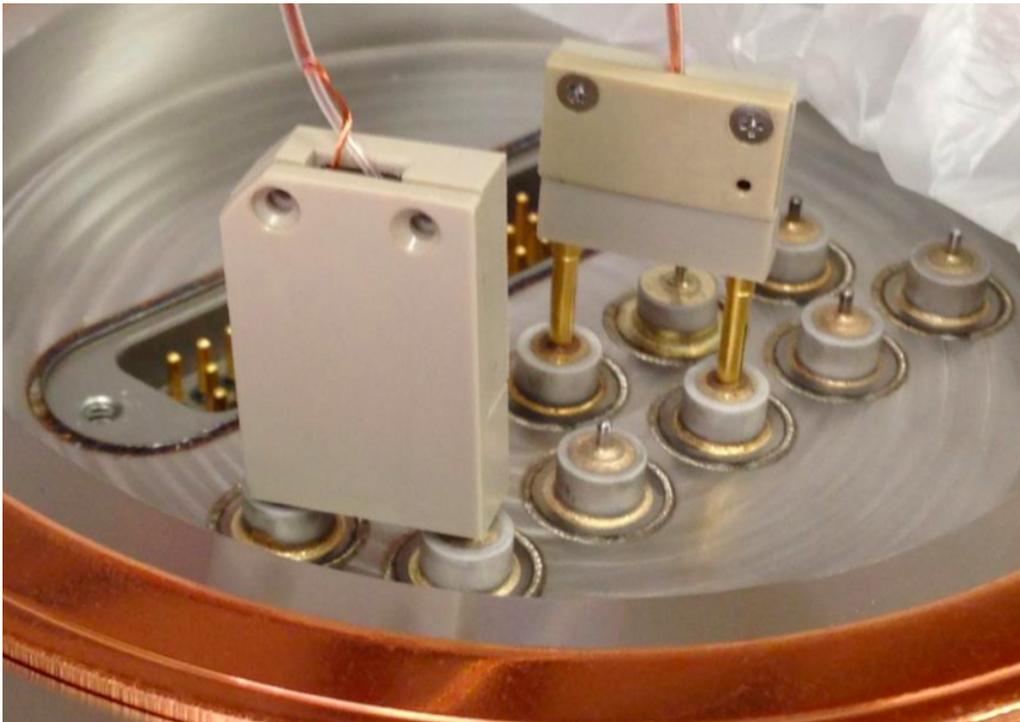
#### A.4.1: Fabrication of Signal Cable Bundles

The signal cable arrived from Axon’ pre-terminated to the nano twist-pin connectors. Therefore, the remaining terminations to fabricate were at the feedthrough connector and LMFE. At both terminations, the assembly team carefully and precisely stripped the cable with a Schleuniger [61] CoaxStrip 5300 RX machine (see Figure A.7a) by way of a programmable rotary cutter. This machine vastly improved the quality of the stripped cables compared to the hand-stripped cables used in the initial installation (see Figure A.7b).

At the feedthrough connector, the assembly team individually crimped the central conductor and ground of the signal cable to the Glenair sockets, and then carefully installed them into the two-body PEEK 50

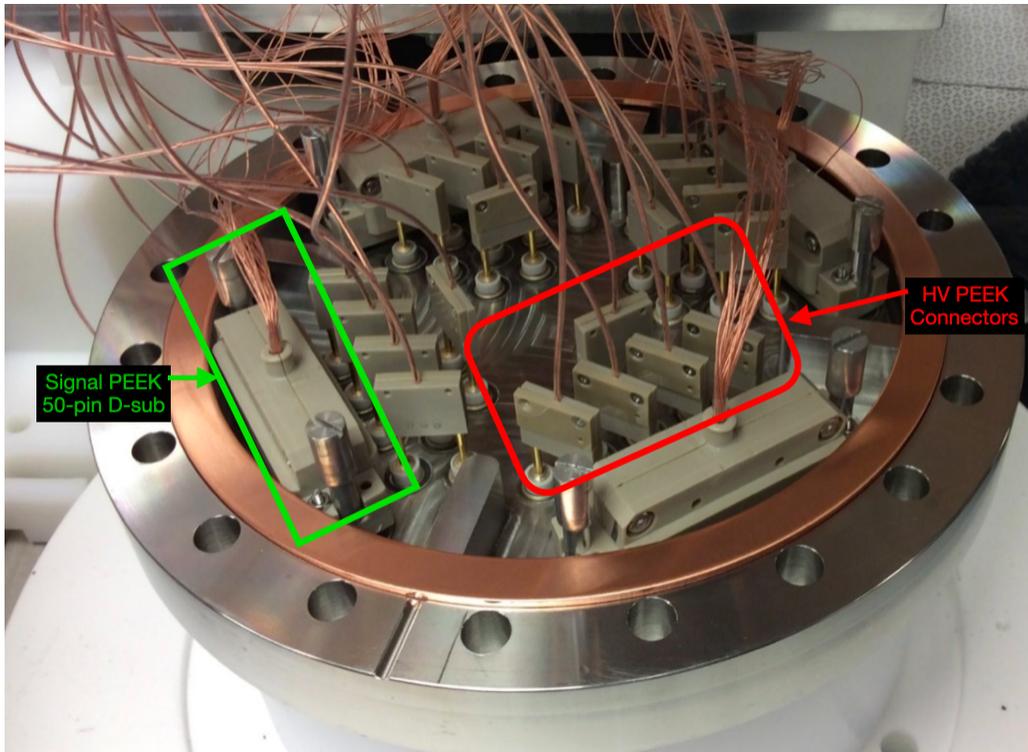


(a) Schematic of the new (left) and old (right) HV Feedthrough Connector designs. The new design features an extended case that covers the crimped sockets (translucent gray) and features a notch on the top that indicates which socket is connected to ground.

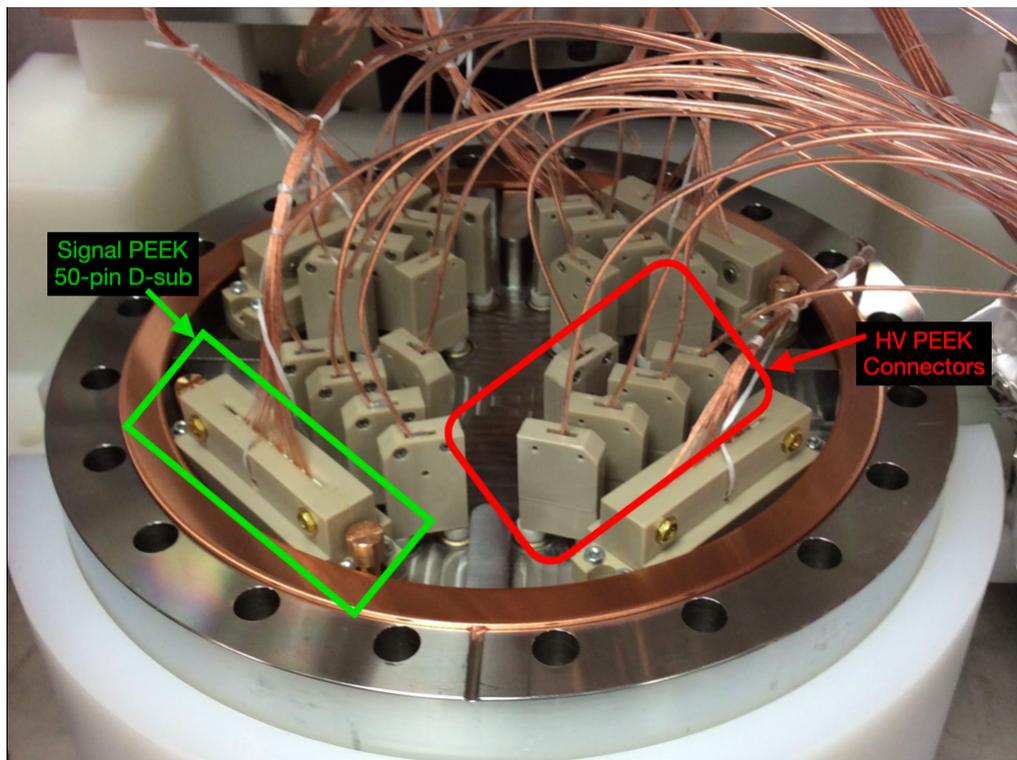


(b) Physical comparison of the new (left) and old (right) HV Feedthrough Connector designs. Note that the new connector in this image is flipped relative to the new connector in the design schematic, as noted by the location of the notch on the top.

Figure A.3: Comparison of old and new HV Feedthrough Connector designs.

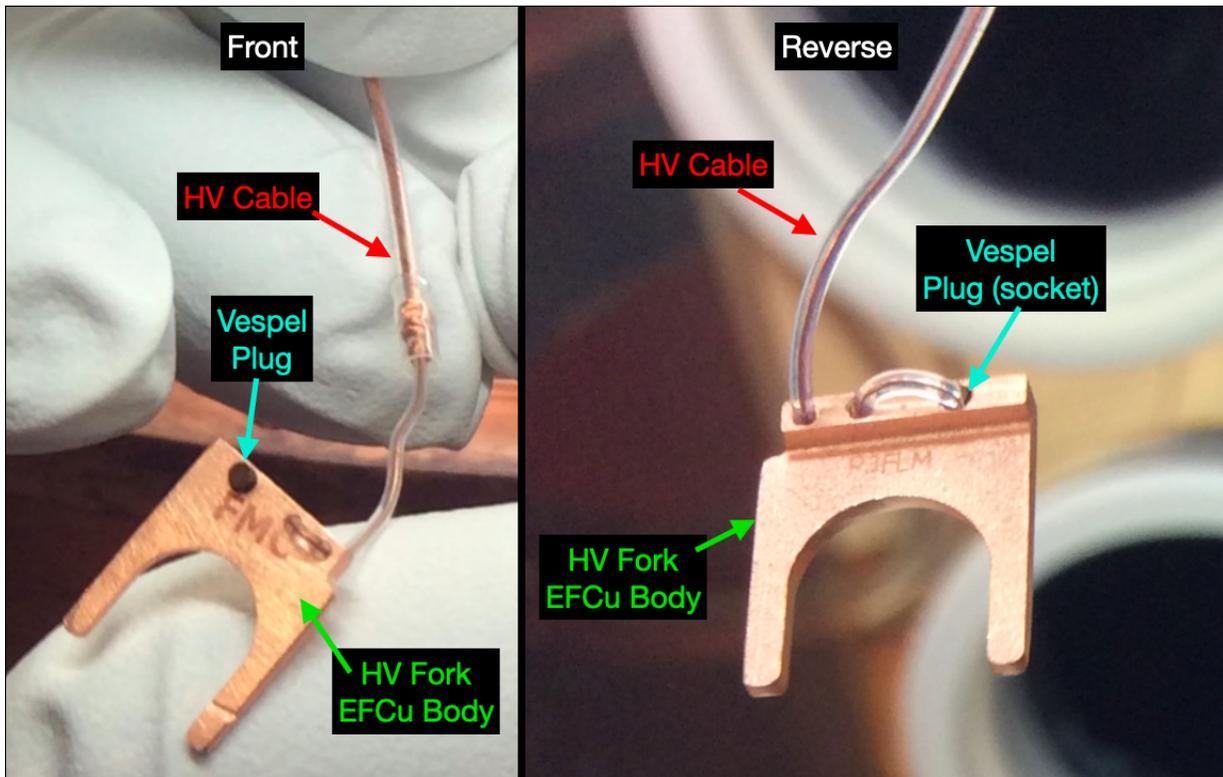


(a) The old setup at the feedthrough flange, made up of HV and signal cable connections through custom built PEEK connectors.

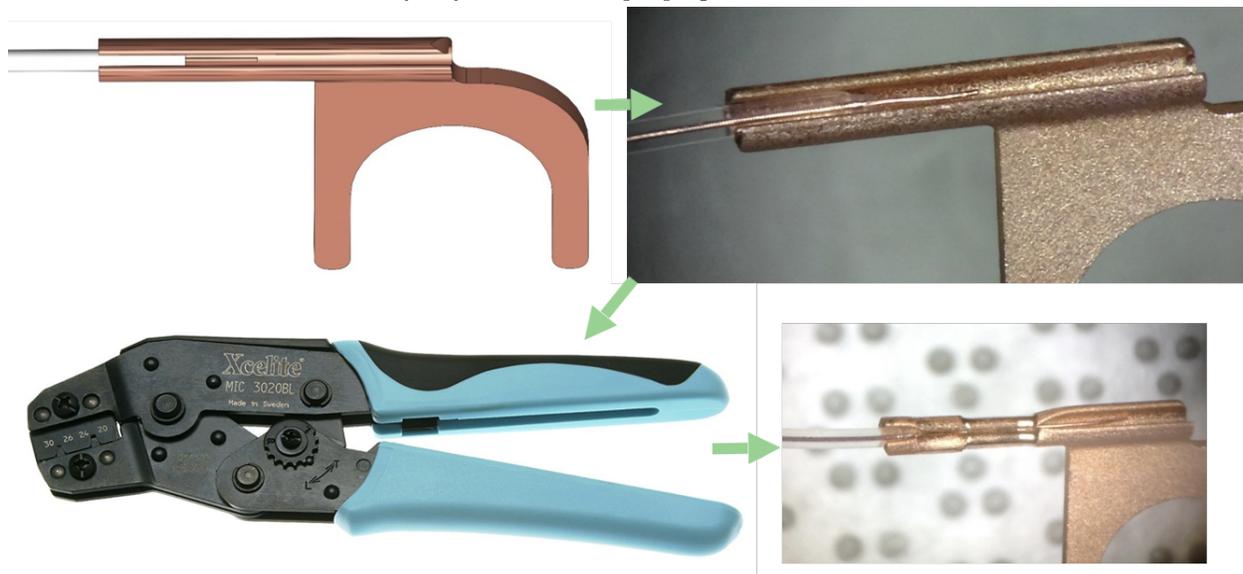


(b) The new setup at the feedthrough flange, made up of HV and signal cable connections through revised, custom built PEEK connectors.

Figure A.4: Cable and connector setup at the feedthrough flange, before and after the upgrade. Notice the improved cable bundling with plain dental floss in the background.

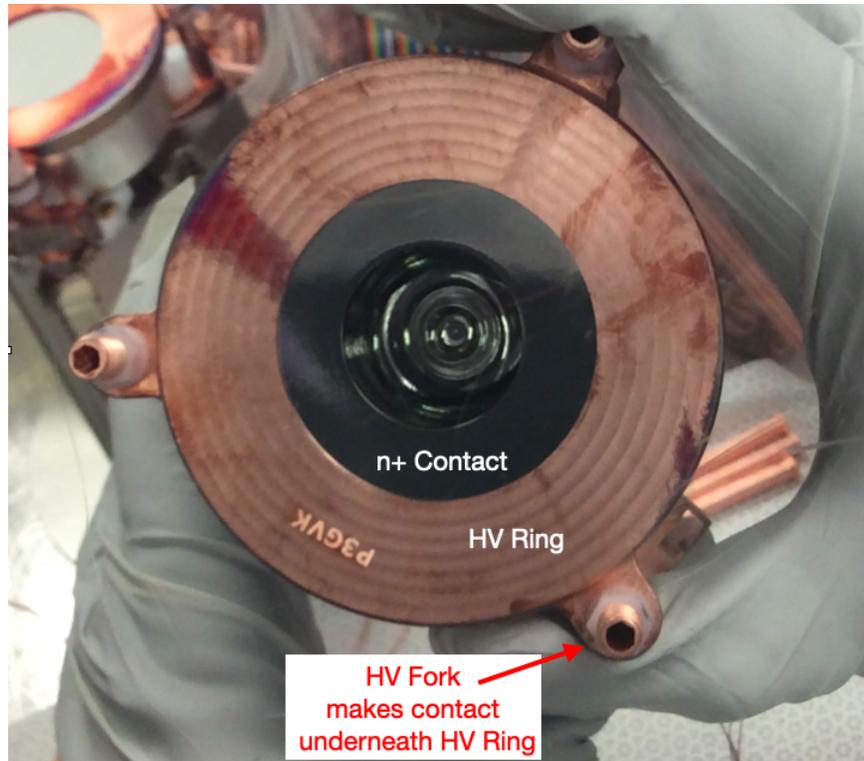


(a) The old design of the HV fork, where the HV cable is threaded through the body of the fork and the central conductor is terminated to the fork by way of a small Vespel plug.

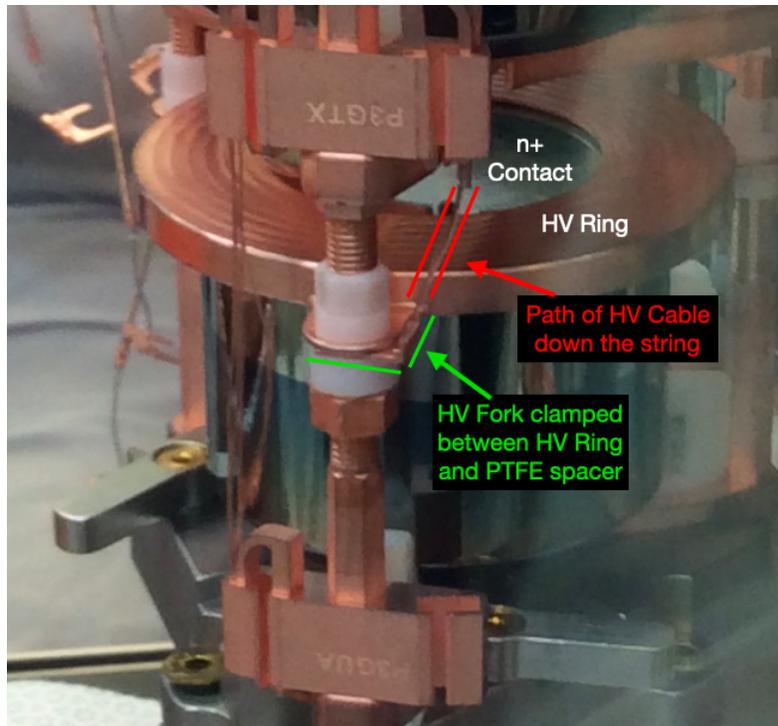


(b) The new design of the HV fork, where the central conductor of the HV cable is crimped to the fork through a copper receptacle. The assembly team performed the crimping using a specially cleaned MIC3020BL Mini-Crimper manufactured by Xcelite.

Figure A.5: Comparison of old and new HV Fork designs.



(a) Top-down view of the HV Ring on the n+ contact of the detector. The HV fork is clamped just below one of the three HV Ring nubs for electrical contact. Note that the crystal in this image is an ICPC detector, but the HV Ring design is the same for a PPC detector.



(b) Diagram showing the location at which the HV Fork is clamped between a PTFE spacer and a copper nub of the HV Ring. The HV cable that runs down the string to the detector is outlined in red. The edges of the HV Fork are outlined in green.

Figure A.6: Overview of HV Fork and HV Ring placement.



(a) The Schleuniger CoaxStrip 5300 RX machine used to strip cables for the upgrade.



(b) Example of the high quality strip-back obtained using the CoaxStrip Machine

Figure A.7: The cable stripping machine used for the upgrade and a stripped cable.

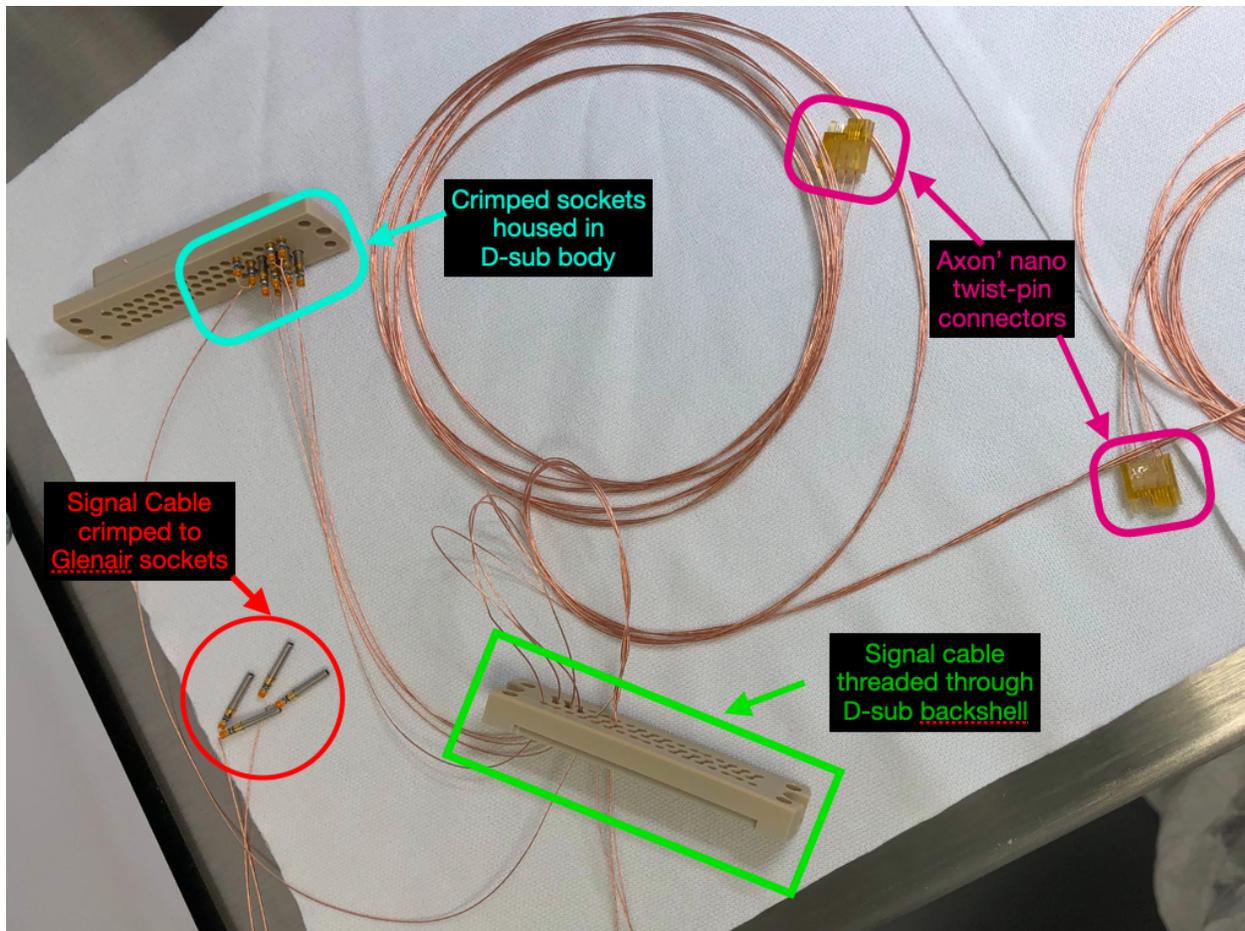


Figure A.8: Fabrication of a signal cable bundle. The ground and central conductor of the signal cable is passed through the back shell of the 50-pin D-sub before being crimped into sockets. Those sockets are then loaded into the main body of the D-sub.

pin D-sub (see Figure A.8). At the LMFE, a team of experts at Lawrence Berkeley National Lab (LBNL) delicately adhered the signal cable to the terminals of the LMFE using silver epoxy (see Figure A.9). This team was also responsible for the production of new LMFEs.

Regarding the crossarm section of signal cable, the assembly team braided together the four wires per nano twist-pin connector by hand, using a braid pattern that repeats every 2 cycles. The team then bundled together several braided groups with plain dental floss according to NASA specifications [58]. The team produced two bundles, with each bundle containing four 50-pin D-sub with 20 signal cables each, for a total of 80 signal cables per bundle. The team tested these bundles for continuity and temperature cycling before shipping them to SURF for installation.

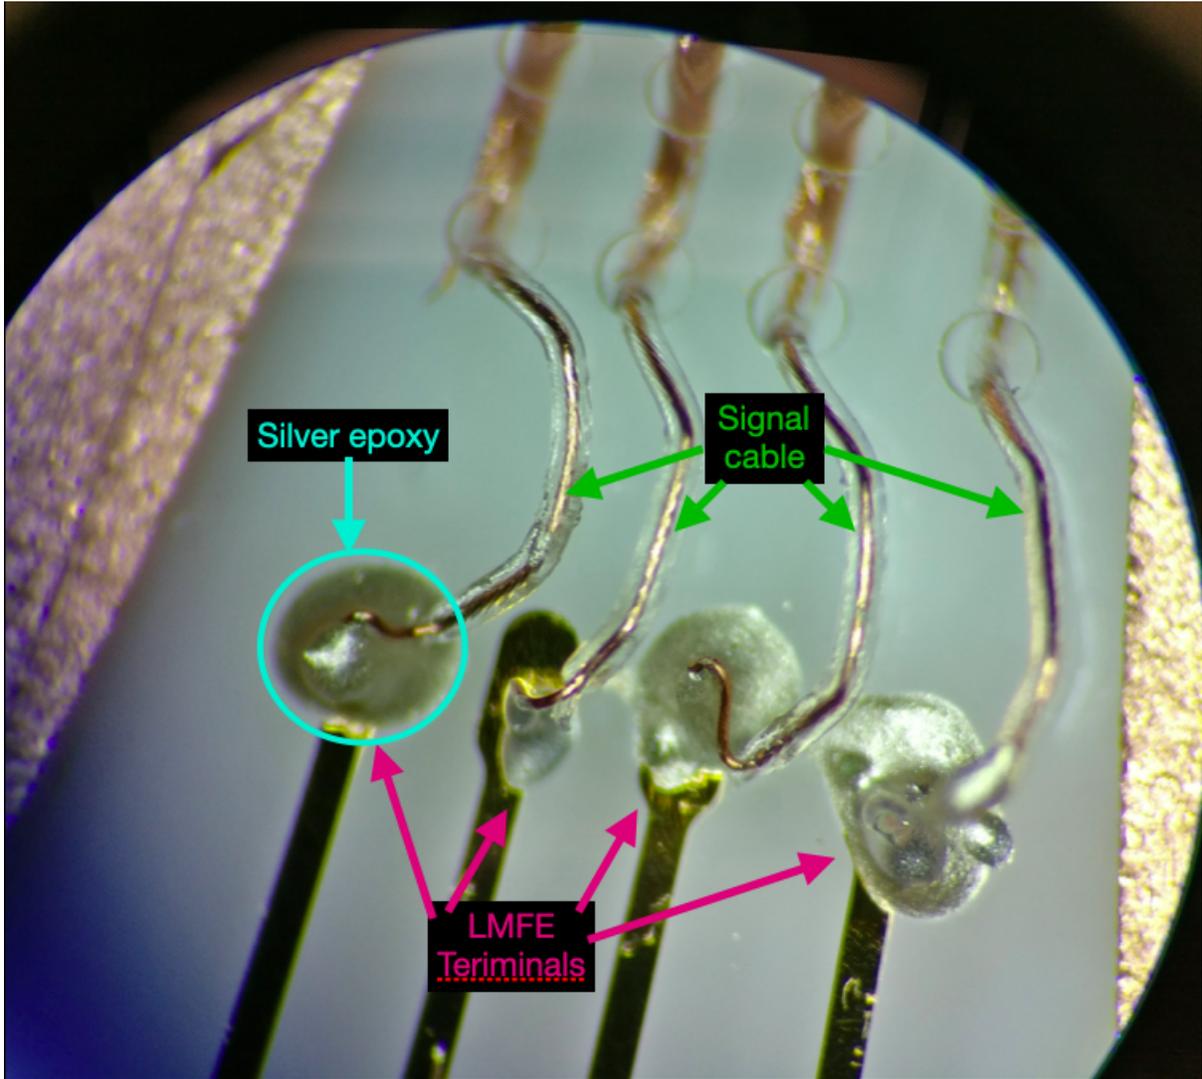


Figure A.9: Termination of signal cable to LMFE terminals. Collaborators perform this work under a microscope with silver epoxy at LBNL.

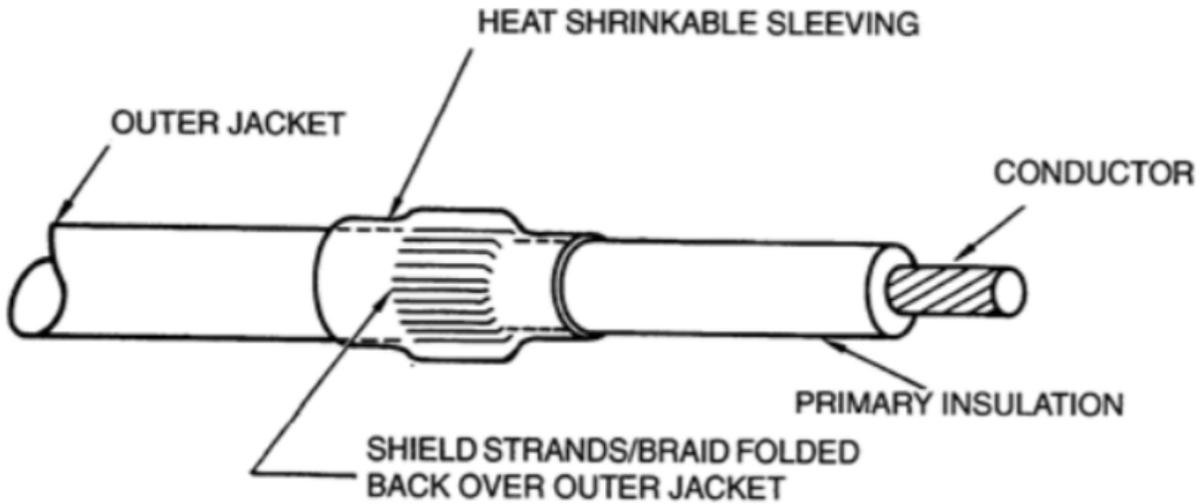


Figure A.10: Configuration of the heat shrink plastic relative to the exposed ground shield wires, according to NASA specifications. [58]

#### A.4.2: Fabrication of HV Cable Bundles

The assembly team sourced the HV cable directly from the 2013 production spool and cut each cable to an appropriate length of about 2m. At both terminations, the team stripped the cable with the Schleuniger CoaxStrip 5300 RX machine. At the feedthrough connector, the team individually crimped the central conductor and ground of the HV cable to Glenair sockets. They epoxied the exposed wire between the crimped location and the length of the stripped outer jacket inside the two-body PEEK case using LOCTITE [62] EA 0151 epoxy to provide strain relief. At the HV fork, the team cut back the shield of the HV cable to an exposed length of 1-2 mm and stripped the insulation layer to reveal 3mm of exposed central conductor. They sealed the exposed shield wires with heat shrink plastic to isolate the central conductor from ground (see Figure A.10) [58]. The team crimped the central conductor into the copper receptacle on the re-designed HV Fork. They then bundled together HV cables with plain dental floss using the same technique as the signal cables. The team produced two bundles, the first containing 23 cables and the second containing 25 cables. They tested these bundles for continuity and temperature cycling with liquid nitrogen. The bundles also underwent a new micro-discharge test to measure reliability and risk for breakdowns. Cables that failed any of these tests had their connectors cut, but remained in the bundle to avoid re-bundling. This left the first bundle and second bundle with 20 operational cables each. The team finally shipped the cables to SURF for installation.

#### A.4.3: Baffle Plate Fabrication

In the original installation, collaborators ran the HV and signal cable along the surface of the thermosyphon within the crossarm. To keep the cables secure and to provide additional shielding within the crossarm, the design used copper baffle plates. Collaborators found that the original design of these baffle plates created a couple of issues. First, the plates had sharp-edged holes of very tight spacing for the passage of cables, leading to an increased probability of cable damage during installation. Second, the plates were not oriented optimally and did not maximize the amount of radiation shielding the region could contain.

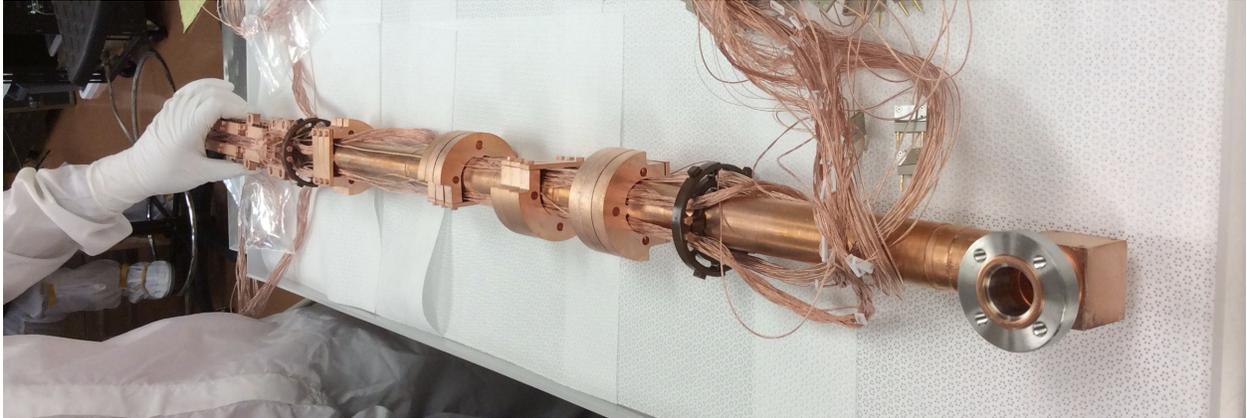
The upgrade used a new baffle plate design that addressed these issues. The design instructed machinists to bore properly sized holes with smooth edges into the plates for the purpose of passing cable bundles through with minimal risk of kinking during installation. Additionally, the design called for the plates to be mounted to the thermosyphon at different angles, creating a spiral effect that maximized the amount of shielding within the crossarm. This design also allowed the cable bundles to “wrap” around the thermosyphon, increasing their stability.

The assembly team test fit the new baffle plates with the fabricated cable bundles on a prototype thermosyphon before shipping them to SURF for installation. See Figures A.11a and A.11b for a comparison of the two baffle plate designs.

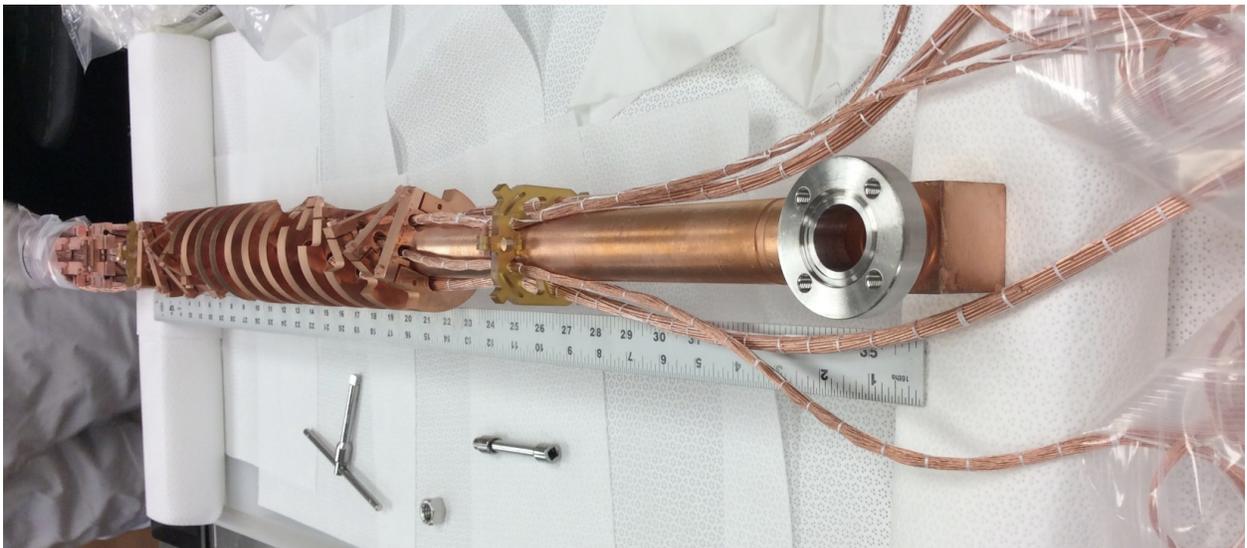
### **Section A.5: Installation at SURF**

A team of experts at SURF deconstructed the polyethylene shield around Module 2 starting in November 2019. Once completed, they removed the cryostat from the DEMONSTRATOR’s lead and copper shield and inserted it into the module glovebox. The team disassembled strings into detector units, and then disassembled those units to allow for the swapping of LMFE boards. They then commenced work on installing new LMFE’s on detectors in December. At the same time, they removed five enriched P-Point Contact (PPC) detectors that were immediately shipped to Italy for the LEGEND PGT.

The team installed new cable bundles and connectors along the Module 2 thermosyphon in January. First, the team removed the thermosyphon and old cables and connectors from the module. Then, they transferred five crossarm-side Axon’ connectors at a time from a 3D-printed Axon’ connector holder (referred to as the “diving board”, see Figure A.12) to a copper structure of the same design. The team then lined HV and signal cable bundles against the thermosyphon to install the baffle plates around them. They installed the baffle plates with 0.75” spacing between plates, and rotated each plate such that every 4th plate was 90 degrees rotated relative to the first. The team installed 16 out of the 17 fabricated plates, as there was insufficient clearance within the crossarm to install the last plate. The team positioned the HV cable bundle



(a) The old baffle plate arrangement on the thermosyphon. 12 baffle plates, split into 4 groups of 3 at orthogonal angles. Cables here are not bundled.



(b) The new baffle plate arrangement on the thermosyphon. 16 baffle plates, each plate at a staggered angle from the previous plate. Cables here are bundled with plain dental floss per NASA specifications.

Figure A.11: Old and new baffle plate configurations.

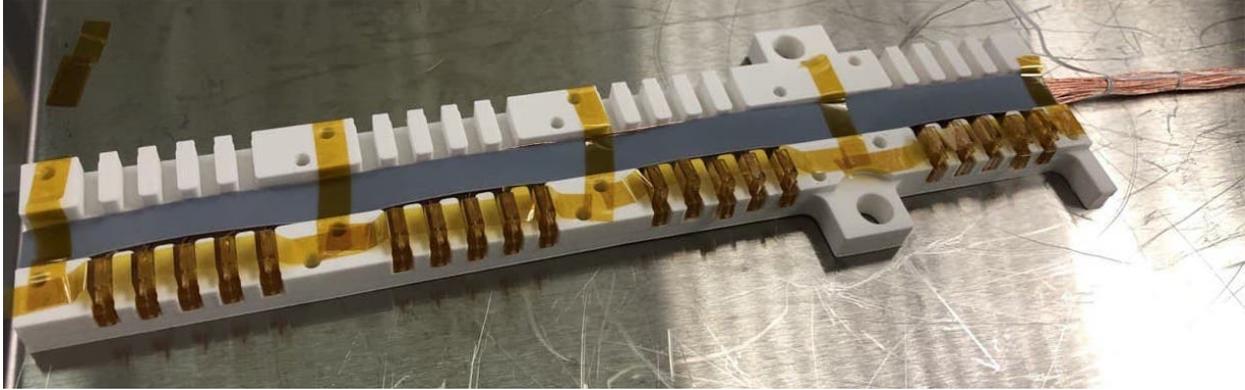


Figure A.12: The 3D-printed Axon’ connector holder, known colloquially as the “diving board” due to its position on the cryostat-end of the thermosyphon (see Figure A.1). The fabrication team at UNC securely positioned the nano twist-pin connectors on the board with kapton tape for the shipment to SURF. During installation, a team of experts carefully removed the tape and transferred five Axon’ connectors at a time to the identical position on a copper structure of the same design.

along the thermosyphon to allow for 24” of cable to be left on the cryostat side. They then re-inserted the thermosyphon into the crossarm and affixed the copper diving board loaded with Axon’ connectors to the cryostat end of the thermosyphon, just above the coldplate. The team then installed the signal 50-pin D-sub connectors and feedthrough flange HV connectors, with the 50-pin D-sub connectors requiring significant insertion force to mate them with the flange. The team performed continuity testing on the cables, with spare cables at the ready if any cables failed. All cables passed continuity testing and the team clipped any spare HV cables of their connectors. The team closed and tightened the feedthrough flange and other flanges of the high-vacuum stainless steel cube and readied the thermosyphon for operation. The team completed LMFE installation and string assembly work, including the installation of HV forks to each detector unit, by February.

From February through July, the Module 2 cryostat remained in the glovebox with detectors cooled in order to test the quality of the installation. In late July the team warmed the cryostat for the arrival of four new ICPC detectors from Oak Ridge National Laboratory (ORNL). They installed these detectors using newly machined electroformed copper structural parts to fit them in their corresponding strings. These detectors replaced the five PPC detectors that the team removed in December.

The team rearranged detectors and strings within the cryostat in order to characterise the effect of “crossarm shine” in the module. Crossarm shine is defined here as radiation from components outside of the cryostat that has a direct line of sight to detectors through the crossarm (see Figure A.1). Figure A.13 maps the position number of each string relative to the crossarm and Table A.5 lists the change in position for each string and detector before and after the upgrade. Detectors in the first row of positions 4 and 5 tend to face the highest exposure to crossarm shine due to their position directly across from the crossarm opening.

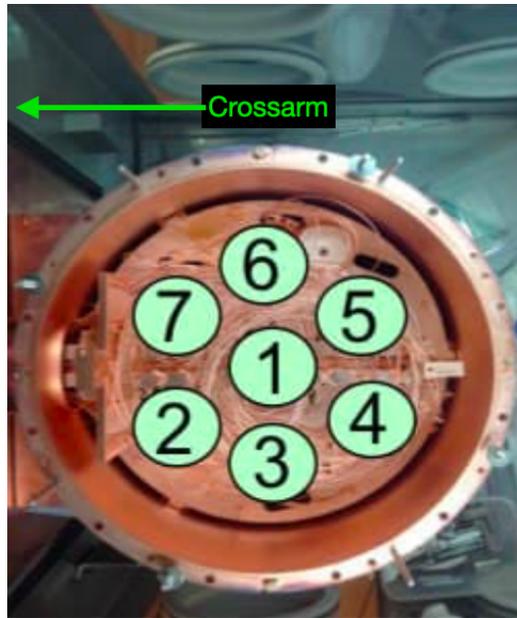


Figure A.13: Configuration of the string position numbers in a cryostat. This is a top-down view of the cryostat with the crossarm running to it from the left.

Therefore, if there is a measurable crossarm shine effect, we expect detectors that moved from positions 4 and 5 to positions 7 and 3 respectively to show a decreased rate. Conversely, we expect the two detectors that moved from position 7 to positions 4 and 5 to show an increased rate. In particular, the detector that moved from the bottom of position 7 to the top of position 4 should exhibit a noticeable rate change. Additionally, all detectors that changed string position (with the exception of the detector that moved from position 7 to 5) have switched from one side of the crossarm to the other, thus potentially exposing rate differences between detectors on the right and left sides of the crossarm opening. Following the installation of strings at their new positions in late July, the team cooled the cryostat in the glovebox to test the quality of the ICPC installation. They then warmed the cryostat, placed it back in the copper and lead shield, and cooled it again in late August. The team re-installed the polyethylene shield and data collection re-commenced in September.

## Section A.6: Results and Impact

As stated in section A.1, Module 2 had 11 inoperable detectors prior to the upgrade. With 29 total detectors, Module 2 had an operational efficiency of 62%. The primary goal of the upgrade was to increase this operational efficiency to greater than 90%. At the conclusion of the upgrade in July, 27 out of 27

Re-arrangement of M2 Detectors and Strings During the Upgrade		
String Number (Det Number)	Position Before the Upgrade	Position After the Upgrade
String 1	Position 1	Position 1
String 2	Position 2	Position 2
String 3	Position 3	Shipped to Italy.
String 4	Position 4	Position 7
String 5	Position 5	Position 3
String 6	Position 6	Position 6 (D4 moved to D3)
String 7 (D1)	Position 7 (D1)	Position 5 (D1)
String 7 (D2,D3)	Position 7	Shipped to Italy.
String 7 (D4)	Position 7 (D4)	Position 4 (D1)

Table A.5: Table of changes made to detector and string positions during the upgrade. If a detector is not specified, then the change applies to the entire string with detector position along the vertical axis unchanged. The four ICPC detectors were installed to the strings in positions 4 and 5, with two detectors at the bottom of each.

detectors<sup>3</sup> operated in Module 2 for an operational efficiency of 100%.

The success of this upgrade has a number of implications. First, the DEMONSTRATOR is on track for a last run of data collection at a higher rate of exposure. Second, the improved shielding down the crossarm from the new baffle plates as well as the rearrangement of detectors and strings in Module 2 will assist in ongoing background studies. These studies are not only important for the MAJORANA DEMONSTRATOR but also the LEGEND experiment as many of the same materials will be used in LEGEND-200. Third, the installation of new enriched ICPC germanium detectors in the DEMONSTRATOR allows for a unique opportunity to characterize these detectors in a vacuum environment with low background prior to their deployment in LEGEND. Finally, the 100% operational status of Module 2 demonstrates the capability of the cables and connectors that are used in the MAJORANA DEMONSTRATOR, given the fabrication and installation techniques that were followed. The nano twist-pin connectors in particular demonstrate an improvement over the previous signal connector design and may be used in future experiments. As many of the same techniques will be used during construction of LEGEND-200, there is an expectation that the operational efficiency of detectors in the initial deployment will be high.

---

<sup>3</sup>The upgrade reduced the total number of detectors in Module 2 from 29 to 27. 4 ICPC detectors replaced 5 PPC detectors, and an additional non-operations PPC detector was removed to be reprocessed by the manufacturer.

## APPENDIX B: COMPONENT GROUPS

Table B.1: List of parts and their corresponding component groups. Note that if the component group begins with M1 or M2, it is assumed that the list of parts are actually pairs of parts - one for each module.

Component Group	Parts	Source Type	Decay Chains	Configs
M1DUStringCopper	String Adapter Plate Adapter Plate Bolt Adapter Plate Nut Split Nut Tie Rod HV Ring Hollow Hex Rod Crystal Mounting Plate			
M2DUStringCopper	Contact Pin Spring Front End Mount LMFE Cover Plate Spring Nut HV Fork Nut HV Fork Inner Cable Guide Cable Guide Insert	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$	DS3-6

Component Group	Parts	Source Type	Decay Chains	Configs
M1DUPTFE	HV Nut	bulk surface	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$	DS3-6
M2DUPTFE	Crystal Insulator			
	Center Bushing			
	Flex Insulator			
M1LMFEs	LMFE Substrate	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$ $^{60}\text{Co}$	DS3-6
M2LMFEs	LMFE Traces			
M1Thermosyphon AndShieldVespel	Thick Support Female Half	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$	DS3-6
	Thick Support Male Half			
	Thick Support Clip			
	Thick Support Short Clip			
	Thin Support Female Half			
M2Thermosyphon AndShieldVespel	Thin Support Male Half			
	Thin Support Clip			
	Thin Support Short Clip			
	Thermal Shield Support			
	Centering Pin			

Component Group	Parts	Source Type	Decay Chains	Configs
M1CryostatCopperNear	Cryostat Top Lid Cryostat Hoop Cryostat Bottom Lid Center Rail Lid Rail Rail Assembly Bolts Rail Assembly Nuts Cold Plate 5/16-24x2.25 Bolts Takeups Clamp Plate Cable Clamp 8.32x0.25 Screws	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$	DS3-6
M2CryostatCopperNear	Cold Plate Interface Tube Adapter Thermal Shunt Clips Connector Clamp Thermal Shield Can Thermal Shield Annulus Thermal Shield Wedge			

Component Group	Parts	Source Type	Decay Chains	Configs
M2CryostatCopperNear (cont.)	8.32x0.625 Screws Top Shield Baffles Cable Guard Top Shield Near Support Top Shield Far Support 8.32x0.5 Screws	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$	DS3-6
M1CryostatCopperNear WeldedCopper	Cryostat Hoop Cold Plate Interface	bulk	$^{60}\text{Co}$	DS3-6
M2CryostatCopperNear WeldedCopper	Tube Adapter			
M1CryostatCopperFar	Crossarm Tube Thermosyphon Tube Corner Piece Thermosyphon Flange	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$ $^{60}\text{Co}$	DS0 DS3-6
M2CryostatCopperFar	Flange Shield Cube Shield Plates Shield Blocks			DS3-6

Component Group	Parts	Source Type	Decay Chains	Configs
M1Connectors	Connector Housing	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$	DS3-6
	Connector Strain Relief			
M2Connectors	Connector Pins			
	Connector Solder			
M1Bellows	Bellows	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ (directed $\gamma$ 's) $^{40}\text{K}$	DS0,3-6
M2Bellows				DS3-6
M1Seals	Top Seal	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$	DS3-6
M2Seals	Bottom Seal			
M1StringCables	String Signal Cables	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$ $^{60}\text{Co}$	DS3-6
M2StringCables	String HV Cables			
M1CrossarmAndCPCables	Coldplate Signal Cables	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$ $^{60}\text{Co}$	DS0
	Coldplate HV Cables			DS3-6
M2CrossarmAndCPCables	Crossarm Signal Cables			DS3-6
	Crossarm HV Cables			
M1EnrGe	Enr Ge PPC Detectors	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{60}\text{Co}$ $^{68}\text{Ge}$ $^{57}\text{Co}$ $2\nu\beta\beta$	DS3-6
M2EnrGe				
M1NatGe	Nat Ge BeGe Detectors	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{60}\text{Co}$ $^{68}\text{Ge}$ $^{57}\text{Co}$ $2\nu\beta\beta$	DS3-6
M2NatGe				

Component Group	Parts	Source Type	Decay Chains	Configs
M1CalSource	Calibration Assembly Track	line- source	$^{228}\text{Th}$	DS3-6
M2CalSource				
N2	Nitrogen Gas Purge Shield	bulk, surface	$^{222}\text{Rn}$	DS0-7
InnerCuShield	Inner Copper Shield	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$	DS1-7
OuterCuShield	Outer Copper Shield	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$ $^{60}\text{Co}$ $^{54}\text{Mn}$	DS0-7
RadShieldPb	<sup>1</sup> Lead Shield	bulk	$^{238}\text{U}$ $^{232}\text{Th}$ $^{40}\text{K}$ Bremstrahlung	DS0-7

<sup>1</sup>PDFs for U,Th,K are sometimes referred to as RadShieldAssembly.001\_RadShieldPb. PDFs of effects from bremstrahlung are sometimes referred to as RadShieldAssembly.001\_RadShieldPb.001



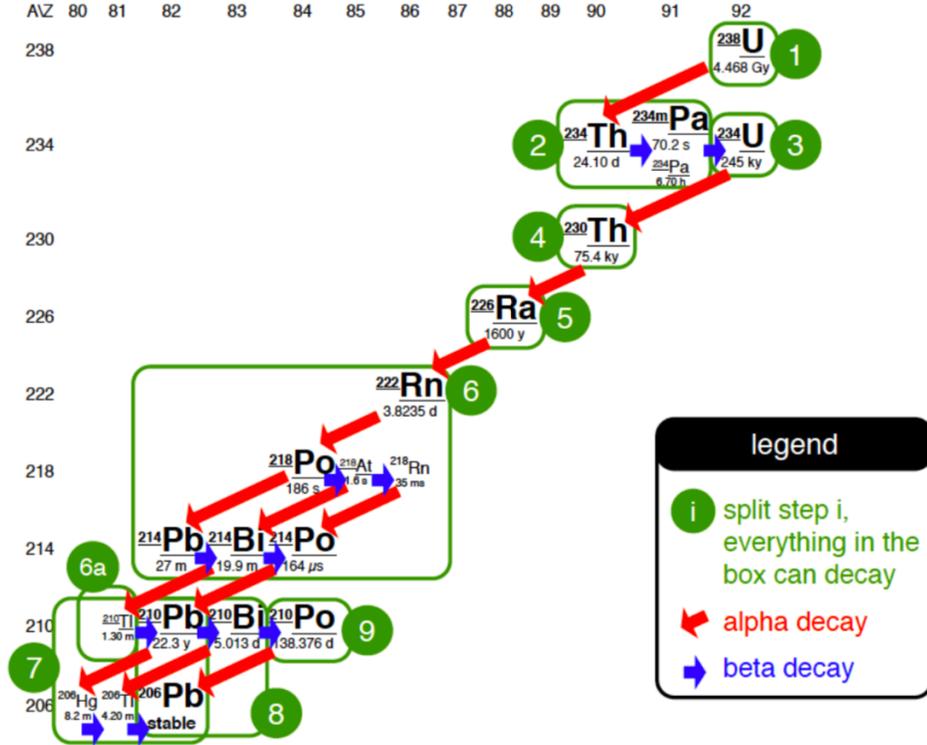


Figure C.2: Ten segments of the  $^{238}\text{U}$  chain. [63]

### Section C.2: Prominent $\gamma$ -decay Peak Energies

Decay Chain	Isotope	Energy (keV)	Intensity (%)
$^{232}\text{Th}$	$^{228}\text{Ac}$	911.2	26.6
$^{232}\text{Th}$	$^{228}\text{Ac}$	969.0	16.2
$^{232}\text{Th}$	$^{212}\text{Pb}$	238.6	43.6
$^{232}\text{Th}$	$^{208}\text{Tl}$	583.2	84.5
$^{232}\text{Th}$	$^{208}\text{Tl}$	2614.5	99.2
$^{238}\text{U}$	$^{214}\text{Pb}$	351.9	37.1
$^{238}\text{U}$	$^{214}\text{Bi}$	609.3	46.1
$^{238}\text{U}$	$^{214}\text{Bi}$	1120.3	15.0
$^{238}\text{U}$	$^{214}\text{Bi}$	1764.5	15.9
$^{238}\text{U}$	$^{214}\text{Bi}$	2204.1	5.0
$^{60}\text{Co}$	$^{60}\text{Co}$	1173.2	99.8
$^{60}\text{Co}$	$^{60}\text{Co}$	1332.5	100.0
$^{40}\text{K}$	$^{40}\text{K}$	1460.8	100.0

Table C.1: Prominent peaks from  $\gamma$ -decays observed by the DEMONSTRATOR. Intensity refers to the intra-isotope intensity.

## REFERENCES

- [1] E. Fermi. “Versuch einer Theorie der  $\beta$ -Strahlen”. In: *Zeits. f. Physik* 88.3-4 (1934), pp. 161–177. DOI: <https://doi.org/10.1007/BF01351864>.
- [2] M. Goeppert-Meyer. “Double beta-disintegration”. In: *Phys. Rev.* 48 (1935), pp. 512–516. DOI: <https://doi.org/10.1103/PhysRev.48.512>.
- [3] L. Cardani. “Neutrinoless Double Beta Decay Overview”. In: *arXiv* (2018). URL: <https://arxiv.org/pdf/1810.12828.pdf>.
- [4] W. H. Furry. “On transition probabilities in double beta-disintegration”. In: *Phys. Rev.* 56 (1939), pp. 1184–1193. DOI: <https://doi.org/10.1103/PhysRev.56.1184>.
- [5] S. Dell’Oro et al. “Neutrinoless Double Beta Decay: 2015 Review”. In: *Adv. High Energy Phys.* 2016 (2016). DOI: <http://dx.doi.org/10.1155/2016/2162659>.
- [6] S. Abe et al. “Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen”. In: *Phys. Rev. Lett.* 130 (5 Jan. 2023), p. 051801. DOI: 10.1103/PhysRevLett.130.051801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.051801>.
- [7] Matteo Agostini et al. “Toward the discovery of matter creation with neutrinoless  $\beta\beta$  decay”. In: *Rev. Mod. Phys.* 95 (2 May 2023), p. 025002. DOI: 10.1103/RevModPhys.95.025002. URL: <https://link.aps.org/doi/10.1103/RevModPhys.95.025002>.
- [8] B. Pontecorvo. “Inverse beta processes and nonconservation of lepton charge”. In: *Soviet Physics JETP* 34 (1958), p. 247.
- [9] Z. Maki, M. Nakagawa, and S. Sakata. “Remarks on the Unified Model of Elementary Particles”. In: *Progress of Theoretical Physics* 28 (5 1962), pp. 870–880. DOI: <https://doi.org/10.1143/PTP.28.870>.
- [10] A. Belley et al. “Ab Initio Neutrinoless Double-Beta Decay Matrix Elements for  $^{48}\text{Ca}$ ,  $^{76}\text{Ge}$ , and  $^{82}\text{Se}$ ”. In: *Phys. Rev. Lett.* 126 (4 Jan. 2021), p. 042502. DOI: 10.1103/PhysRevLett.126.042502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.126.042502>.
- [11] The EXO-200 collaboration. “Search for Majorana neutrinos with the first two years of EXO-200 data”. In: *Nature* 510 (2014), pp. 229–234. DOI: <https://doi.org/10.1038/nature13432>.
- [12] The CUORE Collaboration. “Search for Majorana neutrinos exploiting millikelvin cryogenics with CUORE”. In: *Nature* 604 (2022), pp. 53–58. DOI: <https://doi.org/10.1038/s41586-022-04497-4>.

- [13] The SNO+ collaboration et al. “The SNO+ experiment”. In: *JINST* 16.08 (Aug. 2021), P08059. DOI: 10.1088/1748-0221/16/08/P08059. URL: <https://dx.doi.org/10.1088/1748-0221/16/08/P08059>.
- [14] M. Agostini et al. “Final Results of GERDA on the Search for Neutrinoless Double- $\beta$  Decay”. In: *Phys. Rev. Lett.* 125 (25 Dec. 2020), p. 252502. DOI: 10.1103/PhysRevLett.125.252502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.252502>.
- [15] I.J. Arnquist et al. “Final Result of the MAJORANA DEMONSTRATOR’s Search for Neutrinoless Double- $\beta$  Decay in  $^{76}\text{Ge}$ ”. In: *Phys. Rev. Lett.* 130 (6 Feb. 2023), p. 062501. DOI: 10.1103/PhysRevLett.130.062501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.062501>.
- [16] N. Abgrall et al. “LEGEND-1000 Preconceptual Design Report”. In: *arXiv e-prints* (July 2021), arXiv:2107.11462. DOI: 10.48550/arXiv.2107.11462. arXiv: 2107.11462.
- [17] N. Abgrall, E. Aguayo, and F.T. Avignone III *et. al.* “The MAJORANA DEMONSTRATOR Neutrinoless Double-Beta Decay Experiment”. In: *Adv. High Energy Phys.* 2014 (2014). DOI: <http://dx.doi.org/10.1155/2014/365432>.
- [18] N. Abgrall et al. “New Limits on Bosonic Dark Matter, Solar Axions, Pauli Exclusion Principle Violation, and Electron Decay from the Majorana Demonstrator”. In: *Phys. Rev. Lett.* 118 (16 Apr. 2017), p. 161801. DOI: 10.1103/PhysRevLett.118.161801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.161801>.
- [19] I. J. Arnquist et al. *Exotic dark matter search with the Majorana Demonstrator*. 2022. arXiv: 2206.10638 [hep-ex].
- [20] MAJORANA Collaboration et al. *Search for charge nonconservation and Pauli exclusion principle violation with the Majorana Demonstrator*. 2023. arXiv: 2203.02033 [nucl-ex].
- [21] I. J. Arnquist et al. “Search for Solar Axions via Axion-Photon Coupling with the Majorana Demonstrator”. In: *Phys. Rev. Lett.* 129 (8 Aug. 2022), p. 081803. DOI: 10.1103/PhysRevLett.129.081803. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.081803>.
- [22] S. I. Alvis et al. “First Limit on the Direct Detection of Lightly Ionizing Particles for Electric Charge as Low as  $e/1000$  with the Majorana Demonstrator”. In: *Phys. Rev. Lett.* 120 (21 May 2018), p. 211804. DOI: 10.1103/PhysRevLett.120.211804. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.211804>.
- [23] S. I. Alvis et al. “Search for trinucleon decay in the Majorana Demonstrator”. In: *Phys. Rev. D* 99 (7 Apr. 2019), p. 072004. DOI: 10.1103/PhysRevD.99.072004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.99.072004>.

- [24] I. J. Arnquist et al. “Search for Spontaneous Radiation from Wave Function Collapse in the Majorana Demonstrator”. In: *Phys. Rev. Lett.* 129 (8 Aug. 2022), p. 080401. DOI: 10.1103/PhysRevLett.129.080401. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.080401>. Corrected in “Erratum: Search for Spontaneous Radiation from Wave Function Collapse in the Majorana Demonstrator [Phys. Rev. Lett. 129, 080401 (2022)]”. In: *Phys. Rev. Lett.* 130 (23 June 2023), p. 239902. DOI: 10.1103/PhysRevLett.130.239902. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.239902>.
- [25] C. E. Aalseth, N. Abgrall, and E. Aguayo *et. al.* “Search for Neutrinoless Double- $\beta$  Decay in  $^{76}\text{Ge}$  with the MAJORANA DEMONSTRATOR”. In: *Phys. Rev.* 120 (2018). DOI: <https://doi.org/10.1103/PhysRevLett.120.132502>.
- [26] URL: <https://www.teflon.com/en/products/resins/ptfe-granular/nxt-70-75-85>.
- [27] P S Barbeau, J I Collar, and O Tench. “Large-mass ultralow noise germanium detectors: performance and applications in neutrino and astroparticle physics”. In: *JCAP* 2007.09 (2007), p. 009. DOI: 10.1088/1475-7516/2007/09/009. URL: <https://dx.doi.org/10.1088/1475-7516/2007/09/009>.
- [28] P.N. Luke et al. “Low capacitance large volume shaped-field germanium detector”. In: *IEEE Transactions on Nuclear Science* 36.1 (1989), pp. 926–930. DOI: 10.1109/23.34577.
- [29] G.F. Knoll. *Radiation Detection and Measurement*. 4th. John Wiley & Sons, 2010. ISBN: 9780470131480.
- [30] W. Shockley. “Currents to Conductors Induced by a Moving Point Charge”. In: *J. Appl. Phys.* 9.10 (1938), pp. 635–636. DOI: <https://doi.org/10.1063/1.1710367>.
- [31] S. Ramo. “Currents Induced by Electron Motion”. In: *Proceedings of the IRE* 27.9 (1939), pp. 584–585. DOI: <https://doi.org/10.1109/JRPR0C.1939.228757>.
- [32] B. Shanks. “High precision modeling of germanium detector waveforms using Bayesian machine learning”. In: *Dissertation* (2017).
- [33] N. Abgrall et al. “The Majorana Demonstrator readout electronics system”. In: *Journal of Instrumentation* 17.05 (May 2022), T05003. DOI: 10.1088/1748-0221/17/05/T05003. URL: <https://dx.doi.org/10.1088/1748-0221/17/05/T05003>.
- [34] John Anderson et al. “Data Acquisition and Trigger System of the Gamma Ray Energy Tracking In-Beam Nuclear Array (GRETINA)”. In: *IEEE Transactions on Nuclear Science* 56.1 (2009), pp. 258–265. DOI: 10.1109/TNS.2008.2009444.
- [35] M.A. Howe *et. al.* “Sudbury neutrino observatory neutral current detector acquisition software overview”. In: *IEEE Trans Nucl Sci* 51 (2004), pp. 878–883. URL: <https://doi.org/10.1109/TNS.2004.829527>.

- [36] N. Abgrall, I.J. Arnquist, and F.T. Avignone III *et. al.* “The MAJORANA DEMONSTRATOR calibration system”. In: *Nucl. Instrum. Methods Phys. Res. A* 872 (2017), pp. 16–22. DOI: <https://doi.org/10.1016/j.nima.2017.08.005>.
- [37] LEGEND Collaboration et al. *LEGEND-1000 Preconceptual Design Report*. 2021. arXiv: 2107.11462 [physics.ins-det].
- [38] N. Abgrall, I.J. Arnquist, and F.T. Avignone III *et. al.* “The processing of enriched germanium for the MAJORANA DEMONSTRATOR and R&D for a next generation double-beta decay experiment”. In: *Nucl. Instrum. Methods Phys. Res. A* 877 (2018), pp. 314–322. DOI: <https://doi.org/10.1016/j.nima.2017.09.036>.
- [39] URL: <https://pubs.usgs.gov/of/2004/1050/thorium.htm>.
- [40] URL: <https://pubs.usgs.gov/of/2004/1050/uranium.htm>.
- [41] N. Abgrall, I.J. Arnquist, and F.T. Avignone III *et. al.* “The MAJORANA DEMONSTRATOR radioassay program”. In: *Nucl. Instrum. Methods Phys. Res. A* 828 (2016), pp. 22–36. DOI: <https://doi.org/10.1016/j.nima.2016.04.070>.
- [42] S.I. Alvis, I.J. Arnquist, and F.T. Avignone III *et. al.* “Multi-site event discrimination for the MAJORANA DEMONSTRATOR”. In: *Phys. Rev. C* 99 (6 June 2019), p. 065501. DOI: 10.1103/PhysRevC.99.065501. URL: <https://link.aps.org/doi/10.1103/PhysRevC.99.065501>.
- [43] A. Hostiuc and I. Guinn. “Late Charge Tagging”. In: *Internal Document* (2021).
- [44] I.J. Arnquist, F.T. Avignone III, and A.S. Barabash *et. al.* “ $\alpha$ -event Characterization and Rejection in Point-Contact HPGe Detectors”. In: *Eur. Phys. J. C* (2022). DOI: <https://doi.org/10.1140/epjc/s10052-022-10161-y>.
- [45] I. J. Arnquist et al. “Search for double- $\beta$  decay of  $^{76}\text{Ge}$  to excited states of  $^{76}\text{Se}$  with the majorana demonstrator”. In: *Phys. Rev. C* 103 (1 Jan. 2021), p. 015501. DOI: 10.1103/PhysRevC.103.015501. URL: <https://link.aps.org/doi/10.1103/PhysRevC.103.015501>.
- [46] I. J. Arnquist et al. “An assay-based background projection for the MAJORANA DEMONSTRATOR using Monte Carlo Uncertainty Propagation”. In: *Internal document, to be published*. (2023).
- [47] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8). URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688>.

- [48] Melissa Boswell et al. “MaGe-a Geant4-Based Monte Carlo Application Framework for Low-Background Germanium Experiments”. In: *IEEE Transactions on Nuclear Science* 58.3 (2011), pp. 1212–1220. DOI: 10.1109/TNS.2011.2144619.
- [49] M. Laubenstein. personal communication. Aug. 23, 2023.
- [50] C. Bailer-Jones. *Practical Bayesian Inference: A Primer for Physical Scientists*. Cambridge University Press, 2017. DOI: 10.1017/9781108123891.
- [51] R. McElreath. *Statistical rethinking: a Bayesian course with examples in R and Stan*. 2nd. Chapman and Hall/CRC, 2020. ISBN: 9780367139919.
- [52] J. Salvatier, T.V. Wiecki, and C. Fonnesbeck *et. al.* “Probabilistic programming in Python using PyMC3”. In: *PeerJ Comput. Sci.* 55 (2016). URL: <https://doi.org/10.7717/peerj-cs.55>.
- [53] Pauli Virtanen et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.
- [54] Art B. Owen. *Monte Carlo theory, methods and examples*. <https://artowen.su.domains/mc/>, 2013.
- [55] A. Reine. “An Improved Background Model and Two-Neutrino Double-Beta Decay Measurement for the MAJORANA DEMONSTRATOR”. In: *Dissertation* (2023).
- [56] N. Abgrall, S.I. Alvis, and I.J. Arnquist *et. al.* “Design improvements to cables and connectors in the MAJORANA DEMONSTRATOR”. In: *arXiv* (2017). URL: <https://arxiv.org/pdf/1712.03459.pdf>.
- [57] N. Abgrall, I.J. Arnquist, and F.T. Avignone III *et. al.* “High Voltage testing for the MAJORANA DEMONSTRATOR ”. In: *Nucl. Instrum. Methods Phys. Res. A* 823 (2016), pp. 83–90. DOI: <https://doi.org/10.1016/j.nima.2016.04.006>.
- [58] *NASA-STD-8739.4 - WORKMANSHIP STANDARD FOR CRIMPING, INTERCONNECTING CABLES, HARNESSES, AND WIRING*. URL: <https://standards.nasa.gov/standard/nasa/nasa-std-8739.4>.
- [59] URL: [https://www.axon-cable.com/en/00\\_home/00\\_start/00/index.aspx](https://www.axon-cable.com/en/00_home/00_start/00/index.aspx).
- [60] URL: <https://www.glenair.com>.
- [61] URL: <https://www.schleuniger.com/en/>.

[62] URL: <https://www.loctite.com/en/channel-switch/usa.html>.

[63] T. Gilliss. “An As-Built Background Model for the MAJORANA DEMONSTRATOR”. In: *Dissertation* (2019).